Are Fuel Economy Standards Regressive?

Lucas W. Davis, Christopher R. Knittel

Abstract: Despite widespread agreement that a carbon tax would be more efficient, many countries use fuel economy standards to reduce transportation-related carbon dioxide emissions. We pair a simple model of the automakers' profit maximization problem with unusually rich nationally representative data on vehicle registrations to estimate the distributional impact of US fuel economy standards. The key insight from the model is that fuel economy standards impose a constraint on automakers that creates an implicit subsidy for fuel-efficient vehicles and an implicit tax for fuel-inefficient vehicles. Moreover, when these obligations are tradable, permit prices make it possible to quantify the exact magnitude of these implicit subsidies and taxes. We use the model to determine which vehicles are most subsidized and taxed, and we compare the pattern of ownership of these vehicles between high- and low-income census tracts. Finally, we compare these distributional impacts with existing estimates in the literature for a carbon tax.

GLOBAL OIL CONSUMPTION now exceeds 90 million barrels per day (EIA 2016) fueling the more than 1.2 billion vehicles in use worldwide (BP Energy 2015). Total carbon dioxide emissions from road transportation exceed 5 gigatons annually (IPCC 2015), approximately one-sixth of all anthropogenic emissions. Policy makers are increasingly turning their attention to this important sector, evaluating alternative ap-

Lucas W. Davis is Jeffrey A. Jacobs Distinguished Professor in Business and Technology, Haas School of Business, University of California, Berkeley; Energy Institute at Haas; and National Bureau of Economic Research (lwdavis@berkeley.edu). Christopher R. Knittel is George P. Shultz Professor of Applied Economics, Sloan School of Management, Massachusetts Institute of Technology; director, Center for Energy and Environmental Policy Research, Massachusetts Institute of Technology; and at the National Bureau of Economic Research (knittel@mit.edu).

proaches for reducing carbon dioxide emissions. Economists agree that the most cost-effective approach would be a carbon tax, or equivalently, taxes on gasoline and diesel. Carbon dioxide emissions are proportional to fuels consumption, so either approach would be first-best for reducing carbon dioxide emissions from driving.

Despite widespread agreement that a carbon tax would be more efficient, many countries use fuel economy standards to reduce transportation-related carbon dioxide emissions. The United States and Japan have long histories with fuel economy standards, and similar policies have also been recently implemented by the European Union, China, India, and elsewhere (Anderson and Sallee 2016). The exact format differs between countries, but many programs follow the US Corporate Average Fuel Economy (CAFE) standards in requiring automakers to meet a minimum sales-weighted average for their vehicle fleets.

It can be easier politically to introduce fuel economy standards than taxes, but the two are not equivalent. Probably the single biggest limitation of fuel economy standards is that they do not achieve the efficient level of vehicle usage; to efficiently reduce gasoline consumption you need people to buy more fuel-efficient cars and to drive them less. But economists have pointed out other disadvantages as well. For example, Jacobsen and van Benthem (2015) show that fuel economy standards reduce the incentive for drivers to retire old vehicles, leading fuel-inefficient vehicles to stay on the road longer. Overall, studies have found that fuel economy standards are three to six times more costly than a carbon tax (Austin and Dinan 2005; Jacobsen 2013).

In this paper we ask a different but related question. Are fuel economy standards regressive? How the burden of standards is borne between high- and low- income households is one of the factors that must be considered when comparing standards to alternative policies for reducing carbon dioxide emissions. To answer this question, we pair a simple model of the automakers' profit maximization problem with unusually rich nationally representative data on vehicle registrations. The key insight from the model is that fuel economy standards impose a constraint on automakers that creates an implicit subsidy for fuel-efficient vehicles and an implicit tax for fuel-inefficient vehicles. Moreover, when these obligations are tradable, as they are under US CAFE standards, the permit prices make it possible to quantify the exact magnitude of these implicit subsidies and taxes.

When we consider new vehicles only, we find that US fuel economy standards are mildly progressive. High-income households bear more of the cost as a fraction of income than low-income households. However, this largely reflects that high-income households buy more new vehicles. When we expand the analysis to include used vehicles, standards become mildly regressive.

We then compare these distributional impacts with existing estimates in the literature on the distributional impact of a carbon tax. In general, this literature has found that the regressivity or progressivity of a carbon tax strongly depends on what is done

with the collected revenue. Comparing our estimates with the previous literature, we conclude that CAFE standards are more progressive than a carbon tax that does not recycle the revenue, but more regressive than a carbon tax that recycles the revenue through uniform transfers. Put simply, it is easy to design a carbon tax that is more progressive than fuel economy standards. Thus, we conclude that it is difficult to argue for fuel economy standards on distributional grounds.

This paper fills an important gap in the literature. Previous studies of fuel economy standards have focused almost exclusively on efficiency and overall cost-effectiveness, but the distributional impacts of fuel economy standards have received little attention. An important exception is Jacobsen (2013), which studies the distributional impacts of CAFE using microdata from the 2001 National Household Travel Survey. The paper finds a similar pattern, with the estimated distributional impact flipping from progressive to regressive once used vehicles are incorporated. In the full analysis, Jacobsen (2013) finds that low-income households suffer proportional welfare losses three times as large as high-income households.

Before proceeding we want to be clear about several important limitations of our analysis. First, our maintained assumption throughout is that these impacts are borne entirely by vehicle buyers rather than automakers or retailers. This is a reasonable assumption in market segments that are highly competitive and consistent with at least one study of subsidy incidence in the US automobile sector (Sallee 2011), but it is a strong assumption that we are not able to verify empirically.

Second, our calculations implicitly assume that these taxes and subsidies do not change buyers' vehicle decisions. This is, of course, incorrect. Part of the purpose of fuel economy standards is to move buyers toward more fuel-efficient vehicles. We might see, for example, someone who owns a vehicle that is subsidized under CAFE and conclude they are a "winner" when in fact, without CAFE they would have purchased a different vehicle that because of CAFE is taxed. This buyer is not a winner at all and despite buying a subsidized vehicle might have suffered a significant welfare loss. We do not estimate a demand model so our analysis is silent on this substitution behavior and on the broader welfare impacts of CAFE.

Third, our approach for modeling the impact of CAFE on used vehicles is ad hoc and only a rough approximation. Fuel economy standards apply only to new vehicles but have significant indirect impacts on used vehicle prices. We model this using a strong simplifying assumption that we argue describes the general pattern of impacts but cannot capture all of the complicated cross-price effects.

Fourth, our analysis is short run in that we do not model the impact of fuel economy standards on innovation. Over a longer time horizon fuel economy standards can lead to the development of entirely new vehicle models that might have disproportionate impacts across income groups. There has been some recent work on fuel economy standards and innovation (Klier and Linn 2016), but we are not aware of any work examining the potential long-run distributional impacts.

1. BACKGROUND

The Corporate Average Fuel Economy (CAFE) program was introduced in the United States in 1975 with the objective of reducing gasoline consumption. Under CAFE, automakers are required to meet a minimum sales-weighted average fuel economy for their vehicle fleets. These requirements have been tightened several times, most recently with a significant revision to the program resulting in new program rules which took effect starting in 2012. For a complete regulatory history, see Anderson et al. (2011), Knittel (2012), and Leard and McConnell (2017).

1.1. Footprint-Based Standards

As has always been the case with CAFE standards, automakers are required to meet a minimum sales-weighted average fuel economy for their vehicle fleet. Since 2012, however, this target now depends on the footprint of vehicles in the fleet. Calculated as the product of a vehicle's wheelbase (i.e., length) and track width (i.e., width), the footprint is a simple measure of the overall size of the vehicle.

Each vehicle has a different emissions target based on its footprint and on whether it is a car or truck. Figure 1 shows the emissions targets (in grams of carbon dioxide per mile) for cars and trucks produced between 2012 and 2016. The rules establish increasingly strict requirements on fuel economy for each vehicle model year. Larger vehicles receive larger emissions targets and trucks receive preferential treatment in the form of higher emissions targets for any given footprint and model year.

Just because a vehicle is small does not ensure that it meets the emissions target. For example, the Mini Cooper, with a footprint of 39 square feet in 2012, received an emissions target of 244 grams of carbon dioxide per mile. Actual emissions are 296 grams per mile, significantly above the emissions target. Thus, even though this car is one of the smallest on the road weighing only 2,500 pounds and with 115 horsepower, it is less fuel efficient than its footprint-based target. Thus if BMW wants to sell more Mini Coopers, it also needs to sell more of some other vehicle that is below its target or BMW needs to buy permits from some other automaker.
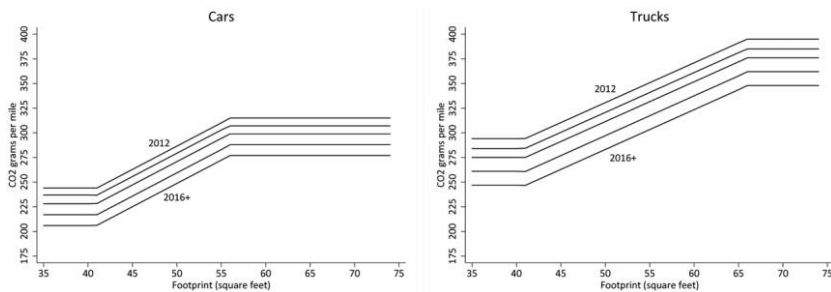

Figure 1. Emissions targets, 2012 to 2016

Herein lies the central problem with footprint-based targets. For a given vehicle footprint, the standards encourage automakers to make their vehicles as fuel efficient as possible. However, the footprint-based standards create no incentive for buyers to choose smaller vehicles. This may make sense from a political perspective in the United States because domestic automakers produce large numbers of SUVs, crossovers, and pickups, but it is an expensive, inefficient approach for reducing greenhouse gas emissions.[1]

Another significant distortion with the standards is the preferential treatment for trucks. For a given footprint, trucks have a less stringent carbon emissions standard than cars, so the standards encourage automakers to sell more trucks and fewer cars. The preferential treatment for trucks also encourages automakers to classify as many vehicles as possible as "trucks."[2] Under current rules "trucks" include not only pickup trucks but also SUVs, crossovers, and minivans. These are some of the largest and fastest growing segments in the US automobile market.[3]

1.2. Permit Trading Rules

Under CAFE, each vehicle sale generates a small surplus or deficit for the automaker depending on whether the vehicle is below or above its target emissions value. The total balance is then evaluated separately for each model year and each automaker. If an automaker is below the total emissions target for all vehicles sold, then it has a surplus for the year and receives permits. If instead an automaker is above the total emissions target, then it has a deficit and must buy permits.

Permits are denoted in tons of carbon dioxide. Emissions rates are converted into emissions levels using an assumption about how many miles each vehicle will be used during its lifetime, with trucks assumed to be driven more total miles over their lifetime. There are 8.887 kilograms of carbon dioxide per gallon of gasoline and 10.180 kilograms of carbon dioxide per gallon of diesel, so given this assumption about total miles driven, there is a simple mapping between vehicle fuel economy and total carbon dioxide emissions in tons.

---

1. See Ito and Sallee (2018) for a broader discussion of the economic costs of attribute-based regulation.

2. The classic example is a vehicle Chrysler manufactured called the PT Cruiser. In the early 2000s, Chrysler was making large profits on its Dodge Ram pickups and wanted to sell more but was running up against the CAFE constraint. Ingeniously, Chrysler responded by introducing the PT Cruiser, which looked like a car but was built on a "truck" platform, thus raising Chrysler's average fuel economy for trucks and allowing Chrysler to sell more fuel-inefficient pickups.

3. The biggest year ever for the US auto industry was 2015, with 17.5 million total vehicle sales nationwide, including large year-on-year increases for trucks, SUVs, and crossovers. See *Automotive News*, "U.S. Auto Sales Break Record in 2015," January 5, 2016.

Automakers can bank permits for up to 5 years and borrow permits for up to 3 years. This flexibility is intended to help automakers smooth over year-to-year fluctuations in demand driven by macroeconomic shocks, changes in gasoline prices, and other factors. The banking and borrowing also provides stability for the permit market, helping to avoid permit price spikes and crashes, and mitigating concerns about market power in permit markets.[4]

Permits may also be traded between automakers. Permit trading increases the efficiency of fuel economy standards. Just as with any cap-and-trade policy, trading equalizes marginal cost across agents, achieving the targeted aggregate level of emissions reductions at lowest total cost. These efficiencies are substantial with fuel economy standards because opportunities for improvements in fuel economy vary widely across automakers. For some automakers there is low-hanging fruit, for example, because they already have relative expertise in producing and marketing fuel-efficient vehicles whereas for other automakers it is much harder.

All automakers have an incentive to improve fuel economy, including those who are well above the fuel economy standard. This was not the case under the old CAFE rules that did not allow trading. For example, Toyota and Honda tend to sell relatively fuel-efficient vehicles so are perennially well above the minimum fuel economy standard. For automakers in this position under the old CAFE rules, it was as if the standard did not exist. There was no penalty, but also no incentive to make further improvements in fuel economy. In fact, these automakers had an incentive to make larger vehicles to pull market share away from other automakers who were constrained by CAFE. Now with permit trading any improvement in fuel economy generates CAFE credits and, thus, profit.

Automakers trade permits through bilateral trades. There is no central clearing house for permit trading, nor is there any system in place for making permit prices publicly available. Leard and McConnell (2017) nonetheless manage to infer permit prices based on information from two different sources: (i) a Department of Justice settlement with Hyundai and Kia resolving overstated fuel economy labels and (ii) Tesla Motors' SEC Filing Form 10-K from 2013 and 2014 reporting earnings from permit sales. These sources yield a permit price of between $35 and $40 per ton. We adopt these permit price values in the empirical results that follow, but it would be straightforward to incorporate updated permit prices as newer information becomes available. We note that because each vehicle faces the same shadow value changes to the shadow value will simply scale up and down the incidence at both the vehicle and consumer-group (e.g., consumers of particular incomes) levels.

Interestingly, these inferred permit prices are close to recent median estimates of the social cost of carbon (see, e.g., Interagency Working Group on Social Cost of Car-

---

4. See Borenstein et al. (2016) for a discussion of similar issues in cap-and-trade programs for carbon dioxide.

bon 2013). Thus, on the margin, automakers would appear to be facing an incentive similar to a, for example, $40 per ton tax on carbon dioxide. (Future carbon dioxide emissions are not discounted under CAFE, so it would actually be equivalent to a somewhat higher tax per ton.) This permit price creates a very different incentive than an equivalent carbon tax, however. Unlike a carbon tax, fuel economy standards do not encourage drivers to use their vehicles less intensively. A carbon tax more efficiently reduces gasoline consumption by encouraging drivers to buy more fuel-efficient cars and to drive less. Fuel economy standards can never address this second margin, nor can they, with footprint-based targets, encourage drivers to buy smaller vehicles.

1.3. Alternative Fuel Vehicles

This section reviews three features of CAFE that provide preferential treatment for alternative-fuel vehicles. First, electric vehicles (EVs), for compliance purposes, are assumed to emit no carbon dioxide. Recent empirical evidence shows that the actual carbon impact from EVs is quite varied depending on where and when the vehicle is charged, but that in many cases the carbon impact of EVs may actually exceed the carbon emissions from gasoline-powered vehicles (see, e.g., Holland et al. 2016). This treatment of EVs was meant as an explicit subsidy and not as an accurate description of the current carbon emissions from EVs.

Second, plug-in hybrids like the Chevrolet Volt and Toyota Prius plug-in also receive preferential treatment. These vehicles have both an electric drive train and internal combustion engine and thus can be operated using either electricity or gasoline. For CAFE compliance purposes, the gasoline component is treated normally but the electric component is assumed to be zero carbon. This is again not an accurate description of current emissions resulting from these vehicles. Even in parts of the country where electricity generation tends to be relatively low carbon, the marginal source of electricity generation is virtually always some form of fossil fuel generation.

Third, the feature that ends up being most important quantitatively during our sample period is the treatment of flexible-fuel (or "flex-fuel") vehicles which can run either on E85 (a blend of 85% ethanol and 15% gasoline) or on regular gasoline. Between 2012 and 2015 when calculating CAFE compliance, flex-fuel vehicles were assumed to be operated 50% using E85 and 50% with gasoline. Moreover, each gallon of E85 was assumed to have the carbon content of only 0.15 gallons of gasoline. That is, the ethanol component of E85 is assumed to be zero carbon.

These assumptions are probably not an accurate description of the carbon dioxide emissions from flex-fuel vehicles. Empirical evidence is not available on the fraction of flex-fuel vehicles that operate using E85, but the 50/50 assumption is an optimistic assumption given that many sales of flex-fuel vehicles occur in parts of the country where there is limited E85 availability (Anderson and Sallee 2011). The assumption about carbon content is also optimistic and hard to reconcile with a substantial scientific literature on the carbon emissions from ethanol. Quantifying the lifetime carbon

impacts of ethanol is challenging because of land use effects and other factors, but most studies find that, at best, ethanol is only marginally less carbon intensive than gasoline (Knittel 2012).

As we show later in the paper, these generous assumptions lead flex-fuel vehicles to be treated by CAFE as if they were extremely fuel efficient. Not surprisingly, automakers have responded by building large numbers of flex-fuel vehicles. For the 2014 model year, for example, there are more than 100 different models of flex-fuel vehicles for sale in the United States. Moreover, even though the preferential treatment for flex-fuel vehicles ended with model year 2015, this feature of CAFE was so lucrative that many automakers were able to generate large stores of surplus credits (EPA 2016). Under CAFE rules these credits can be banked until 2021, so these banked credits will allow automakers to produce lower-MPG vehicles for years to come.

2. CONCEPTUAL FRAMEWORK

In this section we write down the automakers' profit-maximization problem subject to the CAFE constraint. From this problem we derive a first-order condition that quantifies the implicit subsidy and tax for each vehicle which we can then take to the empirical analysis. We first consider the case of perfect competition. The degree to which perfect competition is a reasonable assumption varies across vehicle segments from highly competitive segments like compact sedans and midsize SUVs to less-competitive segments like luxury vehicles where some automakers can influence price considerably. We then relax the perfect competition assumption in the following subsection and show that the first-order condition is similar in both cases.

2.1. Perfect Competition

The automaker chooses quantities to maximize profits,

$$\max_{q_1, q_2 \ldots, q_J} \sum_{j=1}^{J} \left[q_j p_j - c_j(q_j) \right]. \tag{1}$$

Here q_j is total sales of vehicle model j. Revenues are the product of sales (q_j) and prices (p_j), summed over all vehicle models. Profits are total revenues minus total costs, where the cost of producing q_j units of vehicle model j is denoted $c_j(q_j)$. Here we allow production costs to vary between vehicle models but rule out complementarities between models.

The fuel economy standard can be expressed as follows,

$$\sum_{j=1}^{J} \left[\left(\text{emissions}_j - \text{target}_j \right) * VMT_j * q_j \right] \leq 0. \tag{2}$$

In this equation emissions$_j$ is carbon dioxide emissions in grams per mile for vehicle model j. Depending on its footprint, each vehicle model is assigned a target emissions

level, target$_j$, also measured in grams per mile. Thus the first part of equation (2) in parentheses reflects whether each vehicle model is above or below its target. These deviations are then weighted based on assumed lifetime miles traveled (VMT_j) and vehicle sales (q_j).

The automaker maximizes profits subject to the CAFE constraint. The Lagrangian can be written as follows,

$$\mathcal{L} = \sum_{j=1}^{J} \left[q_j p_j - c_j(q_j) \right] - \lambda \sum_{j=1}^{J} \left[\left(\text{emissions}_j - \text{target}_j \right) * VMT_j * q_j \right], \qquad (3)$$

where λ is the Lagrangian multiplier on the CAFE constraint. Differentiating with respect to q_j yields the following first-order condition,

$$p_j = c_j'(q_j) + \lambda \left[\left(\text{emissions}_j - \text{target}_j \right) * VMT_j \right]. \qquad (4)$$

In the first-order condition λ is the shadow value of the CAFE constraint. With permit trading this shadow value equals the permit price. The permit price is the relevant opportunity cost for all automakers, regardless of whether they have a surplus or a deficit. The first-order condition is unchanged by banking or borrowing of permits. For example, we could have included banked permits from previous years in the CAFE constraint but this is not a function of q_j and thus would not have entered the first-order condition.

The first-order condition has an intuitive interpretation. Consider first the case in which the permit price equals zero. In this case the shadow value λ is zero, and the automaker maximizes profit by increasing the quantity sold of each vehicle up until price equals marginal cost. This is the standard first-order condition for perfect competition. For nonzero permit prices, the automaker maximizes profit by adjusting quantities to reflect both marginal cost and the additional cost (or benefit) that accrues because of the standard. For vehicle models that emit more than their target emissions level there is an additional cost for each unit sold, so the optimal quantity is lower. For vehicle models that emit less than their target emissions level there is an additional benefit for each unit sold, so the optimal quantity is higher.

In short, the fuel economy standard creates a tax for fuel-inefficient vehicles and a subsidy for fuel-efficient vehicles.[5] In the empirical analyses that follow we use this insight to calculate the taxes and subsidies associated with all vehicle models sold in the United States. We focus on the period since 2012, which allows us to use permit prices to measure the shadow value of the constraint. In principle, however, it would be relatively straightforward to take parameter estimates from studies like Anderson

---

5. See Kwoka (1983), Helfand (1991), and Holland et al. (2009) for related discussions.

and Sallee (2011) and Jacobsen (2013) to infer the shadow value of λ and use this general approach to perform analogous calculations for the pre-2012 period.

2.2. Imperfect Competition

The problem for an oligopolist automaker is similar except they face a downward-sloping demand curve and in the first-order condition price is replaced with marginal revenue,

$$p_j + \sum_{i=1}^{I}\left[q_i\frac{\partial p_i}{\partial q_j}\right] = c'_j(q_j) + \lambda\left[\left(\text{emissions}_j - \text{target}_j\right) * VMT_j\right]. \qquad (5)$$

Here we have allowed cross elasticities to be nonzero, so changing the quantity sold of vehicle model j can affect sales not only of vehicle model j but also of all other vehicle models in the automaker's fleet. If we assume these cross elasticities are zero, the first-order condition simplifies further,

$$p_j + q_j\frac{\partial p_j}{\partial q_j} = c'_j(q_j) + \lambda\left[\left(\text{emissions}_j - \text{target}_j\right) * VMT_j\right]. \qquad (6)$$

In either case, the first-order condition takes on this same basic form with marginal revenue equal to marginal cost. Just as with the first-order condition with perfect competition, the CAFE constraint enters the first-order condition additively as an additional cost (or benefit) of selling one more unit of model j. Here we have assumed that automakers are choosing quantities; if we instead use Bertrand-Nash the first-order condition is somewhat different but the CAFE constraint enters similarly (Gillingham 2013; Jacobsen 2013).

Thus, overall, pricing behavior with imperfect competition is similar but not identical to pricing under perfect competition. Regardless of market structure, fuel economy standards lead automakers to price fuel-efficient vehicles lower than they would have been otherwise, and to price fuel-inefficient vehicles higher than would have been otherwise. In the sections that follow, we use this insight together with permit prices to quantify these implicit subsidies and taxes.

3. EMPIRICAL ANALYSIS

The key insight from the previous section is that fuel economy standards create an implicit subsidy for fuel-efficient vehicles and an implicit tax for fuel-inefficient vehicles. In this section, we pair this insight with rich microdata from the US automobile market to estimate the distributional impact of fuel economy standards. We proceed in several steps. First, we show how each vehicle model compares to its target emissions level. Second, we use the automakers' first-order condition to quantify the implicit tax (or subsidy) for each vehicle model. Third, we use national data on vehicle registrations by census tract to see how the average impact of CAFE standards varies between high- and low-income tracts.

3.1. Data

We use data from DataOne Software that describe all cars and trucks that were available for sale in the United States during the period 1979–2012. For each vehicle, we know the manufacturer, model, model year, body style, engine type, fuel tank size, and driveline. We also know the wheelbase, front track, and rear track measurements used to calculate each vehicle's footprint, as well as fuel type and vehicle category. We divide fuel types into four categories—gasoline, electric, electric hybrids, and flex-fuel—where gasoline includes vehicles fueled by gasoline, ethanol, natural gas, or diesel. We divide vehicle categories into "trucks" and "cars" following Environmental Protection Agency (EPA) guidelines which treat sports-utility vehicles, pickup trucks, minivans, and vans as trucks and compact, large, midsize, minicompact, subcompact cars, two seaters, and station wagons as cars. For each vehicle we also know the truncated vehicle identification number (VIN), which allows us to merge this information with fuel economy and emissions data from the EPA. For CAFE compliance purposes, the relevant vehicle fuel economy comes from the EPA's City and Highway test procedures (referred to as the "two-cycle" tests), not the measures used for vehicle labels. For summary statistics, see table 1.

We combine this information with data from Polk Automotive on all registered vehicles in the United States. These data are for the calendar year 2012 and provide census-tract level counts by vehicle type, including not only new cars but also the entire stock of older vehicles. Counts distinguish vehicles by manufacturer, model, model year, engine size, cylinders, and fuel type. That is, for each census tract it is not just that we know how many 2012 Toyota trucks there are. We know, for example, how many are flexible-fuel 2012 Toyota Tundra trucks with a 5.7L V8 engine. A limited number of vehicles are further differentiated by various trim levels, such as the Mercedes-Benz C1550, which is manufactured in both standard and "4Matic" models. However, these data do not distinguish between all the available model options that can be installed, such as leather upholstery. We merge this tract-level vehicle registration data from Polk with the vehicle specifications from DataOne using each vehicle type's truncated VIN. Finally, we use tract-level measures of mean and median household income from the American Community Survey (ACS).

3.2. Comparing Actual Emissions to Targets

We first calculate how actual emissions for each vehicle model compare to that vehicle's footprint-based target. For this exercise we focus on all vehicle models from the 2012 model year. Using each vehicle's footprint as well as whether it is a truck or a car, we calculate its emissions target using the official formula from NHTSA (2010, table III.B.2).

Figure 2 shows how each car from vehicle model year 2012 compares to its emissions target. Observations are scaled by each vehicle model's total sales in 2012. The x-axis is each vehicle's footprint, measured in square feet. The green solid line indicates

Table 1. Summary Statistics

	Mean	SD	Min	Max
A. All vehicle characteristics:				
Footprint (square feet)	49.3	8.2	26	111.5
Vehicle classified as "truck"	.49	.50	0	1
Emissions target (CO_2 grams per mile)	331.2	50.6	244	395
CAFE emissions (CO_2 grams per mile)	323.8	79.0	0	743.5
Fuel economy (miles per gallon)	29.2	7.8	12.0	74.3
Implied tax or subsidy ($)	227.0	515.5	−3,270.3	3,178.1
B. New vehicle characteristics:				
Footprint (square feet)	48.0	6.1	26.8	77.6
Vehicle classified as "truck"	.44	.50	0	1
Emissions target (CO_2 grams per mile)	297	39	244	395
CAFE emissions (CO_2 grams per mile)	272	71	0	628
Fuel economy (miles per gallon)	34.8	8.9	14.2	70.8
Implied tax or subsidy ($)	53.2	491.3	−2,237.4	2,556.4
C. Census tract characteristics:				
Population	4,627	2,814	0	55,283
Median household income ($000)	55.3	27.0	2.5	250
Mean household income ($000)	69.2	35.5	5.01	582.3

Note. This table describes vehicle characteristics and demographic characteristics for the United States in 2012. Panel A reports unweighted means across all vehicle models available for sale in the United States in 2012, and panel B only reports unweighted means for vehicles manufactured in 2012. Fuel economy is the unadjusted "combined" miles per gallon from the EPA/DOT label, which is calculated by EPA/DOT as a 55%/45% weighted average of city and highway miles per gallon. A vehicle classified as a "truck" is an indicator variable equal to one for vehicles that classify as a light truck under CAFE. For census tract characteristics we report unweighted means for 2012 across all US census tracts.

the emissions target in grams of carbon dioxide per mile. Most cars are within the upward-sloping part of the emissions target function though there are a sizable number of cars below 41 square feet in the flat portion of the function. Actual emissions vary significantly from the targets. This is particularly true above the line with a nonnegligible fraction of vehicles that emit more than twice as much carbon dioxide per mile as is targeted.

As indicated in the figure, there are several electric vehicles as well as a large number of flex-fuel vehicles. This includes several high-selling vehicle models, including the Ford Focus FFV, Chrysler 300 FFV, and Dodge Charger FFV. We calculated emissions for these vehicles using the standard CAFE assumptions for flex-fuel, so these vehicles show up as having low carbon dioxide emissions.

Figure 3 is the analogous figure for trucks for model year 2012. Most vehicles are within the upward-sloping part of the emissions target function, and again, there is a great deal of variation in emissions for any given footprint. There were no mass-
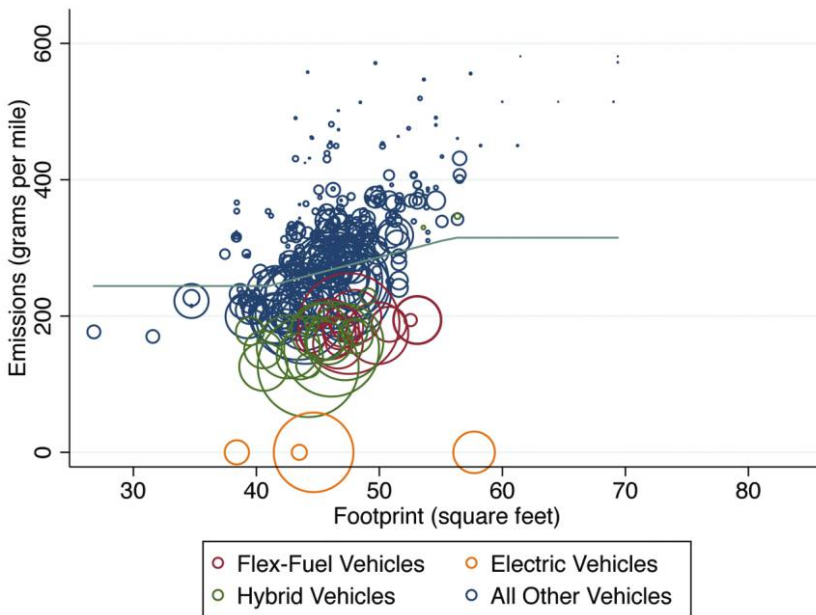
Figure 2. Each vehicle model relative to target, cars. Figure based on fuel economy from the EPA's City and Highway test for the 2012 model year and NHTSA standards. Following EPA guidelines, fuel economy for plug-in hybrids assumes a 50-50 gasoline and electricity split and fuel economy for flex-fuels assumes a 50-50 gasoline and E85 split. Circle sizes correspond to national sales for each vehicle model. The green solid line indicates the emissions target in grams of carbon dioxide per mile.

marketed electric or plug-in hybrid trucks but, again, there are a large number of flex-fuel vehicles. Flex-fuel vehicles are particularly common among trucks with very large footprints and many of these flex-fuel trucks sell in high volumes. The three best-selling flex-fuel trucks in 2012 were the Ford F150, Chevy Silverado, and Mercury Mariner.

In the following section, we take this information to calculate the implicit tax (or subsidy) for each vehicle model. In making these calculations we treat flex-fuel vehicles as regular vehicles. That is, we ignore the preferential treatment of flex-fuel vehicles. We do this because under CAFE there is a maximum limit on how much each manufacturer can use the flex-fuel feature and most automakers operate close to this cap.[6]

---

6. In the 2014 model year, for example, the flex-fuel limit in the CAFE program was 1.2 miles per gallon. See EPA (2015, 22). Thus, flex-fuel vehicles for that model year cannot increase a manufacturer's average fuel economy by more than 1.2 miles per gallon. EPA ensures that no manufacturer exceeds the maximum allowable value of the incentive by calculating fleet average
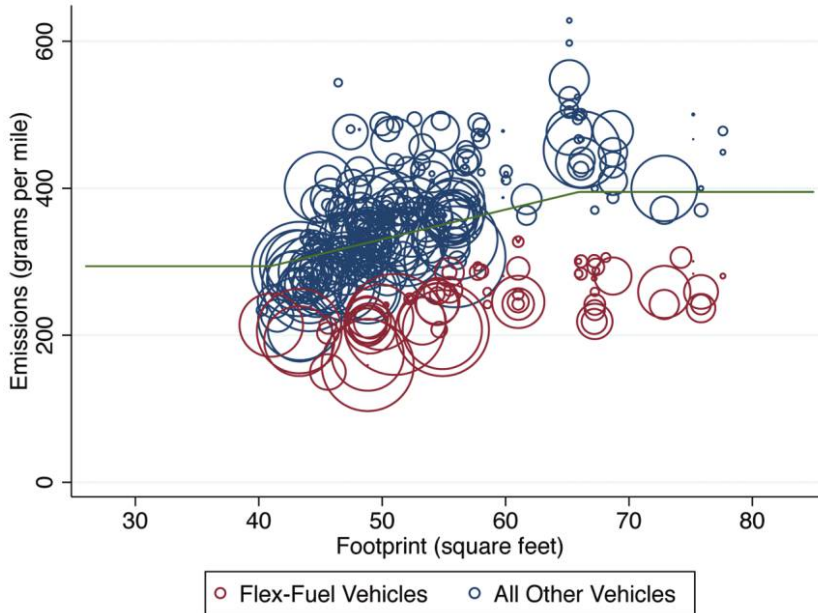
Figure 3. Each vehicle model relative to target, trucks. This figure plots each vehicle's carbon dioxide emissions (in grams per mile) against its footprint (in square feet). The horizontal upward-sloping line indicates the target emissions for each vehicle based on its footprint. Emissions are based on fuel economy from the EPA's City and Highway test for the 2012 model year and NHTSA standards. Following EPA guidelines, fuel economy for plug-in hybrids assumes a 50-50 gasoline and electricity split, and fuel economy for flex-fuels assumes a 50-50 gasoline and E85 split. Circle sizes correspond to national sales for each vehicle model.

Once at the cap, there is no benefit from selling an additional flex-fuel vehicle, so the automaker will not take this into account when pricing vehicles. Moreover, Anderson and Sallee (2011) show that, in practice, flex-fuel vehicles tend to sell for exactly the same price as equivalent regular vehicles.

3.3. Implicit Taxes and Subsidies

CAFE obligations are tradable under the new rules, so the permit prices make it possible to quantify the exact magnitude of these implicit subsidies and taxes. Using the first-order condition in equation (4), we calculate the implicit tax or subsidy imposed by the CAFE constraint as

---

fuel economy both with and without the incentive. Fiat Chrysler, Ford, GM, Jaguar Land Rover, Mercedes, and Volkswagen all maximized the usage of the flex-fuel incentive for model year 2014.

$$t_j = \lambda \left[\left(\text{emissions}_j - \text{target}_j \right) * VMT_j \right], \tag{7}$$

where t_j is the tax or subsidy borne by vehicle j. We use the observed permit price for λ and the EPA's standard assumption for the expected lifetime mileage, VMT_j.

Figure 4 shows the dollar amount of implied tax (subsidy) per car. The x-axis is each vehicle's footprint, measured in square feet. The solid line is a nonparametric estimation of the relationship between a vehicle's footprint and its implied tax using a local polynomial kernel-weighted regression. In general, cars with small footprints are more likely to be subsidized, while those with large footprints are more likely to be taxed. For cars with footprints between 40 and 60, the polynomial line is near flat around zero with a slight upward trend for cars, although most cars lie under the polynomial line and are subsidized for having emissions below their target.

Many of the most popular vehicle models are subsidized. Circle sizes in the figure are proportional to national sales, and many of the most popular vehicles are subsidized. The top three selling cars all have relatively high rates of subsidy; the Toyota Camry 2.5 liter version is subsidized at $260, Honda Civic 1.8 liter is subsidized at
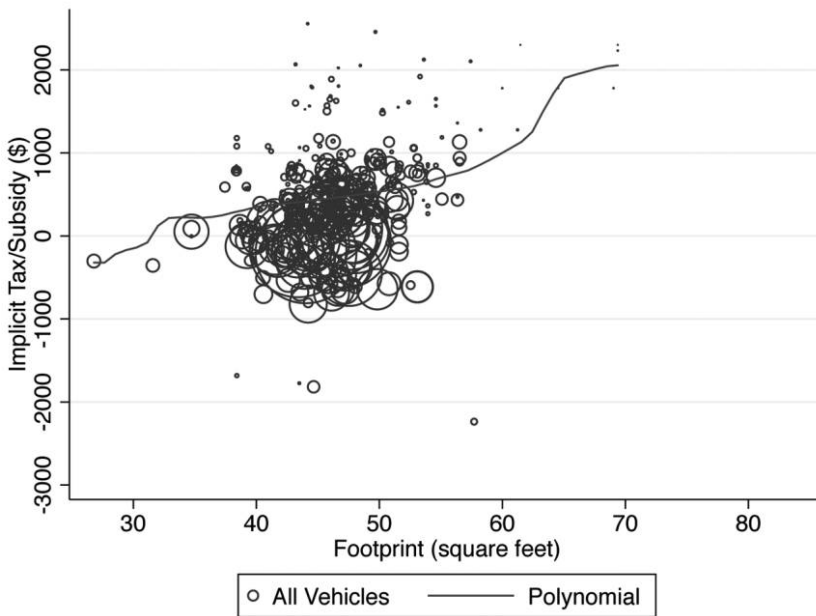

Figure 4. Implicit taxes and subsidies, cars. Implicit tax and subsidy calculations are based on fuel economy data from the EPA's City and Highway test for the 2012 model year, 2012 trading permit price, and vehicle expected lifetime mileage. Implicit taxes and subsidies are measured in 2012 dollars. Circle sizes are proportional to national sales. Color version available as an online enhancement.

$395, and the Nissan Altima 2.5 liter is subsidized at $161. Many of the cars with implicit taxes greater than $1,000 are expensive sports cars, and sales for these vehicles tend to be relatively low, as indicated by the small circle sizes.

Figure 5 is the corresponding plot for trucks. Overall the polynomial is much flatter for trucks. Trucks with small footprints are not more likely to be subsidized, and trucks with large footprints are only slightly more likely to be taxed. Circle size is again proportional to national sales. The top three most popular trucks are all subsidized. The Honda CR-V AWD is subsidized at $380, Chevy Silverado is subsidized at $1,312, and the Dodge Grand Caravan is subsidized at $1,314.

3.4. Distributional Impacts

In this section, we incorporate the vehicle registration data to describe distributional impacts. From the analysis in the previous section, we know the implied tax (or subsidy) for all new vehicles. In addition, fuel economy standards impact used vehicle prices, and we use an ad hoc approach for modeling the impact on used vehicles. We then use the vehicle registration data, and implied tax (or subsidy) for all vehicles, to calculate the average impact of fuel economy standards by census tract. Finally, we
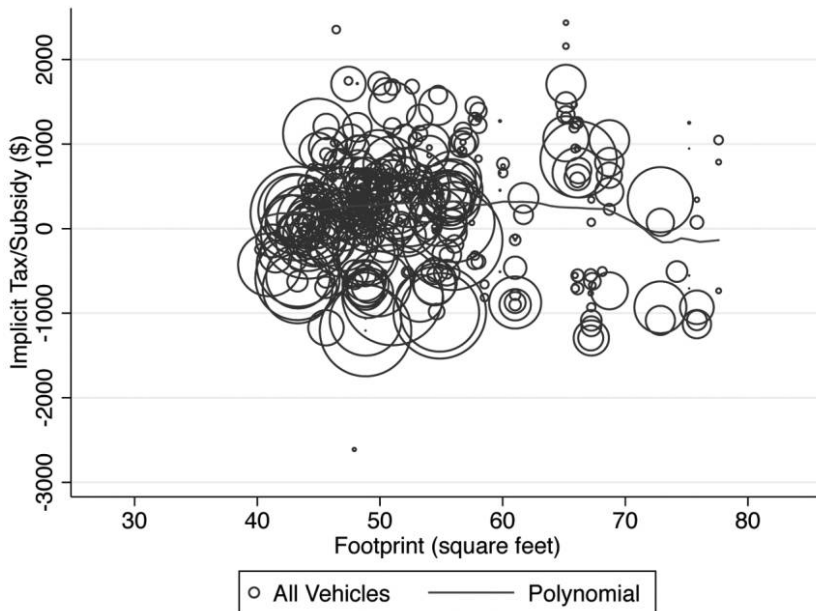

Figure 5. Implicit taxes and subsidies, trucks. Implicit tax and subsidy calculations are based on fuel economy data from the EPA's City and Highway test for the 2012 model year, 2012 trading permit price, and vehicle expected lifetime mileage. Implicit taxes and subsidies are measured in 2012 dollars. Circle sizes are proportional to national sales. Color version available as an online enhancement.

divide tracts by income decile to show how impacts differ across high- and low-income households.

Buyers substitute between new and used vehicles. Thus, when the price increases for a new vehicle, this pushes equilibrium prices upward for similar used vehicles. This same price increase for a new vehicle also means that there will be fewer of that type of vehicle in circulation, so these vehicles will continue to sell at a premium as they enter the used vehicle market. Eventually, however, vehicle scrappage decisions mitigate these impacts, as higher-priced vehicles are scrapped at systematically lower rates than lower-priced vehicles. For this reason and because they are poor substitutes for new vehicles, we would not expect fuel economy standards to have more than a negligible impact on very old vehicles.

Our approach for modeling the impact on used vehicles follows this economic intuition. In particular, we assume that the impact of fuel economy standards attenuates throughout a vehicle's lifetime according to the empirical average scrap rate in Jacobsen and van Benthem (2015). This approach assigns the 100% impact to vehicles that are less than one year old and a 0% impact for vehicles that are 18 years old. Take, for instance, the Toyota Camry, which has been sold continuously throughout our sample period and with approximately the same vehicle footprint. The 2007 model year Toyota Camry (3.5 liter) has an implied subsidy of $117 when new. According to Jacobsen and van Benthem (2015), approximately 7% of cars will be scrapped after 5 years. In order to find the implied subsidy for the 2007 Camry in 2012, the last year of our sample period, we scale the $117 value by 7%, yielding a subsidy of $109. We use separate scrap rates for cars and trucks, and scrap rates vary from year to year; as an example, Jacobsen and van Benthem (2015) estimates that 7%, 22%, and 51% of cars and 8%, 22%, and 46% of trucks will be scrapped after 5, 10, and 15 years, respectively. However, we do not distinguish between vehicle manufacturer or other vehicle characteristics. This attenuation formula is, admittedly, a strong assumption that ignores potentially important differences across vehicle classes, cross-price elasticities, and other factors. We argue that this nonetheless captures the general pattern of price impacts predicted by economic theory.[7]

From our vehicle registration data we have tract-level information on the number of vehicles by model and vintage. We use the implied tax (or subsidy) for each vehicle type, both new and used, to calculate the average tax (or subsidy) per vehicle in each tract. We report this and several related statistics in a series of results below. In all cases, we show results separately for new vehicles only and for new and used vehicles. Moreover, because our objective is to examine the distributional patterns, we divide tracts into deciles using mean household income.

---

7. Jacobsen and van Benthem (2015) report depreciation rates for vehicles over time. If we instead use these functions to adjust the CAFE taxes, the results below with respect to the regressivity of CAFE are exacerbated.

3.4.1. Average Tax per Vehicle

Figure 6 shows box plots of the average tax per vehicle by income decile. As usual with box plots, the middle line indicates the median, a shaded box indicates the inner-quartile range (IQR), and "whiskers" indicate 1.5 times the IQR. Whether we examine new vehicles only or both new and used vehicles, the median is fairly similar across deciles. Between high- and low- income tracts, the median tax per vehicle varies by only about $100.

Although the median is quite similar across deciles, there is a modest increase in the average tax per vehicle in the eighth, ninth, and, particularly, in the tenth decile. This reflects higher-income households buying vehicles that, relative to their footprint, are less fuel efficient on average. High performance luxury vehicles, for example, are in this category. The differences across decile are relatively small, however, relative to the range of implicit taxes across individual vehicle models in the previous analysis, reflecting a relatively weak correlation between income and the implicit taxes from CAFE.

The figure also shows that there is a large amount of variation within decile. The "whiskers" for all 10 deciles include both positive values (taxes) and negative values (subsidies), so there is a wide range of experiences across tracts in any given decile. This wide variation points to fuel economy standards being an imprecise instrument for redistribution. Lower-income households buy a range of different types of vehicles, including very fuel-inefficient vehicles, so CAFE increases costs for these households, even while decreasing costs for others.

It is worth emphasizing again that these calculations assume that the taxes and subsidies imposed by CAFE are borne entirely by vehicle buyers rather than automakers or retailers. This assumption is most reasonable in segments of the vehicle market for which supply is highly elastic or for which demand is highly inelastic. Incomplete pass-through of the taxes and subsidies imposed by CAFE standards would reduce the overall magnitude of our estimated impacts as well as, potentially, by changing the relative burdens borne by different income groups. If lower-income buyers are more price elastic (West 2004), then they would tend to bear a smaller fraction of both the taxes and subsidies from CAFE, thereby making CAFE more progressive.

3.4.2. Average Tax as a Share of Income

Figure 7 plots the average tax per vehicle as a share of median income. For each tract we know the average tax (or subsidy) per vehicle, and for this figure we divide this by the median income in the tract. We report shares as a percentage, so 1 is 1% of income. As before, we then divide census tracts into income deciles. Effectively, we are calculating the average tax as a share of median income for each of the deciles.[8] Income

---

8. This calculation masks within-tract variation in vehicles and incomes. Because the main focus of the paper is how the incidence changes across income deciles, our implicit assumption is that the same slope of incidence to income is present both across and within deciles.
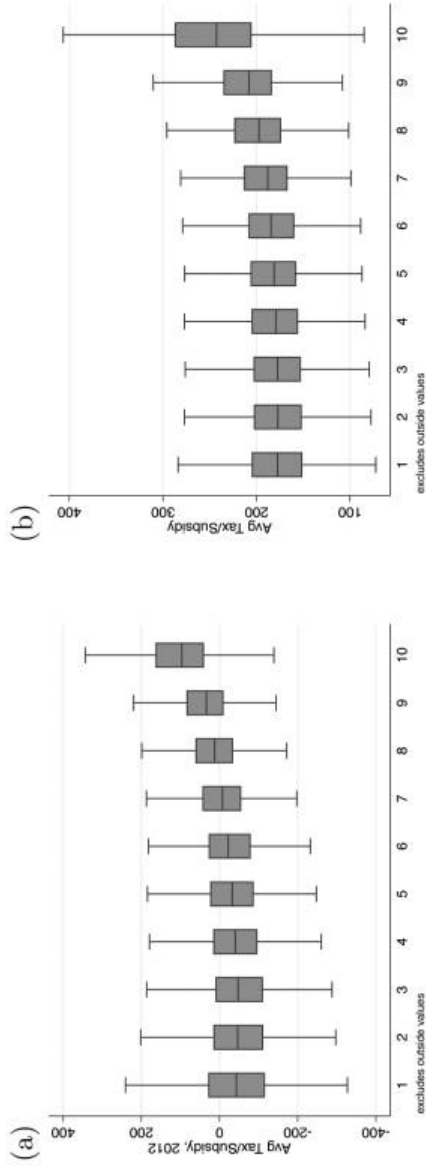
Figure 6. Average tax per vehicle, by income deciles. *a*, New vehicles. *b*, All vehicles.
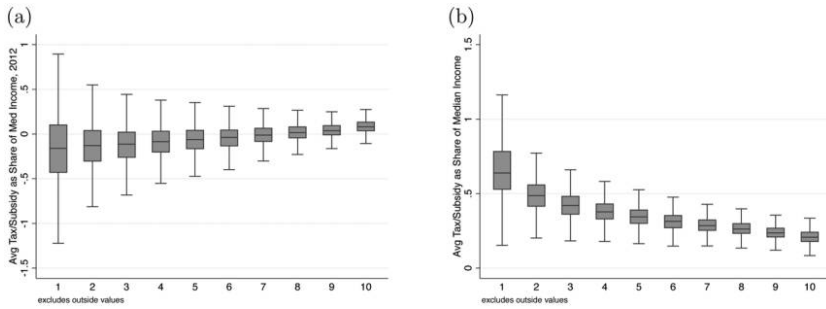
Figure 7. Average tax per vehicle as a share of median income, by income decile. *a*, New vehicles. *b*, All vehicles.

levels vary widely across the deciles, so these scaled results look considerably different from the previous figure.[9]

When we consider new vehicles only, CAFE is slightly progressive. That is, high-income households bear a bit more tax per vehicle as a share of income. However, this pattern reverses sharply once we include used vehicles. The second panel of figure 7 shows that the median household in the bottom decile bears a tax that is about 0.6% of household annual income. This is more than twice the burden experienced by the median household in the top decile. Even though the average tax per vehicle tends to be higher in the top income deciles, it is being divided by a much higher level of income, yielding a regressive pattern overall.

Overall, the median impacts are small when viewed as a fraction of annual income. Even for the bottom decile, the average tax per vehicle is well below 1% of annual income. Keep in mind, moreover, that households buy a vehicle only once every couple of years so the annualized cost of CAFE is considerably smaller. That said, the box plots also show that there is a wide range of experiences across tracts, with some tracts experiencing impacts that are twice as large as the median. There is also a wide range of experiences across households within tracts. While the average tax per vehicle is relatively small as a share of annual income, there are individual households who bear much larger costs.

3.5. Comparison to Other Studies of CAFE

There have been very few previous attempts to measure the distributional effects of fuel economy standards. An important exception is Jacobsen (2013), which studies

---

9. The reader will note the larger amount of variation in this measure for lower-income deciles, while the distributions were similar across deciles for the average tax in fig. 6. The standard deviation of median incomes across tracts within decile divided by the median income is largest for decile 1, driving this large variation for the lower deciles.

the distributional impacts of CAFE using data from the 2001 National Household Travel Survey. The estimates from Jacobsen (2013) are not directly comparable as they are based on older data, and for the older version of CAFE without footprint-based targets. The paper is nonetheless an important point of comparison as it is one of the very few studies to look at distributional impacts, and because the study is unusually well done, with an innovative dynamic model that in many ways goes far beyond the scope of our study.

In Jacobsen (2013) the estimates of distributional impacts come out of a rich equilibrium model in which customers choose vehicles to maximize utility and firms price vehicles to maximize profits. Moreover, used cars are incorporated explicitly into the model, with consumers choosing between new and used vehicles, and with older vehicles being scrapped over time endogenously. While not capturing the full long-run impacts of CAFE (e.g.. endogenous characteristics or innovation), the model does capture how fuel economy standards lead customers to substitute to more fuel-efficient vehicles, behavior that is central to understanding the full welfare impact of CAFE.

Jacobsen (2013) finds that the distributional effects of fuel economy standards depend critically on incorporating used cars. When only the impact on new vehicles is considered, CAFE has approximately the same welfare impact on high- and low-income households as a fraction of income. Higher-income households are more likely to own new cars, but also have much higher incomes. However, once used vehicles are incorporated CAFE becomes sharply regressive, with low-income households experiencing welfare losses that are three times as large as a percentage of income as those experienced by high-income households. It makes sense that fuel economy standards would increase the price and alter the composition of used vehicles, but Jacobsen (2013) was one of the first studies to capture these dynamics, and the first to show the implications for distributional impacts.

The one other paper that estimates the distributional impact of fuel economy standards is Levinson (2019, in this issue). Using data from the 2009 National Household Travel Survey, Levinson (2019, in this issue) shows that high-income households spend on average four times more on gasoline than low-income households, while having about three times as many vehicles of approximately the same average fuel economy. Consequently, fuel economy standards, which are essentially a tax per vehicle, are more regressive than a gasoline tax. Moreover, Levinson (2019, in this issue) finds that high-income households tend to buy higher footprint vehicles, so they do relatively better under footprint-based standards and, thus, footprint-based standards are even more regressive than regular standards.

3.6. Comparison to Gasoline Tax
We can now compare the distributional impacts of CAFE with previous estimates in the literature on the distributional impacts of a gasoline or carbon tax. Previous studies

have examined the distributional effects of both gasoline taxes (Poterba 1989, 1991; West 2004; Bento et al. 2009) and carbon taxes (e.g., Burtraw et al. 2009; Hassett et al. 2009; Williams et al. 2015).

A common theme in the existing literature is that the distributional effects of gasoline and carbon taxes largely depend on how additional revenue from the tax is recycled. For example, gasoline taxes are significantly less regressive when revenues are used to cut labor tax rates than when revenue is discarded (West and Williams 2004; Williams et al. 2015). Bento et al. (2009) show that if returned lump sum on a per capita basis, a gasoline tax could make the bottom four income deciles better off on average, even without incorporating external benefits. Burtraw et al. (2009) analyze five different uses for revenues raised from cap-and-trade auction, including lowering income and payroll taxes, and find significant differences in progressivity. Similarly, Rausch et al. (2010) simulate the distributional effects of carbon taxes under two redistribution bundles with varying amounts of revenue set aside for deficit reductions or cuts for other taxes. They find all scenarios to be progressive in lower-income deciles and proportional in upper deciles; however, the degrees of incidence exhibit significant differences over time. Meanwhile, Metcalf (2009) proposes a distributional-neutral carbon tax by offsetting price increase with capped income tax credit.

Several studies have shown that the regressive implications of carbon taxes can be overstated by overlooking index government transfer programs (Rausch et al. 2010; Dinan 2012; Fullerton et al. 2012). Several government transfer programs, including Supplemental Security Income, are indexed to consumer price measures and thus increase alongside the price of carbon. For example, accounting for transfer program indexing, Rausch et al. (2010) use a computable general equilibrium model of the United States to show that carbon taxes are moderately progressive, even ignoring distribution of after-tax revenue. They find that the lowest two income quintiles are made better off under a carbon tax. In part, the results in Rausch et al. (2010) reflect that a portion of the carbon price is shifted back to the owners of natural resources and capital, which lessens the regressivity of a carbon tax policy. Finally, Fullerton et al. (2012) show that under partial indexing of transfer funds, a carbon tax is progressive for households in the bottom half of the expenditure distribution.

A thorough review of each of these studies focusing on the methods and underlying assumptions is beyond the scope of this paper. However, we attempt to summarize the literature by regressing each paper's incidence measure on the income decile number, where 1 is the lowest decile and 10 is the highest. That is, for each of the papers, we regress the incidence as a share of income on the decile number. For example, our analysis implies that the incidence as a share of median income for deciles 1 through 10, respectively, is: 0.111, 0.167, 0.121, 0.133, 0.121, 0.0926, 0.116, 0.139, 0.107, and 0.108. Regressing these on the decile yields an average slope of −.003; this is slightly regressive. A progressive policy would imply a positive slope.

The coefficient provides a rough idea of the average change in incidence when moving one decile. Given the summary nature of the analysis, we omit standard errors from the table.

Table 2 summarizes this exercise. Our results are listed in the first row. Each study is listed in the first column and studies are separated into three panels based on the type of recycling: no recycling, lump-sum transfers, and tax cuts. Further information on the recycling treatment is listed in the Notes column. For example, under the lump-sum transfers panel, the notes indicate whether lump-sum transfer was uniform or proportional to household income. Similarly, within the tax cuts panel, the notes list whether income, labor, or payroll taxes were lowered using tax revenue. We also track which studies use transfer indexing. The fourth and final column codes each carbon tax treatment as regressive (R), progressive (P), or a combination of the two if the direction of incidence switches. In the case of the latter, the parentheses by each letter contain which income groups are progressive and which are regressive.

We measure the relative progressivity or regressivity of the tax by the slope of the mean welfare impacts across household income group. We report slopes for welfare impacts measured as a per capita share of income as well as the level of incidence when available. The slope captures the direction of the incidence—negative slopes correspond to regressive taxes and positive slopes to progressive taxes—as well as the magnitude. We calculate the slope by running a linear regression of the welfare measure on decile as described above, except Bento et al. (2009), which is measured in quartiles, and West and Williams (2004) and Williams et al. (2015), which are measured in quintiles. We adjust the slopes for these two studies by multiplying the slope from the regression by the number of income categories divided by 10. This makes the slopes we report in the table comparable to a per-decile change.

Table 2 shows that CAFE is more regressive than a carbon tax with lump-sum transfers. The exact magnitude of the difference depends on which study is used for comparison, but in most cases CAFE is considerably more regressive. Moreover, CAFE is significantly more regressive than a carbon tax used to finance an expansion of the earned income tax credit (EITC), or presumably, other tax credits aimed at lower-income households. In contrast, CAFE is more progressive than carbon taxes used to reduce progressive taxes, such as labor and payroll taxes. Again, the exact magnitude differs across studies, but in all cases CAFE is considerably more progressive. Labor and payroll taxes are distortionary, so there are efficiency gains from using carbon tax revenue in this way, but these efficiency gains must be balanced against a pronounced negative impact on equity.

4. CONCLUSION

Economists have long complained that fuel economy standards are an inefficient way to reduce gasoline consumption. In a survey of top economists, 90% answered that

Table 2. Comparing the Distributional Impact of CAFE to a Carbon Tax

Paper	Notes	Slope, Share	Slope, Level	R/P
Davis and Knittel		−.003	6.481	
	Recycling Method: No Recycling			
Burtraw et al. (2009)	None	−.299		R
Fullerton et al. (2012)	Partial transfer indexing	−.217		R
	Full transfer indexing	−.214		R
Hassett et al. (2009)	None	−.296		R
West (2004)	None	−.022		P(1:6), R(6:10)
West and Williams (2004)	None	−.144		P(1:2), R(2:5)
	Recycling Method: Lump-Sum Transfers			
Bento et al. (2009)	Uniform	.083	38.773	P
	Proportional to income	−.017	−6.454	P(1:8), R(8:10)
	Proportional to VMT	.007	9.033	P
Burtraw et al. (2009)	Uniform	.164		P
	Uniform, taxed	.360		P
West and Williams (2004)	Uniform	.331		P
Williams et al. (2015)	Uniform, full transfer indexing	.617		P
	Recycling Method: Tax Cuts			
Burtraw et al. (2009)	Labor tax	−.444		R
	Payroll tax	−.379		R
	Expanded EITC	.474		P
West and Williams (2004)	Labor tax	−.083		R
Williams et al. (2015)	Labor tax, full transfer indexing	.020		R(1:2), P(2:5)
	Capital tax, full transfer indexing	−.116		R

Note. Slope, Share is the slope of mean EV as a share of income across decile and Slope, Level is the slope of mean tax amount across decile. R/P codes if the tax is regressive (R), progressive (P), or a combination of the two. The range is listed in parentheses if the slope is a combination. All slopes calculated using linear regression of welfare measure on income group.

they would prefer a gasoline tax over fuel economy standards (IGM 2016).[10] Fuel economy standards do not achieve the efficient level of vehicle usage, nor do they create efficient incentives for owners to scrap older fuel-inefficient cars, nor do they efficiently distinguish between vehicle models with different average longevities (Jacobsen et al. 2018).

Yet policy makers repeatedly turn to fuel economy standards. In part, this preference reflects a view that gasoline taxes are regressive. A growing literature in economics shows, however, that the regressivity of a gasoline tax depends critically on what is done with the revenues that are generated. If revenues are returned to households progressively, or even through lump-sum transfers, then a gasoline tax can be made to be strongly progressive, with, for example, the bottom four income deciles being made better off, on average (see, e.g., Bento et al. 2009).

Thinking about whether fuel economy standards are regressive or progressive is less intuitive because the costs are less salient. Fuel economy standards do impose costs, however. We show that standards impose a constraint in the automakers' profit maximization problem that imposes an implicit tax (or subsidy) on each vehicle type. These implicit prices can be exactly calculated in our context because under US rules automakers can trade obligations, so permit prices provide a direct measure of the shadow value of the constraint. We then combine this information with rich microdata on vehicle registrations to estimate the distributional impacts.

When we consider new vehicles only, we find that CAFE is mildly progressive. But, of course, fuel economy standards impact not only new vehicles but used vehicles as well. When we include used vehicles, the pattern reverses and we find that CAFE is mildly regressive. High-income households bear less cost as a fraction of income than low-income households. Thus fuel economy standards are more regressive than a gasoline tax with revenues returned lump sum. We conclude, therefore, that it is difficult to argue for fuel economy standards on the basis of distributional concerns.

REFERENCES

Anderson, Soren T., Ian W. H. Parry, James M. Sallee, and Carolyn Fischer. 2011. Automobile fuel economy standards: Impacts, efficiency, and alternatives. *Review of Environmental Economics and Policy* 5 (1): 89–108.

Anderson, Soren T., and James M. Sallee. 2011. Using loopholes to reveal the marginal cost of regulation: The case of fuel-economy standards. *American Economic Review* 101 (4): 1375–1409.

———. 2016. Designing policies to make cars greener. *Annual Review of Resource Economics* 8:157–80.

---

10. There is a strikingly high degree of agreement among economists on this question. In this survey 51% and 39% of economists answered that they "strongly agreed" and "agreed," respectively, that a carbon tax would be a less expensive way to reduce carbon-dioxide emissions than fuel economy standards. This is a high level of agreement compared to other questions in the same survey (Sapienza and Zingales 2013).

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 (3): 562–82.

Bento, Antonio M., Lawrence H. Goulder, Mark R. Jacobsen, and Roger H. Von Haefen. 2009. Distributional and efficiency impacts of increased U.S. gasoline taxes. *American Economic Review* 99 (3): 667–99.

Borenstein, Severin, James Bushnell, Frank A. Wolak, and Matthew Zaragoza-Watkins. 2016. Expecting the unexpected: Emissions uncertainty and environmental market design. Energy Institute at Haas Working paper.

BP Energy. 2015. BP energy outlook 2035. http://www.bp.com/en/global/corporate/energy-economics .html.

Burtraw, Dallas, Richard Sweeney, and Margaret Walls. 2009. The incidence of U.S. climate policy: Alternative uses of revenues from a cap-and-trade auction. *National Tax Journal* 62 (3): 497–518.

Dinan, Terry. 2012. Offsetting a carbon tax's costs on low-income households. Working paper 43713, Congressional Budget Office.

EIA (Energy Information Administration). 2016. International energy outlook 2016. EIA, US Department of Energy. http://www.eia.gov/forecasts/ieo/.

EPA (US Environmental Protection Agency). 2015. Greenhouse gas emission standards for light-duty vehicles: Manufacturer performance report for the 2014 model year.

———. 2016. Draft technical assessment report: Midterm evaluation of light-duty vehicle greenhouse gas emission standards and corporate average fuel economy standards for model years 2022–2025. https:// www3.epa.gov/otaq/climate/documents/mte/420d16900.pdf.

Fullerton, Don, Garth Heutel, and Gilbert E. Metcalf. 2012. Does the indexing of government transfers make carbon pricing progressive? *American Journal of Agricultural Economics* 94 (2): 347–53.

Gillingham, Kenneth. 2013. The economics of fuel economy standards versus feebates. Unpublished manuscript, Yale University.

Hassett, Kevin, Aparna Mathur, and Gilbert E. Metcalf. 2009. The incidence of a U.S. carbon tax: A lifetime and regional analysis. *Energy Journal* 30 (2): 155–77.

Helfand, Gloria E. 1991. Standards versus standards: The effects of different pollution restrictions. *American Economic Review* 81 (3): 622–34.

Holland, Stephen P., Jonathan E. Hughes, and Christopher R. Knittel. 2009. Greenhouse gas reductions under low carbon fuel standards? *American Economic Journal: Economic Policy* 1 (1): 106–46.

Holland, Stephen P., Erin T. Mansur, Nicholas Z. Muller, and Andrew J. Yates. 2016. Environmental benefits from driving electric vehicles? *American Economic Review* 106 (12): 3700–3729.

IGM (Initiative on Global Markets). 2016. IGM economic experts panel. http://www.igmchicago.org/igm -economic-experts-panel.

Interagency Working Group on Social Cost of Carbon. 2013. Technical support document: Technical update of the social cost of carbon for regulatory impact analysis, under Executive Order 12866. http:// www.whitehouse.gov/sites/default/files/omb/assets/inforeg/technical-update-social-cost-of-carbon -for-regulator-impact-analysis.pdf.

IPCC (Intergovernmental Panel on Climate Change). 2015. *Climate change 2014: Mitigation of climate change*. Vol. 3. Cambridge: Cambridge University Press.

Ito, Koichiro, and James M. Sallee. 2018. The economics of attribute-based regulation: Theory and evidence from fuel-economy standards. *Review of Economics and Statistics* 100 (2): 319–36.

Jacobsen, Mark R. 2013. Evaluating U.S. fuel economy standards in a model with producer and household heterogeneity. *American Economic Journal: Economic Policy* 5 (2): 148–87.

Jacobsen, Mark R., Christopher R. Knittel, James M. Sallee, and Arthur A. van Benthem. 2018. The use of regression statistics to analyze imperfect pricing policies. Working paper, University of California, San Diego.

Jacobsen, Mark R., and Arthur A. van Benthem. 2015. Vehicle scrappage and gasoline policy. *American Economic Review* 105 (3): 1312–38.

Klier, Thomas, and Joshua Linn. 2016. The effect of vehicle fuel economy standards on technology adoption. *Journal of Public Economics* 133:41–63.

Knittel, Christopher R. 2012. Reducing petroleum consumption from transportation. *Journal of Economic Perspectives* 26 (1): 93.

Kwoka, John E. 1983. The limits of market-oriented regulatory techniques: The case of automotive fuel economy. *Quarterly Journal of Economics* 98 (4): 695–704.

Leard, Benjamin, and Virginia McConnell. 2017. New markets for pollution and energy efficiency: Credit trading under automobile greenhouse gas and fuel economy standards. *Review of Environmental Economics and Policy* 11 (2): 207–26.

Levinson, Arik. 2019. Energy efficiency standards are more regressive than energy taxes: Theory and evidence. *Journal of the Association of Environmental and Resource Economists* 6 (S1): S7–S36.

Metcalf, Gilbert E. 2009. Designing a carbon tax to reduce US greenhouse gas emissions. *Review of Environmental Economics and Policy* 3 (1): 63–83.

NHTSA (National Highway Traffic Safety Administration). 2010. Light-duty vehicle greenhouse gas emission standards and corporate average fuel economy standards: Final rule. *Federal Register* 40: 25323–728.

Poterba, James M. 1989. Lifetime incidence and the distributional burden of excise taxes. *American Economic Review* 79 (2): 325–30.

———. 1991. Is the gasoline tax regressive? In *Tax policy and the economy*. Vol. 5:145–64. Cambridge, MA: MIT Press.

Rausch, Sebastian, Gilbert E. Metcalf, John M. Reilly, and Sergey Paltsev. 2010. Distributional implications of alternative US greenhouse gas control measures. *B. E. Journal of Economic Analysis and Policy* 10 (2).

Sallee, James M. 2011. The surprising incidence of tax credits for the Toyota Prius. *American Economic Journal: Economic Policy* 79 (2): 189–219.

Sapienza, Paola, and Luigi Zingales. 2013. Economic experts versus average Americans. *American Economic Review Papers and Proceedings* 103 (3): 636–42.

West, Sarah E. 2004. Distributional effects of alternative vehicle pollution control policies. *Journal of Public Economics* 88 (3): 735–57.

West, Sarah E., and Roberton C. Williams. 2004. Estimates from a consumer demand system: Implications for the incidence of environmental taxes. *Journal of Environmental Economics and Management* 47 (3): 535–58.

Williams, Roberton C., 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): 195–214.