The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel-Economy Standards

Koichiro Ito
Boston University and NBER

James M. Sallee*
University of Chicago and NBER

October 22, 2013

Abstract

In many countries, fuel-economy standards mandate that vehicles meet a certain fuel economy, but heavier or larger vehicles are allowed to meet a lower standard. This has the perverse implication of allowing automakers to meet standards either by improving fuel economy or by increasing weight, which lowers fuel economy and increases externalities related to accidents. This is but one example of an attribute-based regulation, in which the subsidy, tax or regulation imposed on a product is a function not just of the amount of an externality that the product generates, but also how each product's externality compares to that of other products deemed to be similar by virtue of a commonality in some other attribute. Such policies are ubiquitous, but the core logic and welfare consequences of their deployment have not been studied by academic economists. This paper develops an analytical framework that captures the central implications of attribute-based policies, characterizes the deadweight loss caused by attribute-basing, and establishes situations in which attribute-basing may be efficient. The paper then empirically examines the consequences of attribute-based fuel economy standards in Japan, where fuel economy standards are an attribute-based function of vehicle weight. We use cross-sectional and panel techniques to demonstrate that attribute-based regulation has significantly altered the distribution of vehicle weight in Japan. We estimate that this alteration generates a welfare loss on the order of $200 per car sold in Japan, which translates into a $1 billion annual loss.

Keywords: fuel economy standards, corrective taxation, notches
JEL: H23, Q48, L62

*Ito: ito@bu.edu. Sallee: sallee@uchicago.edu. The authors would like to thank Kunihito Sasaki for excellent research assistance.

1 Introduction

The goal of this paper is to advance economists' understanding of a feature called attribute-basing that is common to many policies aimed at correcting externalities but is apparently at odds with the basic logic of Pigouvian taxation. Attribute-based policies are policies in which the corrective policy—which could be a regulation, a tax or a subsidy—imposed on a product is not simply a function of how much externality the product creates, but instead depends on how much externality the product creates relative to some benchmark, which itself depends on some other attribute of the product. Attribute-basing is common to corrective policies, in particular policies aimed at improving the energy efficiency of energy consuming durable goods, but there is no academic literature that explores the optimality, merits and problems with such policies.

Attribute-basing is perhaps best explained by example. Corporate Average Fuel Economy (CAFE) standards in the United States became attribute-based as of 2012. Prior to 2012, every automaker had to sell vehicles that met or exceeded a fleet-average fuel economy of some particular level, which was common to all producers. Since 2012, CAFE specifies a schedule of fuel economy targets for each vehicle that is a declining function of the vehicle's "footprint"—the rectangular area inside of the vehicle's tires. Each automaker has to meet a standard equal to the average target of the vehicles that it sells, so that automakers who sell larger vehicles on average are permitted to have less efficient vehicles.

Figure 1 illustrates this schedule for 2012 and shows the large difference in targets across vehicles and automakers. The Honda Fit (a compact car) has a target fuel economy of 36.2 miles per gallon, whereas the Chrysler 300 (a full-size sedan) has a target of 29.4. The steepness of this schedule creates notable differences in the fuel economy required of different automakers who produce cars of different average size. Ford and Toyota, whose targets differ by 1.4 miles-per-gallon in 2012, are labeled in the figure. Vehicle footprint is not immutable, however, so the footprint-based standard not only creates dispersion in requirements across automakers, but also creates an incentive for automakers to enlarge footprint, which runs counter to the goals of the policy by endorsing larger (and hence less energy efficient) vehicles.

Such attribute-basing is ubiquitous in the realm of energy efficiency policies. Appliance standards are typically a function of the size or features of the appliance. The same is true for the

Figure 1: Example of an attribute-based standard: 2012 U.S. new car CAFE

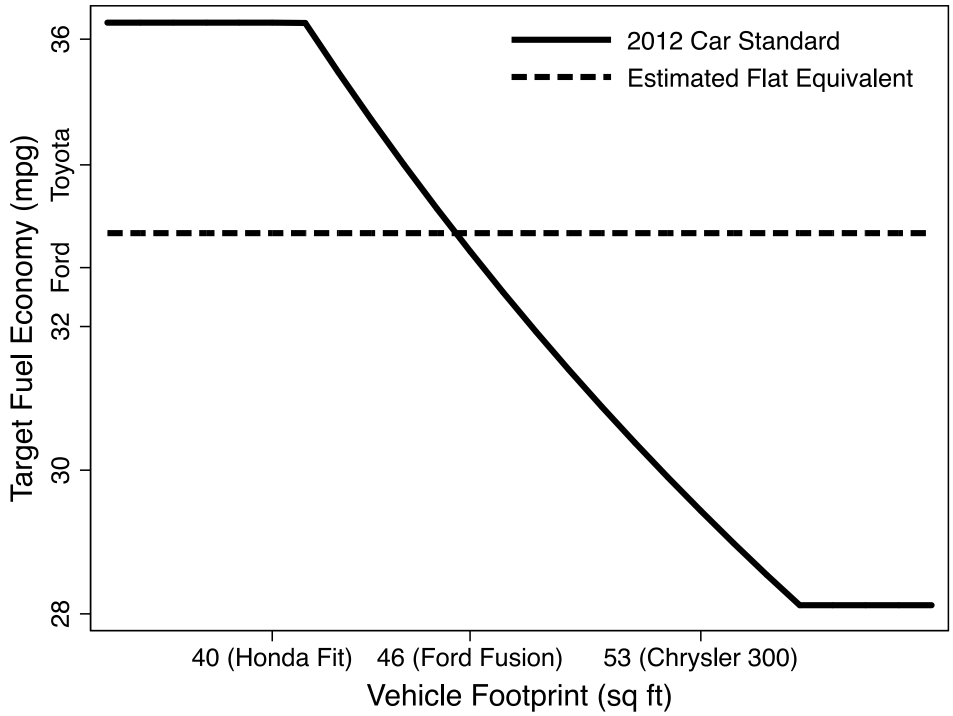

As of 2011, the CAFE target fuel economy for a vehicle depends on its footprint. Three footprint values are labeled with example vehicles that have that footprint. The flat equivalent is the EPA's estimate based on the distribution of footprint in current fleet. Prior to 2011, the new car CAFE standard was a flat standard at 27.5 mpg. The firm specific target values for Ford and Toyota based on model year 2011 average footprint is labeled. Data taken from the Federal Register, 40 CFR Part 85.

cutoff value to qualify an appliance for Energy Star certification, which is frequently also used as the criteria for appliance rebate programs meant to conserve energy. The hybrid vehicle tax credit was attribute-based because the subsidy amount was based on each hybrid's gasoline conservation relative to a benchmark vehicle, determined by its market segment. The U.S. only recently moved to attribute-basing for CAFE, but such policies have long-existed in other countries. For example, the Chinese fuel economy standard is an attributed-based function of engine displacement. In Japan, the regulation and government subsidy programs are sliding scales of a vehicle's weight. Circulation and registration taxes in Germany, France, Sweden and the United Kingdom depend on a vehicle's engine size. Label ratings, which identify a vehicle on a scale of most to least polluting in various places, including Europe and Korea, depend on vehicle size. All of these policies provide a lower standard for products that are larger and more polluting.

In spite of the frequent appearance of attribute-basing in actual policies, the academic literature has provided little scrutiny of its potential costs and benefits. To the extent that these policies are aimed at correcting externalities related to energy consumption, attribute-basing would seem to have no role in an optimal policy. If there is an externality related to a product's energy consumption, then the first-best policy will be a Pigouvian tax on energy consumption and it should not depend on any other attribute of the product. By making the standard or subsidy cutoff lower for larger, more polluting durable goods, attribute-basing typically gives market actors an incentive to comply with regulations by redesigning products to be larger or by pushing sales towards larger models. On its face, this would seem to have a negative welfare implication when compared to the Pigouvian ideal, because this incentive to "upsize" products runs counter to the initial goal of reducing externalities that are positively correlated with size.

Might there be additional considerations that make attribute-basing a good idea? If so, what would characterize the optimal design of an attribute-based policy? If not, so that attribute-basing is inefficient, what distortions does it create and what determines the size of any associated deadweight loss? Is there empirical evidence that distortions exist?

This paper aims to make an important first step in the analysis of attribute-basing by proposing answers to these questions and generally evolving our understanding of this class of public policy. To do so, we first provide additional background regarding a set of attribute-based energy-efficiency policies in section 2. The purpose of this section is to describe the professed logic of policymakers who have created the policies, so that the apparent aims of the policies can be evaluated in our theoretical framework, which we lay out in section 3.

Section 3 establishes an analytical framework that encourages a focus on questions traditionally asked in the public finance literature about corrective tax policy design. We demonstrate conditions under which attribute-basing is distortionary. In those cases, we show that the first-order distortion from attribute-basing is in the allocation of the attribute itself, rather than in the allocation of energy efficiency. This is especially consequential when the attribute used for targeting has externalities itself, which is the case for vehicle weight or footprint. We also show that attribute-basing could be justified when there are technological spillovers across producers or when the attribute-regulation applies to each product in the market as opposed to an average across multiple products.

In neither case, however, does the shape of the attribute function resemble the shape common to actual policies.

The paper then focuses on empirical evidence from Japanese fuel economy standards, which are particularly amenable for analysis, in section 4. Japanese fuel economy standards require that automakers meet a minimum fuel economy standard, but the standard depends on each vehicle's weight, with heavier vehicles needing to meet laxer standards. Moreover, the minima are discrete functions of weight class, so that vehicles on either side of a weight threshold are required to meet discretely different standards. These "notches" provide a straightforward way to identify distortions in the attribute in response to policy.

We show that Japanese automakers increase the weight of a substantial fraction of their vehicles in response to the incentives implicit in the weight-based fuel economy regulations. With cross-sectional data, we provide transparent evidence of distortions in the distribution of vehicle weight. Using panel data and a regression discontinuity design, we estimate an elasticity of the change in weight with respect to the fuel economy standard's stringency and conclude that a one percent increase in the stringency of the standard leads to a .16 percent increase in vehicle weight. We estimate that attribute-basing raises weight by around 10 to 20 kilograms across all vehicles in Japan. Based on recent estimates of the value of externalities related to vehicle weight, we estimate that this creates welfare losses on the order of $200 per car sold, which aggregates to roughly $1 billion annually across the new car market in Japan.

To the best of our knowledge, we are the first paper to take an analytical approach to assessing the merits and distortions caused by attribute-basing in policies aimed at correcting externalities. Our work is nevertheless related to several strands of literature. The literature on the move to footprint-based CAFE rules in the U.S. include Whitefoot and Skerlos (2012) and Whitefoot, Fowlie, and Skerlos (2013), which use engineering estimates of design costs and a discrete-choice economic model to predict how much automakers will manipulate footprint in response to a tightening of CAFE standards. Both studies conclude that an increase in footprint will be a major source of adjustment, which is broadly consistent with our findings of considerable manipulation of weight in the Japanese car market. Gillingham (2013) discusses the implicit incentive for the expansion of footprint in a broader discussion of CAFE policies. Jacbosen (2013) addresses the safety impacts of footprint-based standards in the U.S., which our work will also touch on. Our consideration of notched attributed-based policies relates to the literature on notched corrective taxation, which began with Blinder and Rosen (1985), and includes prior analysis of automobile fuel economy in Sallee and Slemrod (2012).

2 Key facts about attribute-based policies

Before proceeding to our analytical model and empirical analysis, it is useful to specify some of the features of attribute-based policies that we aim to better understand through our analysis. This gives us a sense of policymakers' intended consequences against which to judge policies, as well as aiding in identification of the most important research questions.

First, policymakers appear to (at least sometimes) determine the slope of the attribute-based function in order to mimic an isocost frontier; that is, they trace out the trade-off between efficiency and the attribute for a given level of technological deployment. For example, to determine the footprint-target function in the reformed CAFE, government officials used engineering information to estimate the footprint and fuel economy of a fleet of vehicles that used a suite of technological improvements deemed to be feasible in the near future. They then fit a line (minimizing the median absolute deviation) to the data to determine how quickly fuel economy declined as footprint rose. They used this estimated slope as the slope of the attribute-based target.

A policy thus constructed eliminates one of the potential margins of adjustment that automakers can use to improve fuel economy, that of downsizing vehicles. The government appears to recognize this: "there should be no significant effect on the relative distribution of different vehicle sizes in the fleet, which means that consumers will still be able to purchase the size of vehicle that meets their needs" (Federal Register Vol. 75, No. 88, p. 25338). This appears to have been interpreted as a desirable feature of the policy, perhaps because it was deemed important to ensure that fullsize vehicles continued to be affordable. Or, it was perceived to be fair to require improvements from all vehicles; as the government states, the footprint function implies that "all vehicles, whether smaller or larger, must make improvements" (Federal Register Vol. 75, No. 88, p. 25338). To understand the welfare impacts of this design, we examine incentives when policymakers draw the attribute slope so that it it mimics the efficiency versus attribute cost trade-off.

In Japan, the formulation is somewhat different, but it has similar implications. When considering a tightening of fuel economy standards, the Japanese government defines weight categories. Within each weight category it identifies the vehicle with the highest fuel economy. It then determines the required fuel economy rating for vehicles in that weight category in some future target year by multiplying the fuel economy of this "front runner" by some factor, requiring a percentage improvement over the leader in the current period. On average, this will draw an attribute-based slope that is closely related to an isocost curve.

Second, the attributes that determine the efficiency target are sometimes themselves associated with unpriced externalities. In Japan and China, for example, the attribute used to determine targets for automobiles is vehicle weight. In a collision, heavier vehicle are safer for their occupants, but they increase the risk of injury or death to those outside the vehicle, be they pedestrians or occupants of another vehicle. Safety to occupants is a private good that should be internalized in vehicle prices, but the risk posed to others is an unpriced externality. To better understand the implications of policy design, we examine the welfare implications of attribute-basing when the attribute that determines the target generates a separate externality.

Third, attribute-based regulation commonly features notches. For example, in Japan and China, the fuel economy targets are discrete (step) functions of vehicle weight categories. Within a category, a marginal increase in weight has no impact on the fuel economy target, but for vehicles close to a cut-off point, a marginal increase in weight can lead to a discrete drop in the fuel economy required. In our theoretical section, we consider how notched versions of the policy affect welfare and what the endogenous response to notches reveals. In the empirical section, we use the notches as a key identification strategy. In Japan, there is significant bunching around weight notches. This reveals the automakers' response to attribute-basing in a more transparent way than is possible in the U.S. or other contexts where the attribute-slope is smooth.

3 Theoretical framework

The theoretical portion of this paper explores the use of attribute-based regulations for correcting externalities. We want to establish how attribute-basing helps or hinders a policymaker who aims to correct market failures stemming from externalities. In particular, we wish to compare attribute-basing to a standard Pigouvian tax. To facilitate our analysis, we make several modeling assumptions—about market structure, product differentiation, the nature of the externality in question and how regulatory and tax policies relate. Before proceeding to the model, we explain these key assumptions and their implications.

3.1 Key assumptions used in the model

3.1.1 We model a tax instead of a regulation

A majority of the attribute-based policies in the real world are regulations, but we wish to study them through the lens of optimal tax policy in order to facilitate comparison to Pigouvian taxes. Thus, our first step is to demonstrate that the regulatory policies that we wish to study can be recast as equivalent tax policies under certain conditions.

To show this, we first specify the regulation in the standard way. Consider a firm that produces goods, indexed by $j = 1, ..., J$ with energy consumption g_j, the sales-weighted average of which is required to exceed a standard σ. Firms can choose each good's price p_j, energy consumption g_j and some other attribute a_j. The firm solves a Bertrand competition problem over all of the products in its portfolio:

$$\max_{g_j, a_j, p_j} \pi = \sum_{j \in J} (p_j - c_j) q_j$$
$$\text{s.t.} \sum_{j \in J} \frac{q_j}{Q} g_j \leq \sigma,$$

where q_j is the quantity sold, which is a function of g, a and p for all vehicles sold in the market; c_j is the marginal cost of vehicle j, which is a function of g_j and a_j but assumed to be constant in q_j, and $Q \equiv \sum_{j \in J} q_j$ is the total sales of the producer across all models. The Lagrangean for this problem is written:

$$\max_{g_j, a_j, p_j} \mathcal{L} = \sum_{j \in J} (p_j - c_j) q_j + \lambda Q \left(\sigma - \sum_{j \in J} \frac{q_j}{Q} g_j \right) = \sum_{j \in J} (p_j - c_j) q_j - \lambda \sum_{j \in J} q_j g_j + \lambda Q \sigma \qquad (1)$$

where we multiply the constraint by Q so λ is interpreted as the per-unit shadow cost of tightening the regulatory constraint. This is the standard representation of a fleet efficiency standard.

Alternatively, suppose that there was a tax (or subsidy) equal to $t(g - k)$ per unit sold. That is, there is a marginal tax on fuel consumption g, but we allow also that this marginal tax is shifted uniformly by k so that the net tax on a vehicle could be positive or negative. The profit maximization problem for the firm in this case is:

$$\max_{g_j,a_j,p_j} \pi = \sum_{j \in J}(p_j - c_j)q_j - \sum_{j \in J} q_j(tg_j - k) = \sum_{j \in J}(p_j - c_j)q_j - t\sum_{j \in J} q_j g_j + tQk. \tag{2}$$

If $t = \lambda$ and $k = \sigma$, so that the tax on a product j is $\lambda(g_j - \sigma)$, then the maximization problem for the regulation (equation 1) and the tax problem (equation 2) are identical. Thus, the non-attribute-based regulation is identical to a tax on g.

In an attribute-based regulation, the target efficiency level of each vehicle is a function of the attribute, which we denote as $\sigma(a)$. These standard that the firm must meet is the average target of its products, which is equivalent to saying that the firm's products must exceed their targets, on average. The Langragean under attribute-based regulation can be written as:

$$\begin{aligned}\max_{g_j,a_j,p_j} \mathcal{L} &= \sum_{j \in J}(p_j - c_j)q_j + \lambda Q \left(- \sum_{j \in J} \frac{q_j}{Q} \left(g_j - \sigma(a)\right) \right) \\ &= \sum_{j \in J}(p_j - c_j)q_j - \lambda \left(\sum_{j \in J} q_j \left(g_j - \sigma(a)\right) \right). \end{aligned} \tag{3}$$

Alternatively, the attribute-based tax policy would be one in which each vehicle faced a tax rate of $t(g - \omega(a))$. In this case, the optimization problem for the firm is:

$$\max_{g_j,a_j,p_j} \pi = \sum_{j \in J}(p_j - c_j)q_j - t\sum_{j \in J} q_j(g_j - \omega(a)). \tag{4}$$

It is immediately apparent then that, when $t = \lambda$ and $\omega(a) = \sigma(a)$, the attributed-based regulation (equation 3) and tax (equation 4) problems are equivalent. We will use this equivalence to justify an examination of tax policies, which facilitates comparison between Pigouvian taxes and maps into a perfectly competitive framework more naturally than regulation.

We specified here the problem faced by a single firm. With multiple firms, a common policy function $\sigma(a)$ will correspond to different shadow prices λ, and therefore a different equivalent tax policy. In fully tradable systems (such as CAFE), in which a firm that exceeds the standard can sell its excess credits to another firm, a single shadow price will prevail for all firms, provided that firms do not fail to trade for competitive reasons or because of illiquid markets or transaction costs. At times, our tax model is therefore implicitly assuming that there is full tradability and no competitive considerations prevent efficient exchange of credits across firms.

3.1.2 We assume perfect competition

In a market with imperfect competition, attribute-based regulation may have welfare effects through its impact on mark-ups, to the extent that it shifts product placement and induces more (or less) competition in market segments characterized by more or less elastic demand. This is an area well worth exploring, but we believe it is a distraction from the necessary first step of understanding the welfare impacts related directly to the externality itself (which is the professed goal of the policies we consider), rather than secondary effects the policy has on market failures related to competition. We thus assume perfect competition on the supply side.

Specifically, we posit a total cost of production function for model j, $TC(a_j, e_j, q_j)$, which has constant marginal cost per unit q_j and no fixed cost, so that $TC(a_j, e_j, q_j) = C(a_j, e_j) \cdot q_j$. Constant returns to scale (in q_j) and the absence of fixed costs implies that if a firm charges a price over $C(a_j, e_j)$ for a good j, another firm will enter and price below. Thus, in equilibrium, consumers will be able to buy any bundle of attributes at a price equal to marginal cost.

This assumption also implies that the firm's portfolio of vehicles is irrelevant to their pricing considerations. The standard way of modeling firms with multiple products under imperfect competition implies that firms must worry about how increasing some attribute of one of their products will affect the demand for all of their other products. Our assumption about perfect competition abstracts from those concerns. This has the further implication that, even if a given firm produces several goods, we could study their decision about each one of them in isolation and arrive at identical conclusions because there are no cross-product demand effects on profit. This enables us to study a model of product demand from the point of view of a consumer of a given preference function, and we can then simply model how the existence of policies shifts the attribute bundle that the consumer demands.

This assumption regarding perfect competition is a natural starting place, but it does obscure our ability to study one margin of adjustment in the product market, which is the shift in market share between existing products (sometimes called "mix shifting" in the literature) holding product attributes constant. In that sense, our model is best understood as a "long run" model in which no product characteristic is fixed.

3.1.3 We assume that energy efficiency directly causes an externality

We also assume that energy consumption g—or energy efficiency, which we will denote as $e = 1/g$—directly generates an externality. This is not the case in reality for energy efficiency. Rather, energy used creates an externality, and total energy use of a good is the product of its energy consumption rate g and total utilization. Utilization is an endogenous function of g, which generates a rebound effect.[1] Moreover, the policies under consideration here generally target only the new product market, and the used market may adjust in ways that also impact the externality.[2] For these reasons, it is well understood that energy efficiency policies—whether they are regulatory policies like CAFE or equivalent taxes on vehicles—represent significant departures from a first-best Pigouvian tax on the externality itself.[3]

Our focus is on the marginal efficiency impacts of attribute-basing. For that reason, we will abstract from these considerable concerns and assume that there is some externality related directly to energy efficiency e and that a tax on that feature of a product represents a first-best solution. It is possible that attribute-basing could provide flexibility in the policy that would allow it to address some of the inefficiencies of energy efficiency programs, but this is not the professed motivation for the use of attribute-basing and is not therefore our primary concern.

3.2 Baseline model

We build our baseline model by specifying a welfare function for a consumer who consumes a durable good that has two attributes, e for energy efficiency and a for some other attribute (e.g., footprint), and some other good x that represents all other consumption. We are not particularly interested in complementarities of consumption between the energy-consuming durable and the remainder of the consumer's consumption bundle, so we will assume that utility is separable in the

[1] The margin of adjustment described here is sometimes called the "direct rebound" effect. There are also other effects that work through income. See Borenstein (2013) for a recent discussion.

[2] See Jacobsen and van Benthem (2013) for an exploration of this effect for the case of automobiles.

[3] See Anderson, Parry, Sallee, and Fischer (2011) for a recent review of CAFE that describes related evidence. For the equivalence of vehicle taxes and CAFE, see Sallee (2011) and Gillingham (2013).

utility derived from x, denoted $\lambda(x)$, and the flow of services from the durable, denoted $U(a, e)$. We assume that both a and e are goods ($\partial U/\partial a > 0$ and $\partial U/\partial e > 0$) with diminishing marginal returns ($\partial^2 U/\partial^2 a < 0$ and $\partial^2 U/\partial^2 e < 0$) and that they are complements ($\partial^2 U/\partial a \partial e > 0$).

There is also an externality, denoted $\phi(E)$, where E is the sum of e over all consumers. The consumer has an exogenous amount of income I and receives a lump-sum demogrant (which can be positive or negative) from the government, G, which the government will use to recycle revenue and balance its budget. The price paid for the durable good is the price charged by the producer $P(a, e)$, which is a function of both attributes, as well as any tax (or subsidy) levied by the government, $t(a, e)$.

The consumer's problem is the following, which differs from a standard optimization problem over the three goods a, e and x only in that the price and tax paid for a and e may be related to each other:

$$\begin{aligned} \max_{a,e,x} W =& U(a, e) + \lambda(x) + \phi(E) \\ \text{s.t. } & I + G \geq P(a, e) + t(a, e) + x, \end{aligned}$$

where W denotes consumer utility from all sources and the price of x has been normalized to one, making x the numeraire. We assume that there is a continuum of identical consumers of total measure 1, so that $E = e$ but consumers do not recognize their individual impact on the total externality E when making individual choices.[4]

As discussed above, we assume perfect competition so that prices are directly determined by firm cost. Specifically, we assume that production involves no barriers to entry, there are many potential firms, and production costs involve a constant marginal cost so that total production cost is $C(a, e) \cdot q$. Under these assumptions, equilibrium prices will be $P(a, e) = C(a, e)$; there will be no producer surplus; and consumers will bear the full burden of any taxation.

We assume that $C(a, e)$ is rising and convex in both attributes. In all real world policies that we know of, the attribute a—when written as a good as opposed to a bad—is negatively related to energy efficiency e. (For cars, a is either footprint or weight, both of which correlate negatively

[4]Given our assumptions about perfect competition, constant marginal cost in q and no fixed costs, the problem we describe and the results we derive would be identical if there were many types of consumers, each of whom selected a different vehicle. Describing only one consumer type greatly simplifies notation.

with fuel economy.) Thus, we assume that $\partial^2 C/\partial a \partial e \geq 0$. This ensures the convexity of the budget set.[5]

In this framework, attribute-based policies represent a particular functional form of the tax function $t(a, e)$. As shown in section 3.1.1, a non-attribute based regulation can be modeled as a tax function $t(a, e) = t \cdot (e - k)$. For the traditional Pigouvian tax $k = 0$ and $t =$ marginal benefits (or damages). Under our assumption about quasi-linearity, the choice of a and e will be independent of k and thus it will have no effect on choice.[6] An attribute-based tax like the one imposed implicitly by footprint-based CAFE standards will be a function $t(a, e) = t \cdot (e - \sigma(a))$.

3.2.1 First-best solutions do not involve attribute-basing

The first step in our analysis is to derive the first-order conditions for the consumer facing a tax policy and compare them to the planner's first-best conditions when the planner is allowed to choose the allocation directly. We can then see what function $t(a, e)$ would make the consumer's conditions equal to the first-best. As we proceed, we make two functional form assumptions purely for notational ease. First, we assume that $\lambda(x) = \lambda \cdot x$, which means that there is a constant marginal utility from consumption equal to λ. Second, we assume that $\phi(E) = \phi \cdot E$, which means that there is a constant marginal benefit (or damage) from the externality equal to ϕ.

Under these assumptions, we now write the consumer's problem as an unconstrained optimization problem by substituting the budget constraint into the objective function:[7]

$$\max_{a,e} W = U(a, e) + \lambda \cdot [I + G - P(a, e) - t \cdot (e - \sigma(a))] + \phi \cdot E.$$

This is a standard choice problem, and the first-order conditions can be derived directly by differ-

[5]Empirical estimates in Knittel (2011) indicate that the log of fuel economy is concave down in the log of weight, which is the specific application we have in mind.

[6]When $k \neq 0$, regulatory policies differ from the Pigouvian tax in that they provide different incentives on the extensive margin. That margin is assumed away in this context, though it has been shown to be relevant in other settings, for example in Holland, Knittel, and Hughes (2009). This margin is why taxes on a product's energy consumption are not equivalent to subsidies for its energy efficiency (Metcalf 2009).

[7]Our assumptions regarding $C(a, e)$ ensure that the budget set is convex and thus that the consumer will exhaust their budget.

entiation:

$$\frac{\partial W}{\partial a} = \frac{\partial U}{\partial a} - \lambda \frac{\partial P}{\partial a} + \lambda t \sigma' = 0 \tag{5}$$

$$\frac{\partial W}{\partial e} = \frac{\partial U}{\partial e} - \lambda \frac{\partial P}{\partial e} - \lambda t = 0, \tag{6}$$

where σ' is the derivative of $\sigma(a)$ with respect to a.

First-order conditions for the first-best allocation can be found by differentiating the planner's problem, which is:

$$\max_{a,e} W = U(a,e) + \lambda \cdot [I - C(a,e)] + \phi \cdot e.$$

This differs from the consumer's problem in directly inserting the cost function in place of prices (though these will be equal under perfect competition) and taxes, and in recognizing how the choice of e affects the total externality. The first-best optimization conditions are found by differentiation:

$$\frac{\partial W^*}{\partial a} = \frac{\partial U}{\partial a} - \lambda \frac{\partial C}{\partial a} = 0 \tag{7}$$

$$\frac{\partial W^*}{\partial e} = \frac{\partial U}{\partial e} - \lambda \frac{\partial C}{\partial e} + \phi = 0. \tag{8}$$

Under perfect competition, prices reflect unit costs, so $\frac{\partial P}{\partial a} = \frac{\partial C}{\partial a}$ and $\frac{\partial P}{\partial e} = \frac{\partial C}{\partial e}$. Then, it is immediately apparent that the only way to make the consumer's first-order conditions equal to the first-best (that is, to make equation 5=equation 7 and equation 6=equation 8) is to set $\sigma' = 0$ and $t = -\phi/\lambda$. That is, to have a Pigouvian subsidy on e with no attribute-basing, where dividing by the marginal utility of income transforms the externality in utils into the externality in dollars.

This is the first result of our theoretical analysis. The result is not surprising, but it is relevant for evaluating policy, as it establishes a set of conditions under which attribute-basing is counter-productive. If the only market failure is an externality associated with e, then the optimal policy will not involve attribute-basing. The optimal attribute-based slope is uniquely zero. If there is a role for attribute-basing, it must be due to some additional factor, such as a technology-spillover or a constraint on the choice of t, which we discuss below.

3.2.2 The first-order deadweight loss from attribute-basing is the distortion in a

Suppose that the true welfare model is as specified above and that the tax rate is set at the Pigouvian level, $t = -\phi/\lambda$. The optimal attribute slope is zero, but suppose the policy is attribute-based, $\sigma' \neq 0$, for some outside reason, such as political economy or simple confusion on the part of policymakers. What distortions does attribute-basing cause in this setting?

A qualitative characterization of distortions can be made by examining the first-order conditions of the consumer, again comparing this to the first-best. Let a^* and e^* denote the first-best choices. Let a' denote the value of a that is chosen under the attribute standard. Then the following are the first-order conditions for e from the two problems:

$$\begin{aligned} \text{First-best:} \quad & \frac{\partial U(a^*, e^*)}{\partial e} + \phi = \lambda \frac{\partial P(a^*, e^*)}{\partial e} \\ \text{Consumer's choice:} \quad & \frac{\partial U(a', e')}{\partial e} - \lambda t = \lambda \frac{\partial P(a', e')}{\partial e}. \end{aligned}$$

If $t = -\phi/\lambda$, as would be the case if policy-makers used the Pigouvian tax level in the attribute-based problem, then the consumer's first-order condition is exactly the same as the first-best, except that the level of a' is distorted. This creates a *second-order* distortion in e.

This is illustrated in Figure 2a, which shows the consumer's choice of e by plotting the marginal cost and marginal benefit of (willingness to pay for) e, conditional on $a = a^*$. (We do not need to condition on x because of the quasi-linearity of $U(a, e)$.) The marginal cost of increasing e to the consumer is the derivative of the cost function with respect to e, given a level of a. The private marginal benefit of increasing e is the derivative of U with respect to e, for a given level of a. The solid black lines plot hypothetical curves for these functions when $a = a^*$. The corrective tax rate t shifts the willingness to pay curve up by amount t. Thus, the consumer will choose e^* when there is no attribute-basing and the tax rate is t.

Suppose that, under attribute-basing, the consumer's choice of a, call it a', exceeds a^*. In the diagram, this will shift both the marginal cost and the marginal benefit curves up. The marginal cost curve will shift up by $\approx (a' - a^*)\frac{\partial^2 P(a^*,e^*)}{\partial a \partial e}$. The willingness to pay curve will shift up by $\approx (a' - a^*)\lambda^{-1}\frac{\partial^2 U(a^*,e^*)}{\partial a \partial e}$. If $\frac{\partial^2 P(a^*,e^*)}{\partial a \partial e} = \lambda^{-1}\frac{\partial^2 U(a^*,e^*)}{\partial a \partial e}$ (which would occur only by coincidence) then $e^* = e'$ and attribute-basing would cause no change in the allocation of e. In general, which shift

Figure 2: Distortions in the choice of a and e from attribute-basing

(a) Choice of e (second-order distortion) **(b)** Choice of a (first-order distortion)

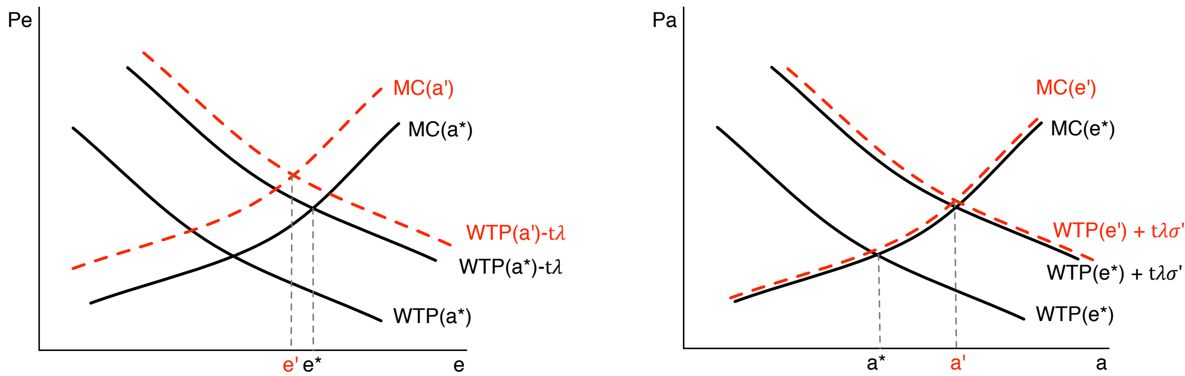

is larger will determine whether the choice of e falls or rises relative to the first-best, but in either case the change will be the result of two offsetting second-order shifts.

The same is not true in the analysis of a. The first-order conditions with respect to a are:

$$\begin{aligned} \text{First-best:} \quad & \frac{\partial U(a^*, e^*)}{\partial a} = \lambda \frac{\partial P(a^*, e^*)}{\partial a} \\ \text{Consumer's choice:} \quad & \frac{\partial U(a', e')}{\partial a} + \underbrace{\lambda t \sigma'}_{>0} = \lambda \frac{\partial P(a', e')}{\partial a}. \end{aligned}$$

When $\sigma' \neq 0$, attribute-basing introduces a wedge between the marginal benefit and marginal cost functions for a. The difference between e^* and e' will cause a change in the choice of a, but this will be a second-order effect, of magnitude $\approx (e' - e^*) \left(\frac{\partial^2 P(a^*, e^*)}{\partial a \partial e} - \lambda^{-1} \frac{\partial^2 U(a^*, e^*)}{\partial a \partial e} \right)$, where we expect that $e' - e^*$ will be small. The direct effect of attribute-based standards is in creating the wedge $-\lambda t \sigma'$, which equals $\phi \sigma'$ when t is set at the Pigouvian level. This wedge will cause the consumer to choose too high of a value of a, and the distortion will be rising as σ gets steeper (σ' gets more negative). The wedge is also larger (in absolute value) as t gets larger (or as ϕ/λ gets larger assuming the Pigouvian rate for t is used).

This is shown in Figure 2b, which plots the marginal cost and benefit of a, conditional on a given level of e. The introduction of the attribute-based wedge shifts the benefit curve out linearly with σ'. This creates what would be labeled as a Harberger triangle in a standard partial equilibrium graph, showing the lost utility from choosing the too large value of a'.

This analysis has important implications for empirical studies of attribute-based policies. Take CAFE for example. The introduction of attribute-basing may have only a very limited impact on fuel economy as compared to a policy with the same implicit subsidy for fuel economy that has no attribute-basing. Fuel economy might be higher or lower under the attribute-based policy than under the shadow-price equivalent attribute-free policy. The first-order distortions and welfare losses of the footprint-based standards are associated with the distortion in the footprint, not the distortion in fuel economy. The size of the externality enters the distortion formulas because it determines the size of the tax wedge that distorts the attribute, not because the primary welfare loss is from too little externality reduction. Altogether, this suggests that the place to look for distortions in the data is in the distribution of the attribute upon which standards are based, not in the distribution of energy efficiency.

3.2.3 Mispricing of externality does not justify attribute basing

The analysis above focuses on the first-best allocation, in which attribute-basing has no role to play. But, might there be a second-best argument that makes attribute-basing useful? It is a standard intuition in optimal tax design that it is efficient to introduce a small tax in one market that is not distorted in order to reduce a pre-existing distortion in a second market because the initial distortion creates small welfare losses compared to marginal improvements in pre-existing distortions.

To see whether this intuition applies to attribute-basing, suppose that the optimal tax rate t is set at a rate too low, so that the e is below the social optimum and the policymaker cannot increase it (perhaps due to political constraints) by changing t. The externality presents a first-order pre-existing distortion in the market for e, so the net marginal benefit of increasing e from the private optimum is $\approx \Delta e\phi/\lambda$. Introducing an attribute slope will create a distortion in the market for a, but the starting point is efficient. The Harberger triangle in a is $\approx 1/2\Delta a(\sigma')^2$, which is ≈ 0 if the starting point is $\sigma' = 0$.

This logic breaks down, however, because the change in e induced by a change in a is itself second order. The attribute slope does not have a direct impact on the choice of e, but rather shifts both the marginal cost and willingness to pay curves for e by amounts proportional to the cross-partial (second order) derivatives of the utility and cost functions. Thus, the "small" distortion created in the a market will correspond to a "small" improvement in the e market because, even though changes in e create "large" welfare gains, e only changes by a "small" amount in response to a. We are currently developing a formal demonstration of this intuition.[8]

There are many real world policies in which it appears that policymakers are constrained by political factors from setting regulations or taxes at the right level. It is tempting to think that the introduction of attribute-basing into regulations would give the policymakers greater flexibility and therefore help circumvent constraints on setting the right level. This is unlikely to be the case, however, because attribute-basing is a very unwieldy tool for affecting the choice of the product characteristic that creates the externality; it primarily acts as an implicit subsidy for the attribute itself, not the externality generating characteristic.

3.2.4 Distortions in a are more consequential if a has an unpriced externality

In some real world attribute-based policies, the attribute upon which the target is based itself generates externalities. Take the example of Japanese fuel economy standards, in which case e is fuel economy and a is vehicle weight. Heavier vehicles pose greater risk to drivers of other vehicles in the case of a collision. This risk is unpriced (whereas the increased safety of those in the heavier car is a private benefit and therefore should be priced in the market) and represents a negative externality associated with weight.

The model can accommodate this possibility by adding a second externality to the welfare function equal to $\omega \cdot a$, so $W = U(a, e) + \lambda x + \phi e - \omega a$, with $\omega > 0$. To see the implications of this for the distortion related to the use of attribute-based standards in the baseline model, we can consider an approximation of the welfare loss due to a movement of choices from the Pigouvian benchmark (the choice when $t = -\phi/\lambda$ and $\sigma' = 0$), denoted (a^*, e^*), to that which is chosen by the consumer under the attribute-based policy, denoted (a', e'). The first-order Taylor expansion

[8] A related, but distinct, question is whether or not a policymaker should "exchange" the introduction of an attribute slope for an increase in t towards the optimum. This characterizes the move to footprint-based CAFE policies, in which the Detroit Three automakers agreed to go along with a tightening of standards as long as attribute-basing, which would place a greater burden on Japanese competitors relative to a flat standard, was introduced. The answer to this obviously depends on the size of the initial mispricing of t versus the steepness of the attribute basing that is required.

for this welfare loss, when a has an externality of ω, is:

$$\begin{aligned} W(a', e') - W(a^*, e^*) =& (a' - a^*) \left(\frac{\partial U(a^*, e^*)}{\partial a} - \lambda \frac{\partial P(a^*, e^*)}{\partial a} - \omega \right) \\ &+ (e' - e^*) \left(\frac{\partial U(a^*, e^*)}{\partial e} - \lambda \frac{\partial P(a^*, e^*)}{\partial e} + \phi \right) \\ =& - (a' - a^*)\omega. \end{aligned} \tag{9}$$

Note that $\frac{\partial U(a^*,e^*)}{\partial a} = \lambda \frac{\partial P(a^*,e^*)}{\partial a}$ and $\frac{\partial U(a^*,e^*)}{\partial e} = \lambda \frac{\partial P(a^*,e^*)}{\partial e} - \phi$, and several components therefore cancel. (This is just an application of the envelope theorem.) What remains is that, to a first-order approximation, the welfare loss from attribute-basing will be equal to the change in a times the externality associated with a.

Intuitively, what this tells us is that, when the attribute a has an unpriced externality associated with it, the welfare loss from the introduction of a small attribute slope will be dominated by the externality associated with a. This is relevant for our empirical evaluation, where we use this reasoning to justify a focus on how Japanese fuel economy standards cause an increase in average vehicle weight.

Conversely, if the policymaker were aware of the externality associated with a, the attribute-based policy could be designed to correct both externalities by setting $t = -\phi/\lambda$ and $\sigma' = \omega/\phi > 0$. That is, if there are two externalities, then the attribute-based standard can affect a Pigouvian tax on each. This would imply an attribute slope which has the opposite sign as what we observe in reality, suggesting that policymakers are not using the degree of freedom in the policy to correct the related externality.

3.2.5 Real world attribute slopes preserve relative prices of a and e

As discussed in section 2, some policymakers have explicitly determined the attribute slope to match estimates of the cost frontier. That is, they have picked σ' so that the net subsidy $t(e - \sigma(a))$ is constant along an isocost curve, so $\sigma' = \frac{\partial C/\partial e}{\partial C/\partial a}$. In terms of first-order conditions, this preserves the relative prices of a and e that prevail when there is no subsidy at all. This eliminates the incentive to divert expenditure on a towards e that is induced by a Pigouvian tax on e. Instead, the relative prices of a and e are both distorted in equal amounts relative to the numeraire x.

Figure 3: Sample isocost curve under Pigouvian and attribute-based policies

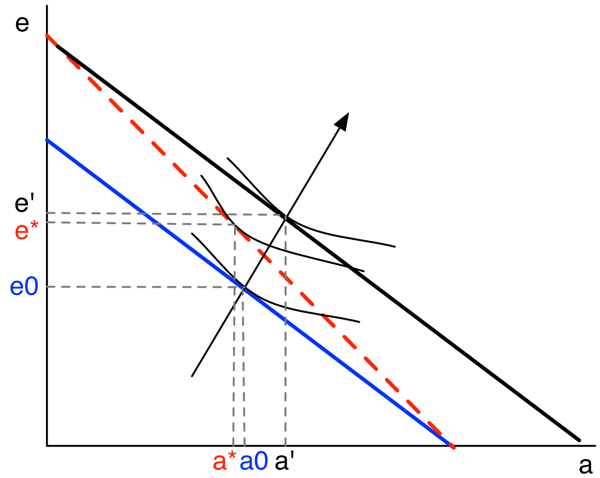

This can be shown visually by plotting an isocost curve, along which a consumer spends a constant amount on the vehicle $P(a, e) + t(a, e)$. Figure 3 plots such a cost curve. When there is no subsidy at all, the isocost curve (isoconsumption of x curve) is the innermost solid blue line, and the consumer will choose the point labeled $(a0, e0)$. The red dashed line represents the same isocost curve when there is a Pigouvian subsidy $t \cdot e$. This lowers the cost of e but has no impact on the cost of a, though it may alter the choice of a depending on the shape of the utility and cost curves. (Note that cost curves are drawn here to be linear, which aids illustration. Generally, the cost curves will bend outwards.) Under this subsidy, the consumer will choose (a^*, e^*).

Finally, the outer black solid line represents the isocost curve for a consumer facing a tax $t \cdot (e - \sigma(a))$, when $\sigma' = \frac{\partial C/\partial e}{\partial C/\partial a}$. This shifts the final budget set to be parallel to the original by making the relative price of a and e the same as in the no-subsidy case. Both are now cheaper relative to x, however, which pushes the consumer to choose a higher bundle of both a and e that represent a greater (social) expenditure on the vehicle relative to x.

The arrow drawn represents the expansion path of a and e for the original relative price curve; that is, the arrow shows how the bundles of (a, e) evolve as total expenditure $P(a, e)$ grows but relative marginal prices of a and e are held constant. The original and final points are on that curve, but the Pigouvian allocation (a^*, e^*) is not.[9] A Pigouvian tax gives the consumer an incentive to

[9] This same result can be shown from the first-order conditions from the consumer's problem, in which setting $\sigma' = \frac{\partial C/\partial e}{\partial C/\partial a}$ causes the ratio of marginal utilities of a and e to equal the unsubsidized ratio of marginal costs.

"downsize" by picking a lower value of a in order to boost e. When attribute standards are drawn to mimic the cost frontier, as has been explicitly done in several automobile policies, this incentive is eliminated. All improvements in e come from increases in the overall cost of the good (movements up the technology frontier) in a way that is consistent with the original expansion path. To the extent that downsizing is a significant component of the optimal product adjustment in response to standards, attribute-basing will have a large welfare impact by muting that margin of adjustment.

3.3 Technology spillovers create a role for attribute-basing, but not the one given to it by policymakers

Attribute based standards eliminate (or at least limit) the ability of producers to improve energy efficiency by downsizing in a, instead requiring them to make improvements in the products by dialing down attributes other than a that are correlated with e, or by adding technology that improves e while holding a constant. This approach could be appealing if there are positive externalities related to technological advancement. If firms cannot fully appropriate the value of their technological improvements, then technological improvements may represent a spillover benefit to other firms. In this case, there will be too little technological innovation in the absence of policy, and attribute standards might be beneficial in spurring innovation (instead of downsizing).

A direct way to model technological spillovers is to allow some secondary externality to directly enter the welfare function. We develop here a very reduced form model of these spillovers by simply specifying that costs incurred over some reference level, meant to represent the current state of the art or the current production possibility frontier, induce a positive external benefit. We are developing a more complete model of this spillover process, but we strongly suspect that the basic insights of the simplistic model represented here will carry forward in a more rigorous framework.

Denote the frontier cost function (some reference cost level) as $C(\bar{a}, \bar{e}) \equiv \bar{C}$. When the firm chooses a bundle inside of the frontier, that is $C(a, e) \leq C(\bar{a}, \bar{e})$, then the firm creates no technology spillover. The idea in this case is that the good is being produced with already established techniques. If the firm must push costs beyond the current frontier to make a product with a "higher" bundle of a and e, this requires some additional investment that costs $\rho(C(a, e) - C(\bar{a}, \bar{e}))$. We make the functional form assumption that this extra cost is linear and increasing in its argument, so that the extra cost can be written as $\rho(a,e) = \rho \cdot \max\{C(a,e) - C(\bar{a},\bar{e}), 0\}$ and $\rho > 0$. Conversely, there is some externality that benefits all of society from the technological advancement, which we write as $\gamma(a,e) = \gamma \cdot \max\{C(a,e) - C(\bar{a},\bar{e}), 0)\}$

The consumer's and planner's problems can be written to include this spillover. There will be first-order conditions for each of two cases, where $C(a,e) > C(\bar{a},\bar{e})$ and where $C(a,e) \leq C(\bar{a},\bar{e})$. In either case, the consumer will fail to recognize the externality γ and will recognize only the costs. The only difference in this problem from the baseline model is that there will be a kink in the cost function at $C(\bar{a},\bar{e})$ which will induce "bunching"—not at a particular level of a or e, but along a particular isocost curve of a and e.

When the first-best choice involves values of a and e below $\bar{C}$, then the consumer's and planner's problems will be identical to the baseline. It is only when the first best allocation lies above $\bar{C}$ that this problem yields different conditions, so we focus on that here. In this case, the planner's problem becomes:

$$\max_{a,e} W = U(a,e) + \lambda \cdot [I - C(a,e) - \rho \cdot (C(a,e) - C(\bar{a},\bar{e}))] + \phi \cdot e + \gamma \cdot (C(a,e) - C(\bar{a},\bar{e})).$$

The planner's first-order conditions of this problem are:

$$\begin{aligned}\frac{\partial W}{\partial a} &= \frac{\partial U}{\partial a} - (\lambda + \lambda\rho - \gamma)\frac{\partial C}{\partial a} = 0 \\ \frac{\partial W}{\partial e} &= \frac{\partial U}{\partial e} - (\lambda + \lambda\rho - \gamma)\frac{\partial C}{\partial a} + \phi = 0.\end{aligned}$$

The consumer's first-order conditions are:

$$\begin{aligned}\frac{\partial W}{\partial a} &= \frac{\partial U}{\partial a} - (\lambda + \lambda\rho)\frac{\partial C}{\partial a} + \lambda t\sigma' = 0 \\ \frac{\partial W}{\partial e} &= \frac{\partial U}{\partial e} - (\lambda + \lambda\rho)\frac{\partial C}{\partial e} - \lambda t = 0.\end{aligned}$$

To make the consumer's first-order condition for e equal to the planner's, the tax on e must be set at $t = -\lambda^{-1}(\phi + \gamma\frac{\partial C}{\partial e})$. This again is directly consistent with the Pigouvian tax result, where now the total externality related to e includes both the energy externality ϕ and the spillover externality γ. If $t = -\lambda^{-1}(\phi + \gamma\frac{\partial C}{\partial e})$, then to make the consumer's first-order condition for a equal to the planner's, the attribute slope must be:

$$\sigma' = -\frac{\gamma \frac{\partial C}{\partial a}}{\phi + \gamma \frac{\partial C}{\partial e}}. \tag{10}$$

This result indicates that attribute-basing does have a potential role in the first-best solution when their are technology spillovers. Specifically, at the first-best, the slope of the attribute function should be equal to the ratio of the externalities associated with a and the externalities associated with e.

As we argued in section 2, policymakers appear to commonly choose attribute-based schedules that follow the production frontier. That is, they choose a slope of $\frac{\partial C}{\partial a}/\frac{\partial C}{\partial e}$. This is optimal only if there is no externality related to e, and instead there is only a technological spillover, in which case $\phi = 0$ and equation 10 collapses to be only the ratio of marginal costs of the two attributes. More generally, the larger is the technology spillover γ relative to the energy externality ϕ, the closer will be the attribute slope to the production ratio. In contrast, as the energy externality is greater relative to the technology spillover ($\phi >> \gamma$), the attribute slope will approach zero.

Importantly, this result does not imply that attribute-basing is essential. Instead, attribute-basing provides a second instrument in addition to the tax on e. With two instruments and two externalities (one for e and one for a) it is possible to achieve the first best. That the two instruments are related via attribute-basing is not at all essential. The first best could equally well be achieved by subsidizing e at rate $\lambda^{-1}(\phi + \gamma \frac{\partial C}{\partial e})$ and subsidizing a at rate $\lambda^{-1}(\gamma \frac{\partial C}{\partial a})$.

3.4 There is a role for attribute-basing when the standard applies to each good

Up to now we have focused on the case when regulation requires firms, or the industry as a whole when trading is allowed, to meet an average of some parameter e across a variety of products. This characterizes a number of regulations, including fuel economy standards, low-carbon fuel standards, and cap-and-trade systems, as well as policies that are implemented as taxes, including automobile feebates or carbon taxes.

Some attribute-based regulations apply to each and every product in a market, however. Key examples are appliance product standards, which specify a minimum energy efficiency that all products are required to meet. Refrigerators in the U.S., for example, must meet a minimum energy efficiency that is a continuous function of its "adjusted volume" (equal to fresh volume plus 1.63 times freezer volume) and a discontinuous function of features, including door type (french or not), location of freezer (top, bottom or side) and whether there is through the door ice. If all refrigerators were required to meet the same minimum, the minimum might affect the total prohibition of some models. We thus conjecture that, in this setting, attribute-based standards can be welfare improving compared to a common standard.

We are currently building a version of our model that better accommodates this scenario. Our preliminary work suggests that in this setting, there is a role to be played by attribute-based standards. The optimal slope of the attribute function, however, is not simply a reflection of technological costs (which appears to be what policymakers traditionally have used) but instead is a function of the technology cost curve relative to the utility function and the size of the externality.

3.5 Notched policies have similar welfare properties

Many of the attribute-based regulations that we consider feature notches. That is, the function relating the fuel economy target to the attribute, $\sigma(a)$, is a step function with discrete jumps. Sallee and Slemrod (2012) study notched fuel economy policies and conclude that the use of notches in place of a smooth Pigouvian tax is welfare decreasing because it provides overly large incentives for fuel economy improvements for some agents and too small incentives for others. Our context is different in that we have two dimensions, and either could have notches.

We first consider the case where the notches exist only in the attribute function σ, but the tax on e is smooth. That is, conditional on e, the tax policy function $t(a, e)$ takes discrete jumps at certain values of a. How would such a policy function affect consumer choice? We can get some initial intuition graphically. Figure 4 shows an isocost curve, that is, the set of values for which a consumer spends a constant amount on the durable, $P(a, e) + t(a, e)$. The figure is drawn with several notches, at a', a'' and a'''. As in Figure 3, the solid blue line (drawn to be linear for the sake of illustration) shows the isocost curve before any policy intervention; and the dashed red line shows the modified isocost curve for the same expenditure on the good when there is a Pigouvian subsidy on e that has no attribute slope.

Next, the dashed grey line represents the isocost curve that would exist under a smooth attribute policy. In the diagram, the grey line is drawn parallel to the original blue line, which represents

Figure 4: Isocost curve with notched attribute-based tax

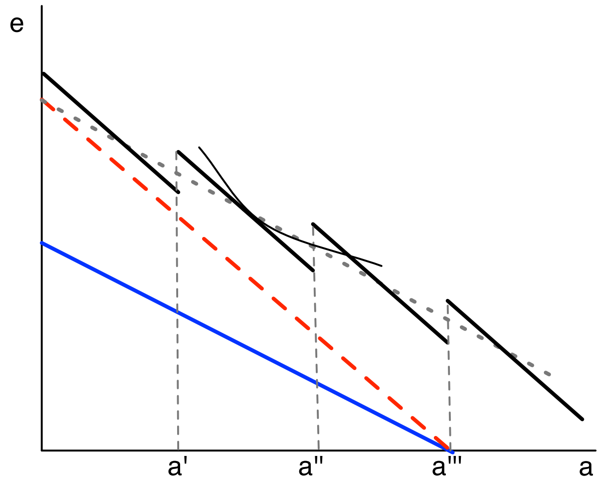

the case when policy makers draw the attribute slope to match existing isocost curves, thereby preserving the original relative prices of a and e. This grey dashed line is not the final isocost curve, however, when $\sigma(a)$ is notched. In that case, the solid black lines represent the isocost curve for the consumer.

Importantly, the line segments on the final isocost curve are parallel to the red dashed line representing the Pigouvian tax. As in the smooth case, the existence of the attribute function does not distort the price of a relative to x, which means that the distortion in the choice of e will be second-order—it will be driven only by the utility and cost interactions of the optimal choice of e and the distorted choice of a. Furthermore, because the line segments are parallel in slope to the original Pigouvian line (and because we assume quasi-linearity) the choice of a will not be distorted if the consumer is choosing an interior point along one of the line segments. All of the distortion is due to cases where a consumer chooses a'. That is, all of the distortion is evident from those who "bunch" at the notch points. This guides our empirical work below.

We now provide algebraic analysis to flesh out the graphical intuition. For notational ease, we focus on the case with only one notch, at a', above which the tax subsidy jumps by amount $\tau > 0$. Then, the tax function can be written as:

$$t(a,e) = \begin{cases} t \cdot e & \text{if } a < a' \\ t \cdot e + \tau & \text{if } a \geq a'. \end{cases} \tag{11}$$

As above, denote by (a^*, e^*) the bundle chosen by a consumer facing a Pigouvian tax of $t \cdot e$. If the consumer's choice under the smooth attribute policy had $a^* > a'$, then the addition of the notch τ is purely an income effect. It has not changed the marginal price of a or e relative to each other or relative to x. Given quasi-linearity, this means that the durable choice of a consumer with $a^* > a'$ is unaffected by the introduction of a notched attribute policy.

When $a^* < a'$, the consumer will face a discrete choice of maintaining their original allocation or switching to a' exactly. They will not choose $a > a'$. To see why, suppose that they chose a value under the notched policy, call it $\tilde{a}$ strictly greater than a'. Then their optimization problem can be written $\mathcal{L} = U(a,e) + \lambda[I + G - P(a,e) + t \cdot e + \tau] + \mu[a - a']$, where there is a budget constraint as well as an inequality constraint that $a \geq a'$. If $\tilde{a} > a'$, then the shadow price on the latter constraint, μ, is zero. In that case, the first-order conditions of the problem will be exactly the same as in the benchmark case with no attribute notch, which by construction featured an optimal choice of $a^* < a'$.

Thus, the consumer with $a^* < a'$ will either choose $\tilde{a} = a^*$ (and not receive τ) or will choose $\tilde{a} = a'$ exactly. This has the empirical implication that all bunching should come "from the left"—changes in a in response to the notched incentives should always be *increases* in a.

If a consumer chooses a', then their choice of e will solve:

$$\max_{e} = U(a',e) + \lambda[I + G - P(a',e) - t \cdot e - \tau], \tag{12}$$

which has the first order condition $\frac{\partial U(a',e)}{\partial e} = \lambda \frac{\partial P(a',e)}{\partial e} + \lambda t$. This compares to $\frac{\partial U(a^*,e)}{\partial e} = \lambda \frac{\partial P(a^*,e)}{\partial e} + \lambda t$ for the Pigouvian case, where $a' > a*$. Thus, just as in the smooth case, the distortion in e depends only the relative second-order curvature of P and U, which will be working in offsetting directions. If they offset perfectly, then $\tilde{e} = e^*$, but this would owe to a coincidence regarding the shape of P and U. More generally, the consumer's choice of e could rise or fall relative to the first-best choice, but this distortion will be second-order, operating through an interaction with the distortion in a.

The distortion in a will be first-order and rising in τ. The consumer will choose $\tilde{a} = a'$ if and only if:

$$-\tau > P(a', \tilde{e}) - P(a^*, e^*) - \lambda^{-1}(U(a', \tilde{e}) - U(a^*, e^*)), \tag{13}$$

that is, whenever the tax benefit is larger than the cost increase from moving from (a^*, e^*) to $(a', \tilde{e})$ minus the increase in utility from that change. The welfare loss can be written as a Taylor expansion, which again has the same intuition as a traditional Harberger triangle.

For our purposes, the main point of this analysis is that, even when the attribute function is notched, the focus of welfare analysis should be on how the policy distorts the choice of a relative to the Pigouvian baseline, and that we should expect the distortion to result in bunching at exactly the notch points in a. For empirical purposes, notched policies are useful in revealing the distortion because it is generally easier to detect bunching at specific notch points than shifts over time in an entire schedule. For that reason, the Japanese fuel economy standards are an especially fruitful context for study.

All of this applies to situations where the (implicit or explicit) tax on e is smooth, but the attribute-basing function is notched. This describes fuel economy regulations in Japan. It is also possible, however, that both the tax on e and the attribute function σ have notches. That is, $t(a, e)$ is discontinuous in both a and e. This case describes the Japanese tax subsidy programs.

The simplest version of this policy is one with a single cutoff for a, call it a' and a pair of cutoffs for e, call them e' and e''. The tax for such a system can be described algebraically as:

$$t(a, e) = \begin{cases} \tau_1 & \text{if } e > e'' \\ \tau_2 & \text{if } e'' > e > e' \text{ and } a > a' \\ 0 & \text{otherwise.} \end{cases} \tag{14}$$

An isocost curve for this case is shown in Figure 5. The unsubsidized budget constraint is drawn as a faint line. The final budget constraint is represented by the bold black line segments, which overlap in parts with the unsubsidized line. Allocations in the yellow shaded area receive some subsidy. The subsidy is equal to τ_1 for any allocation above e''. Note that there are large regions

Figure 5: Isocost curve with notches for both a and e

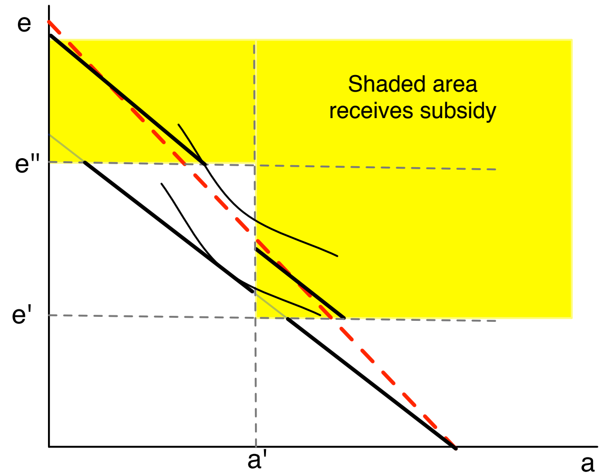

of dominance in this diagram, where a subsidized point that has more of a and more of e has the same cost to the consumer as an unsubsidized bundle.

In the diagram, the red dashed line represents the simple Pigouvian tax. The values of τ_1 and τ_2 are chosen in this case to match the average Pigouvian subsidy for the relevant line segments, but this need not be the case. Note that, if it is the case, then $\tau_1 \neq \tau_2$. In many policy examples, $\tau_1 = \tau_2$, which may be suboptimal.

When there are notches in both dimensions, there can be bunching in the distribution of e, at e' and e''. Above we argued that any change in a caused by attribute-basing relative to the Pigouvian optimum would come from *increases* in a. But, in cases with notches in both dimensions, it is possible that responses to the policy will lower a by inducing bunching at e' or e''. This would occur for cases like those represented by the sample utility curve in Figure 5, where a consumer's response to the notched subsidy is to bunch at e''. In that example, the indifference curve that is tangent to the unsubsidized budget constraint features a higher initial choice of a than at the bunch point.

4 Empirical Analysis and Results

In this section, we investigate how the market responds to attributed-based regulation by analyzing the automobile market in Japan. The Japanese context provides several key advantages. First,

Figure 6: Fuel Economy Standard in Japan

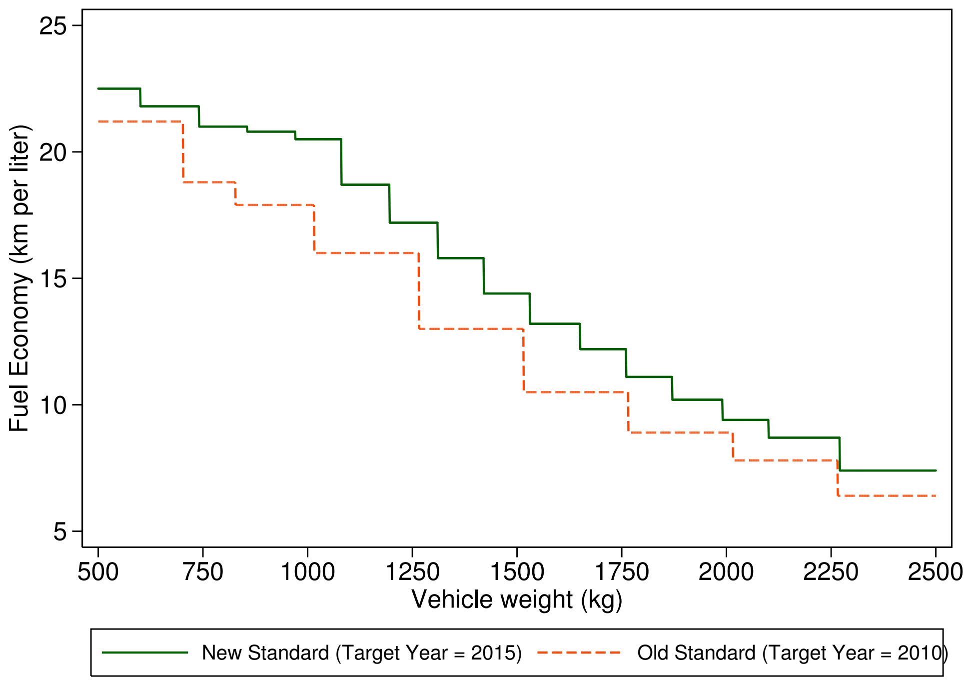

Note: The dashed line shows the fuel economy standard in Japan until 2010. The solid line shows the new fuel economy standards whose target year is 2015.

the Japanese government has been using attribute-based fuel economy regulation since the 1970's. This gives us a long window for analysis and also allows us to use several policy changes for identification. Second, the Japanese fuel economy regulations feature notches in the attribute (weight) that determines the fuel economy target. These notches provide substantial variation in regulatory incentives and allow us to use empirical methods that are similar to the recent literature on nonlinear income tax schedules (Saez 1999 2010; Chetty, Friedman, Olsen, and Pistaferri 2011; Kleven and Waseem 2013). In contrast, the U.S. instituted attribute-based regulation for automobiles only in 2012, which leaves us little data for analysis, and its attribute-based target function is smooth, making identification more challenging.

4.1 Fuel Economy Standards in Japan

The Japanese government introduced the first attribute-based fuel economy standard in 1979. Since then, there have been four different schedules. Our data, which start in 2004, come from years spanned by the two most recent regimes.[10] We summarize those two policies in Figure 6, which plots the fuel economy target for vehicles as a function of their weight. The downward slope in the figure shows that the required fuel economy is a decreasing function of weight. That is, heavier vehicles face a less stringent regulation. Furthermore, the regulation is notched. The declining function is not smooth, but rather is discontinuous at many weight cutoff points. Thus, for vehicles that have a weight near a cutoff point, marginal changes in weight lead to discrete changes in the fuel economy standard. This attribute-basing implicitly subsidizes vehicle weight, which we expect will lead to a distortion by leading to an increase in weight among firms attempting to meet the standard.

Automakers have two incentives to meet these fuel economy standards. First, they are obligated to have their sales-weighted average fuel economy exceed the minimum standard. Technically, this obligation extends to each weight segment separately, but firms are allowed to apply excess credits from one weight category to offset a deficit in another. Thus, in the end, the policy is functionally similar to the U.S. CAFE program, where there is one fleetwide requirement. Automakers have to pay fines if they cannot meet the standard by the target year of the regulation.

Second, in some years there are subsidies and tax exemptions that apply to individual cars if their fuel economy exceeds the standard for their weight class by a certain percentage. In 2009, the government introduced these subsidies as part of an economic stimulus package. Since then, the government has changed the eligibility requirements and the amount of the subsidies.

An important difference between the fuel economy regulation and the subsidy policy is that the later provides an incentive for an individual car to meet the fuel economy standard (as opposed to the fleet), and it thus creates a subsidy system that is notched in both fuel economy and weight. By contrast, the fuel economy regulation does not provide a direct incentive for an individual car to meet the target because what automakers have to do is to meet the target at the sales-weighted average. In the current version of this paper, we recognize that automakers have only the first incentive before 2009 and that they have both incentives after 2009. but we do not currently exploit the notches in the fuel economy dimension. We are currently developing strategies for exploiting the variation in subsidies and tax exemptions in recent years to conduct further empirical analysis.

[10] We are in the process of collecting data for the years before 2004, and we plan to extend our analysis when those data become available.

Table 1: Summary Statistics

Year	N	Fuel Economy (km/liter)	Weight (kg)	Displacement (liter)
2004	1,558	14.2	1,257	1.8
2005	1,529	14.3	1,252	1.8
2006	1,287	13.2	1,352	2.1
2007	1,289	13.3	1,368	2.1
2008	1,157	13.4	1,389	2.1
2009	1,230	13.4	1,393	2.2
2010	1,238	13.4	1,433	2.2
2011	1,156	13.6	1,444	2.2
2012	1,384	14.6	1,399	2.0
2013	1,488	14.1	1,484	2.3

Note: This table shows the number of observations, means of fuel economy, weight, and displacement for each year. Data are not sales weighted.

4.2 Data

We analyze data from the Japanese Ministry of Land, Infrastructure, Transportation, and Tourism (MLIT) that records fuel economy data for all new vehicles sold in Japan for each year between 2004 and 2013.[11] The record includes each vehicle's model year, model name, manufacturer, engine type, displacement, transmission type, weight, fuel economy, fuel economy target, estimated carbon dioxide emissions, number of passengers, wheel drive type, and devices used for improving fuel economy. The Ministry data does not include sales. Table 1 presents summary statistics of these data. There are between 1,100 and 1,600 different vehicle configurations sold in the Japanese automobile market each year. This includes both domestic and imported cars. The average fuel economy (not sales-weighted) declined slightly in 2006 and increased again around 2012. There is an overall positive trend in weight and displacement over time.

Although most of our analysis uses cross-sectional data, we do conduct panel data analysis in Section 4.5. To do so, we match data between different years in the following way. In our data set, each observation has a product identifier (ID). The product ID is a narrower definition than model name. For example, a Honda Civic may have several product IDs in the same year because there are Civics with different transmissions, displacements and drive types. We first match on product ID across years, which is often, but not always, constant over time. If automakers change the product

[11] The Ministry may have some data from before 2004. We are in the process of verifying availability of this earlier data.

ID between years, we match by using model name, displacement, drive type (e.g. four-wheel drive), and transmission (manual or automatic). That is, we consider two cars sold in two different years to be the same if they have matching IDs, or if they have the same model name, displacement, drive type, and transmission. We have experimented with matching on differing sets of characteristics and found that our key results are robust to the different criteria.

4.3 Excess Bunching at Notches in Fuel Economy Standard

Our theoretical model indicates that the main welfare distortions resulting from attribute-based standards come from distortions in weight, and that for notched incentives, the distortion should be evident from excess mass in the weight distribution at cutoff values. We thus begin our analysis by looking at the distribution of the raw data to look for evidence of bunching. Figure 7 plots the fuel economy schedule for 2006, along with a scatter plot of each vehicle's fuel economy and weight in that year. We use 2006 as an example; the pictures for all other years reveal similar patterns.

The diagram reveals three things. First, there is a negative relationship between fuel economy and vehicle weight. The negative relationship comes from the fact that it is harder to achieve high fuel economy for heavier vehicles. Second, we observe "hugging" of many observations on the fuel economy standard schedule. That is, many vehicles have exactly the same fuel economy as the fuel economy standard, or just slightly higher fuel economy than the standard. Finally, and most prominently, there are excess numbers of observations at the notch points in vehicle weight, which occur at 710, 830, 1020, 1270, 1520, 1770, 2020, and 2270 kg.

To see the excess numbers of observations at the weight notches more clearly, we next provide histograms of vehicle weight superimposed on the notched schedules. Panel A of Figure 8 is the histogram of vehicles sold between 2004 and 2008. The bin size of the histogram is 10 kg. That is, each bar represents the density of the number of observations for each 10 kg bin. In this period, all vehicles had an old fuel economy standard schedule, which is presented in the same figure. There is visually clear excess bunching at each notch of the fuel economy standard schedule. The magnitude of the bunching is substantial. It is visually clear that the number of observations at each notch point is more than twice as large as the number of observations around the notch points.

Panel B of Figure 8 provides evidence that the excess bunching is the result of automakers' response to the notched schedule. After a policy change in fuel economy standard, automakers

Figure 7: Scatter Plot of Vehicle Weight (kg) against Fuel Economy (km/liter) in 2006

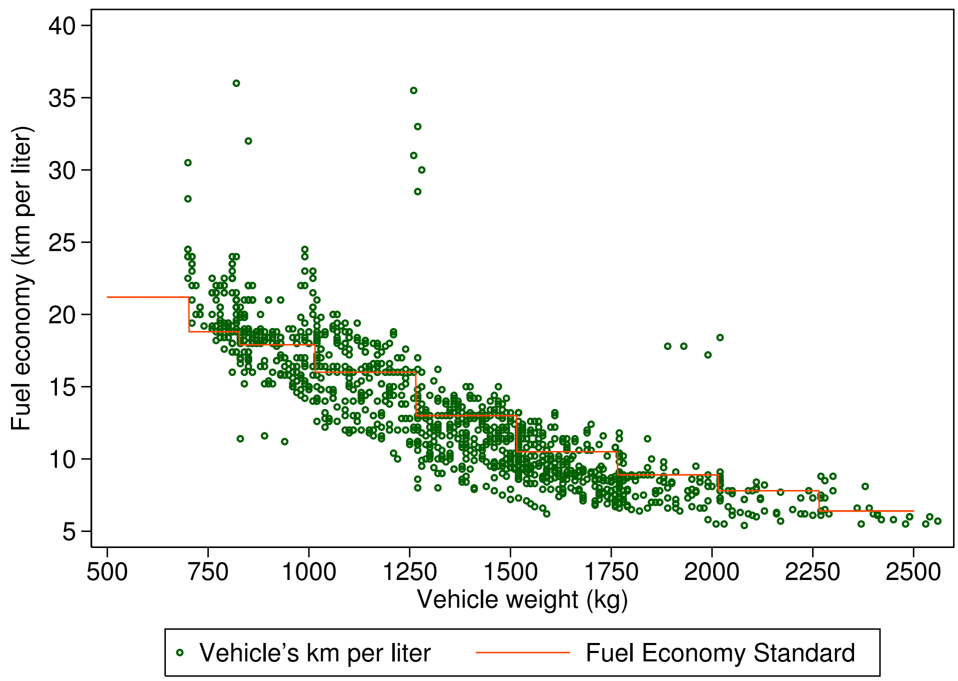

Note: This figure plots each vehicle's fuel economy (km per liter) in 2006 against the vehicle's weight. Each dot represents one observation.

faced the new fuel economy standard schedule in 2013 as we present in Panel B. By comparing Panel A and B, one can see that the mass points shifted precisely in accordance with the change in the locations of the notch points. In sum, the raw data show strong evidence that the market responded to the attribute-based regulation by changing vehicle weight. Our theory suggests that, on its face, this is evidence of distortion. In the next section, we use an econometric method similar to that used in Saez (1999 2010); Chetty, Friedman, Olsen, and Pistaferri (2011); Kleven and Waseem (2013) to estimate the magnitude of the excess bunching and discuss the implications of this bunching for welfare.

4.4 Estimation of Excess Bunching at Notches

Econometric estimation of excess bunching in kinked or notched schedules is relatively new in the economics literature. Saez (1999) and Saez (2010) estimate the income elasticity of taxpayers in the U.S. with respect to income tax rates and EITC schedules by examining excess bunching around

Figure 8: Histogram of Vehicles and Fuel Economy Standard

Panel A. Years 2004 to 2008 (Old Fuel Economy Standard Schedule)

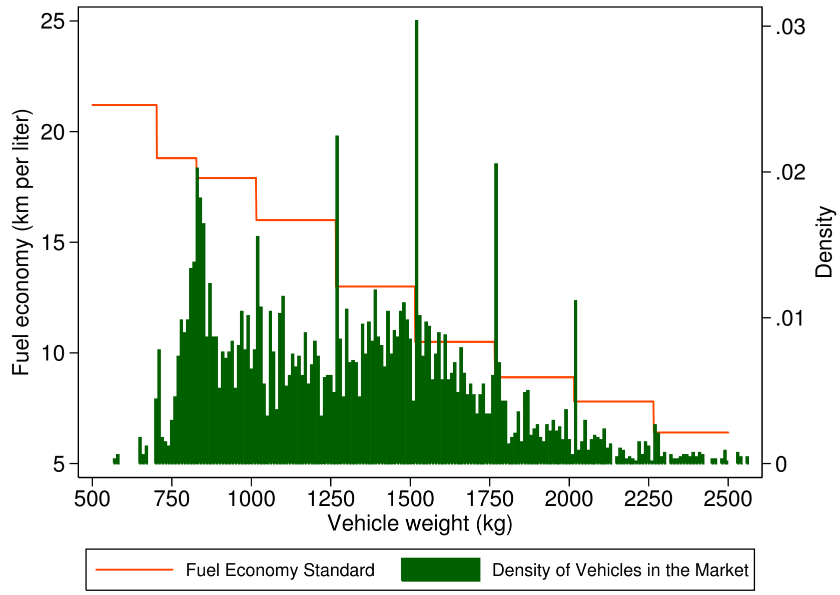

Panel B. Year 2013 (New Fuel Economy Standard Schedule)

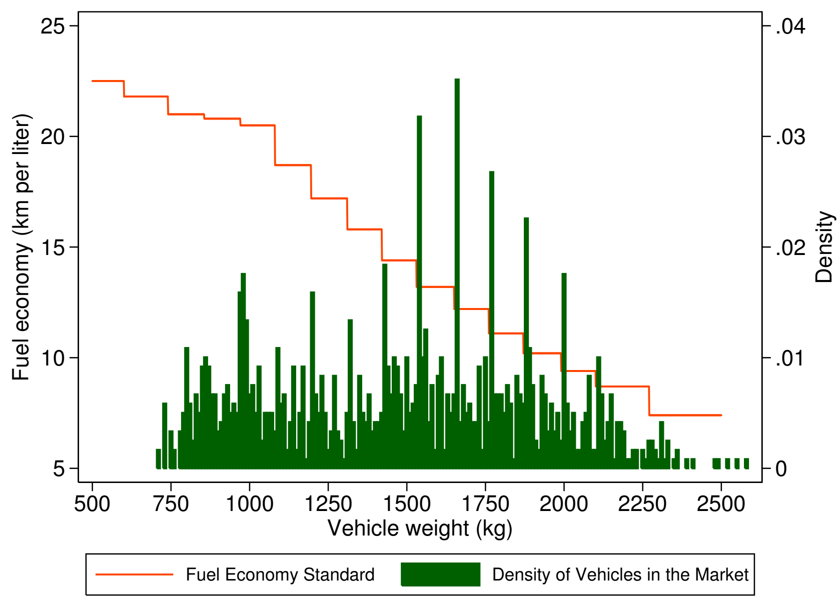

Note: Panel A shows the histogram of vehicles from 2004 to 2008, where all vehicles had the old fuel economy standard. In 2013, essential all vehicles had the new fuel economy standard, which is presented in Panel B.

kinks in the U.S. personal income tax schedule. Similarly, Chetty, Friedman, Olsen, and Pistaferri (2011) estimates the income elasticity of taxpayers in Denmark with respect to income tax rates by examining the excess bunching in the kinked tax schedules there. In Pakistan, the income tax schedule has notches instead of kinks. That is, the *average* income tax rate is piecewise linear. Kleven and Waseem (2013) uses a method similar to Chetty, Friedman, Olsen, and Pistaferri (2011) to estimate the elasticity of income with respect to income tax rates using bunching around these notches. Our approach is closely related to these papers, although our application is a fuel economy regulation, not an income tax.

To estimate the magnitude of the excess bunching, we first estimate the counterfactual distribution as if there were no bunching at the notch points. We start by grouping vehicles into small weight bins (10 kg bins in the application below) and denote by c_j the number of vehicles and by w_j the weight level in bin j. We fit a polynomial of order q to the bin counts in the empirical distribution, excluding bins in the range (w_L, w_U). We use $w_L = 10$, and $w_U = 10$ in the applications below. To minimize specification errors of the polynomial fit, we locally estimate the counterfactual distribution separately for each notch point. For each estimation, we include bins whose weight is larger than w_j - 200 and less than w_j + 200. For observations around each notch point k, we estimate:

$$c_j = \sum_{p=0}^{q} \beta_p \cdot (w_j)^p + \sum_{h=w_L}^{w_U} \gamma_h \cdot 1\{w_j = h\} + \varepsilon_j, \tag{15}$$

where γ_h is a bin fixed effect for each bin in the excluded range (so that the regression gives a perfect fit in that range). The counterfactual distribution is then estimated as the predicted values from (15) omitting the contribution of the dummies in the excluded range, i.e. $\hat{c_j} = \sum_{p=0}^{q} \hat{\beta}_p \cdot (w_j)^p$. Excess bunching is the difference between the observed and counterfactual bin counts in the part of the excluded range at the notch point: $\hat{B} = \sum_{j=w_L}^{w_U} (c_j - \hat{c_j}) = \sum_{h=w_L}^{w_U} \hat{\gamma_h}$. We use $q = 3$ for the polynomial fit. Because we estimate the counterfactual distribution locally, using polynomials higher than third-order does not change our results. To provide precise estimates, we pool the data from 2004 to 2008 to estimate equation (15), although we still find fairly precise estimates of excess bunching even if we run the estimation separately for each year.

Figure 9 shows our estimation procedure graphically for two notch points. In Panel A, we plot the actual distribution and estimated counterfactual distribution around the 1520 kg notch point. Graphically, our estimate of excess bunching is the difference in height between the actual and counterfactual distribution at the notch point. The estimate and standard error of excess bunching

Figure 9: Graphical Illustration of Estimation of Excess Bunching at Each Notch Point

Panel A. Notch at 1520 kg

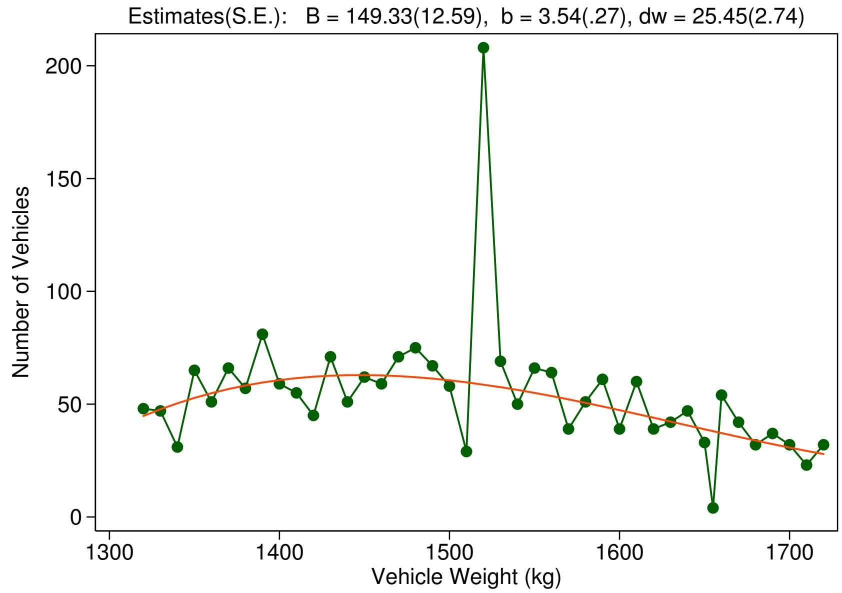

Panel B. Notch at 1770 kg

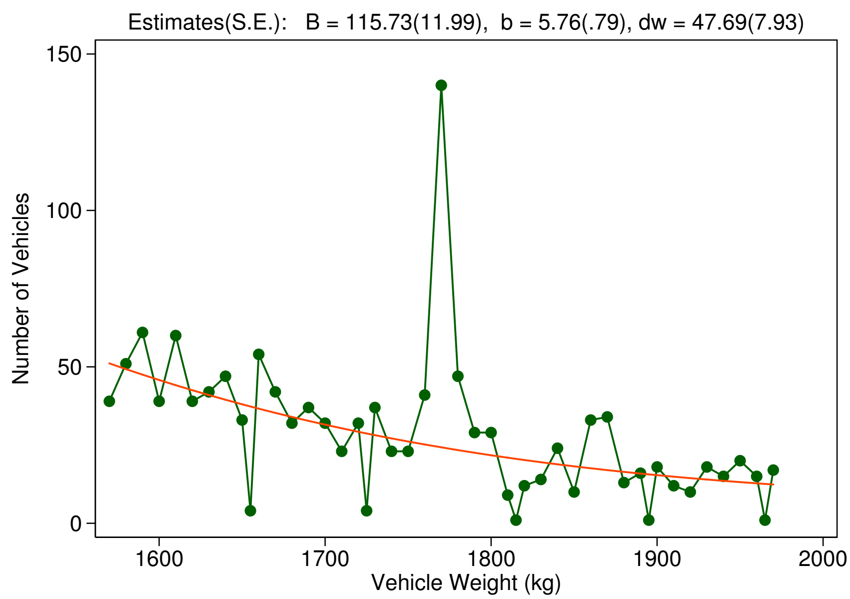

Note: This figure graphically shows the estimation in equation (15). The counterfactual distribution is estimated by the third-order polynomial of weight in the local area within 200 kg from each notch point. The figure also lists the estimates of B (excess bunching), b (proportional excess bunching), and dw (the range of weight affected by the notch). See the main text for details on these statistics.

Table 2: Estimation of Excess Bunching at Each Notch Point

Panel A. Notches at 710, 830, 1020, and 1270 kg

	(1) 710kg	(2) 830kg	(3) 1020kg	(4) 1270kg
Excess Bunching	37.00	73.88***	58.21***	106.83***
	(23.71)	(24.55)	(17.70)	(13.53)
Estimate of b (ratio)	3.31	2.15	2.21	3.31
(S.E.)	(2.12)	(.42)	(.4)	(.36)
Estimate of Δw (kg)	23.11	11.52	12.18	23.13
(S.E.)	(21.24)	(4.21)	(4.08)	(3.64)

Panel B. Notches at 1520, 1770, 2020, and 2270 kg

	(1) 1520kg	(2) 1770kg	(3) 2020kg	(4) 2270kg
Excess Bunching	149.33***	115.73***	64.59***	13.42***
	(12.59)	(12.00)	(7.01)	(3.91)
Estimate of b (ratio)	3.54	5.76	6.66	3.93
(S.E.)	(.27)	(.79)	(1.08)	(1.15)
Estimate of Δw (kg)	25.45	47.69	56.6	29.3
(S.E.)	(2.74)	(7.93)	(10.88)	(11.59)

Note: This table shows the regression result in equation 15. The dependent variable is the number of vehicles in each of 10 kg bins. We include a constant term and first-, second-, and third-order polynomial of weight in the regression. Standard errors are in the parentheses. ***, **, and * show 1, 5, and 10% levels of statistical significance from zero.

B is 149.33 (12.59). We also define the proportional excess bunching, b, which is the ratio between the actual and counterfactual distribution at the notch point. In this case, $b = 3.54$ (0.27), which means that the observed distribution has 3.54 times more observations than the counterfactual distribution at the notch point. Similarly, we illustrate our estimation around the 1770 kg notch point. At this notch, $B = 115.73$ (11.99) and $b = 5.76$ (0.79), meaning that the height of the excess bunching is 5.76 times higher than the counterfactual distribution at the notch point.

Table 2 presents our estimation results for each notch. Excess bunching is statistically significant from zero for all notches except for the first notch point, around which there are very few observations (see Panel A of Figure 8). The results show that proportional excess bunching, b, is particularly large for the notch points for heavier vehicles.

What do the estimates of B and b tell us about welfare? Our theoretical model indicated that, under our assumptions about market structure, excess bunching should come from the "left"; that is, automakers should increase vehicle weight to reach notch points, not decrease it. Based on that assumption, we can use the estimates of B and b to estimate how much automakers increase vehicle weight in response to the attribute-based fuel economy regulation. Consider a notch point at $w*$. Suppose that vehicles with weight $\in (w^* - \Delta w, w^*)$ bunch at w^* because of the notch schedule at w^*. Denoting the counterfactual distribution of vehicle weight by $h_0(w)$, excess bunching B at notch point w can be characterized by:

$$B = \int_{w^*-\Delta w}^{w^*} h_0(w)dw \approx h_0(w^*) \cdot \Delta w^*. \tag{16}$$

The approximation comes from the fact that the counterfactual distribution is fairly close to a uniform distribution at the local area of each notch point. Because our estimation provides estimates of B and $h_0(w^*)$, we can calculate estimates of Δw^* as $\Delta w^* = B/h_0(w^*)$.

The last rows of Table 2 shows our estimates of Δw^* for each notch point. The estimate is largest at the 2020 kg notch point, where $\Delta w^* = 56.6(10.88)$. Given the assumption that all bunching comes from the left of the notch point, the estimate implies that vehicles with weight $\in (2020 - 56.6, 2020)$ bunch at the 2020 kg notch point. If the distribution of vehicles in this area is close to a uniform distribution, the estimate implies that the average increase in weight is 56.6/2 = 28.3 kg. Given similar assumptions, the results imply that the average increase in vehicle weight for each notch point is between 6 kg and 28 kg.

Overall, the results in this section provides evidence that automakers significantly respond to the attribute-based fuel economy regulation by changing their vehicle weight. We show that given the assumption that bunching comes from the left (lower weight) of each notch, our estimates of excess bunching imply that automakers increase vehicle weight between 6 kg and 28 kg to reach less stringent fuel economy regulations. In the next section, we exploit our panel data and a policy change in fuel economy regulation to test the assumption about the direction of bunching.

4.5 Panel Data Analysis: Where does bunching come from?

The previous section provided visually clear and statistically significant evidence of bunching at the notch points in the weight dimension. Theory suggests that bunching should come from the left (weight should increase), but we wish to test this assumption with the data. To do so, we exploit

Figure 10: Changes in Vehicle Weight from 2008 to 2013 for Samples that Bunch in 2013

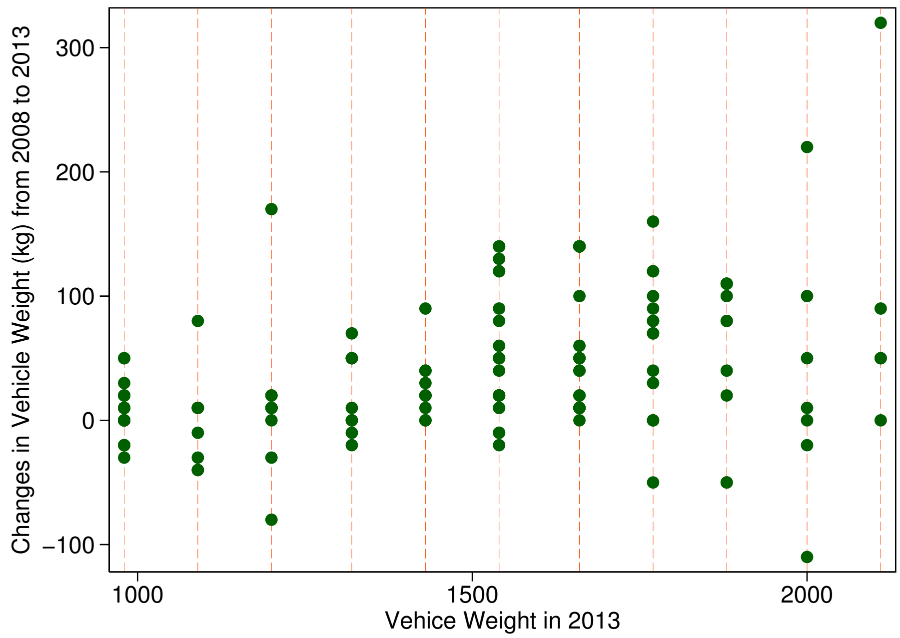

Note: This figure plots the raw data of individual vehicles that bunch in the notch points in the fuel economy standard in 2013. The horizontal axis shows the vehicle weight in 2013, the vertical axis shows the change in weight from 2008 to 2013, and the dashed vertical lines present the notch points.

a policy change in the Japanese fuel economy standard that changed the location of the notched schedules between 2008 and 2013.

The two policy regimes are shown in Figure 6. Before 2008, automakers faced the old fuel economy standard represented by the dashed line. In 2009, the new standard was introduced. The years between 2009 and 2012 were a transition period, during which time automakers were allowed to choose which schedule they wished to use. In 2013, all automakers were required to be on the new standard. We focus our comparison on 2008 and 2013 because the automakers did not have a choice about the schedule in those years. For these two years, we use our panel data, which links the same car type over time, to see how vehicle weight and fuel economy changed.

We take two approaches to investigating whether bunching comes from the left or right. Our first approach is to examine whether vehicles that we observed bunching in 2013 featured an increase or a decrease in their weight between 2008 and 2013. In Panel B of Figure 8, we see significant bunching within 10 kg of the notches at 970, 1090, 1200, 1320, 1430, 1540, 1660, 1770, 1880, 2000,

Figure 11: Kernel Density of Changes in Vehicle Weight from 2008 to 2013

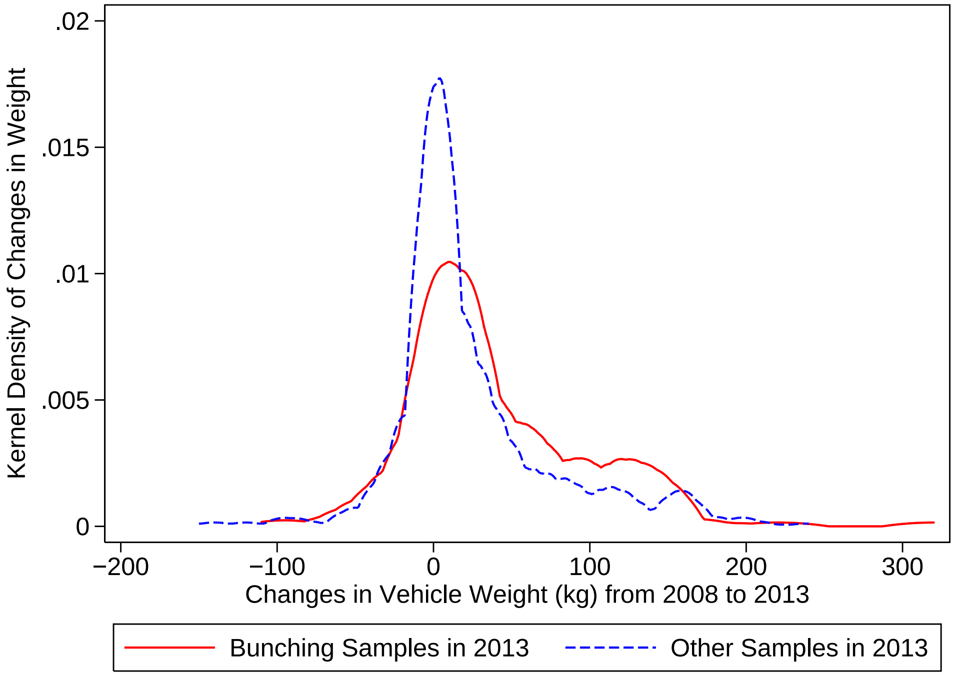

Note: This figure shows the kernel density of the changes in weight from 2008 to 2013 for the bunching samples and other samples in 2013.

and 2110 kg. In Figure 10, we take the sample of those vehicles and plot their change in weight between 2008 and 2013 against their final weight in 2013. The horizontal axis shows the vehicle weight in 2013, the vertical axis shows the change in weight from 2008 to 2013, and the dashed vertical lines show where there are weight notches. For most vehicles in the diagram, the change in weight is positive, which suggests that bunching comes from weight increases.

There was an overall increase in weight, however, in vehicles between those years, so that Figure 10 may be partly the product of a secular trend. To account for any secular trend in weight, we use the vehicles that do not bunch in 2013 as a control group to establish a counterfactual. We focus on vehicles that have a weight between 960 and 2020 kg, which represents a significant majority of the market, as a control group because light vehicles (those weighing less than 800 kg) appear to have had a significantly different weight trend.[12]

Figure 11 shows the kernel density of the changes in weight from 2008 to 2013 for the vehicles

[12]Note that adding the excluded light cars does not change results in our regression analysis, below, in which we include a polynomial that allows the trend to differ vehicles of different weights.

Table 3: Changes in Vehicle Weight from 2008 to 2013: Does Bunching Come from Left or Right?

	(1) ΔWeight	(2) ΔWeight	(3) ΔWeight	(4) ΔWeight
1{Bunching in 2013}	13.79***	13.23***	13.62***	12.48**
	(5.25)	(5.04)	(4.98)	(5.07)
(Weight in 2013)/1000		46.14***	261.34***	-131.18
		(6.83)	(54.79)	(346.21)
(Weight in 2013)/1000)^2			-71.47***	185.45
			(18.06)	(224.48)
((Weight in 2013)/1000)^3				-54.03
				(47.05)
Constant	24.56***	-42.32***	-196.26***	-4.05
	(2.81)	(10.26)	(40.19)	(172.16)
Observations	531	531	531	531
R^2	0.013	0.091	0.118	0.120

Note: This table shows the regression result in equation 17. The dependent variable is the change in vehicle weight (kg) from 2008 to 2011. 1{Bunching in 2013} is a dummy variable or samples that bunch in the fuel economy standard schedule in 2013. Standard errors are in the parentheses. ***, **, and * show 1, 5, and 10% levels of statistical significance from zero.

that bunched in 2013 and those that did not. The sample of those who bunched is pushed towards the right relative to the control group of non-bunchers. In particular, the distribution indicates that there are many vehicles who bunched at a weight notch after increasing weight by 50 to 150 kg.

To estimate the difference in the change in weight between the bunching and non-bunching vehicles, we estimate the following equation by OLS:

$$\Delta weight_i = \alpha + \beta \cdot 1\{bunching_i\} + f(weight_{i,2013}) + \varepsilon_i, \tag{17}$$

where $\Delta weight_i$ is the change in weight (kg) from 2008 to 2013 for vehicle i and $1\{bunching_i\}$ equals one if vehicle i is located within 10 kg of a weight notch in 2013. Regressing $\Delta weight_i$ on $1\{bunching_i\}$ estimates a difference in mean weight between the two groups. This simple regression potentially produces a bias if the change in weight is systematically larger for lighter or heavier vehicles. In our case, any potential bias is likely small because the bunching locations are widely spread in the weight dimension so that the bunching and non-bunching samples have fairly similar weight on average. Nevertheless, we aim to eliminate any bias that might exist by controlling for a polynomial in $weight_{i,2013}$.

Table 3 presents regression results. Column 1 shows that the simple mean difference in the change in weight between the bunching samples and others is 13.79 kg. That is, vehicles that bunched increased their weight by 13.79 kg more than other vehicles. In column 2, we control for any linear relationship between the change in weight and final weight by adding a linear control for weight in 2013. As expected, this yields a positive coefficient on weight, which means that heavier vehicles had slightly larger increases in weight, but this has very little impact on our coefficient of interest, as expected. We include the second- and third-order polynomials in columns 3 and 4, which has no significant impact on our coefficient of interest.

These regressions and figures indicate that, on average, bunching comes from the left of the notch points. That is, automakers *increase* their vehicle weight in response to the attribute-based regulation notches, and our estimates suggest that the mean of this weight increase is about 13 kg. We go further in the next section by using variation in incentives created by the policy change to estimate a weight response elasticity.

4.6 Panel Data Analysis: Elasticity of Attributes with Respect to Regulation

Figure 6 shows that the new regulation increased the fuel economy requirement for vehicles in all weight classes, but the increase was of a different magnitude for vehicles in different weight classes. For example, compare a vehicle that weighed 1420 kg in 2008 to one that weighed 1430 kg in that year. The former vehicle faced an increase in its fuel economy standard from 13 to 15.8 km/liter (a 22% increase), while the latter vehicle faced an increase in its fuel economy standard from 13 to 14.3 km/liter (an 11% increase).[13] In this section, we exploit the discontinuities in the variation in the changes in fuel economy standards to estimate the elasticity of attributes with respect to fuel economy regulation.

We use a regression discontinuity design (RDD) that is similar to the identification strategy in Saez (2003). Saez uses discontinuous variation in income tax rates created by "bracket creep" in

[13]Note that this variation is driven by the front-runner in each category, which we treat as plausibly random variation.

the U.S. income tax schedule. In his study, there is no policy change in the income tax schedule in nominal terms. However, inflation is high enough to shift the income tax schedule towards the right in real terms, causing taxpayers to face differing real income tax schedules. In our context, we have changes in 1) the locations of the kink points and 2) the stringency of the fuel economy regulation, which provides powerful variation enabling us to estimate how variation in the stringency of regulation affects attribute changes.

Our estimation is based on the following OLS regression:

$$\Delta weight_i = \alpha + \beta \cdot \Delta standard_i + \gamma \cdot X_i + \varepsilon_i, \tag{18}$$

where $\Delta weight_i$ is the change in weight (kg) from 2008 to 2013 for vehicle i, $\Delta standard_i$ is the change in the fuel economy standard (km/liter) for i, and X_i is a vector of control variables. OLS estimates will be biased because $standard_{it}$ is a function of $weight_{it}$ for $t \in (2008, 2013)$. If there are unobservable shocks in ε_i, they will affect $\Delta standard_i$ as well as $\Delta weight_i$. Because the fuel economy standard is a decreasing step function of weight, we expect that OLS estimation produces a downward bias for β.

To address the endogeneity, we use a policy-induced change in the standard as an instrument, as in Saez (2003); Saez, Slemrod, and Giertz (2012); Ito (Forthcoming). Specifically, we use $\Delta standard_i^{PI} = standard_{i,2013}(weight_{i,2008}) - standard_{i,2008}(weight_{i,2008})$. This instrument, which is sometimes called a simulated instrument, computes the predicted change in the standard at a weight level $weight_{i,2008}$. The instrument thus captures the change induced by the policy change for the weight level $weight_{i,2008}$. To be a valid instrument, $weight_{i,2008}$ has to be uncorrelated with ε_i, because the instrument is a function of $weight_{i,2008}$. This condition can be violated if lighter cars and heavier cars in 2008 have different underlying changes in weight. The advantage of our identification strategy is that we can include any smooth controls for $weight_{i,2008}$ in X_i to account for such concerns. Because the instrument's variation in the change in standard is discontinuous in weight, including any smooth controls for $weight_{i,2008}$ does not eliminate our ability to identify the coefficient. In our estimation, we use $\Delta standard_i^{PI}$ as an instrument for $\Delta standard_i$ and include a third-order polynomial of $weight_{i,2008}$ in X_i.[14]

[14] Including yet higher orders of the polynomial does not change our results.

Table 4: Elasticity of Attributes (weight) with Respect to Fuel Economy Standards

	ΔWeight(kg)			Δln(Weight)	
	(1)	(2)	(3)	(4)	(5)
ΔFuel Economy Standard	-8.67***	18.70**	18.86***		
	(1.69)	(7.37)	(7.26)		
Δln(Fuel Economy Standard)				0.16***	0.16***
				(0.05)	(0.05)
(Weight in 2008)/1000		479.58***	313.30	0.24***	0.18
		(76.40)	(434.38)	(0.04)	(0.26)
(Weight in 2008)/1000)^2		-146.39***	-36.23	-0.08***	-0.04
		(24.04)	(278.79)	(0.01)	(0.17)
((Weight in 2008)/1000)^3			-23.40		-0.01
			(58.19)		(0.04)
Constant	42.00***	-371.45***	-291.45	-0.18***	-0.14
	(3.51)	(66.05)	(224.59)	(0.03)	(0.13)
Observations	531	531	531	531	531
Estimation	OLS	RD(2SLS)	RD(2SLS)	RD(2SLS)	RD(2SLS)

Note: This table shows the regression result in equation 18. The dependent variable is the change in vehicle weight (kg) from 2008 to 2011 or log of the change in weight. Δ Fuel Economy Standard is the change in fuel economy standard. Standard errors are in the parentheses. ***, **, and * show 1, 5, and 10% levels of statistical significance from zero.

Table 4 presents estimation results. As expected, the OLS estimate in column 1 is lower than our RD estimates in columns 2 and 3. This is likely due to the downward bias stemming from the endogeneity of $\Delta standard_i$. In columns 2 and 3, we show our RD estimates. The first stages are very strong because the policy-induced change in fuel economy standard strongly affects the actual change in standard. After we include the first- and second-order of polynomials in $weight_{i,2008}$, our estimates are not sensitive to the inclusion of higher-orders of polynomials. The estimates imply that a one unit increase in the fuel economy standard (km/liter) results in an increase in vehicle weight of19 kg. In columns 4 and 5, we use a log-log RD specification to estimate a constant elasticity. The estimate implies that an 1% increase in the fuel economy standard results in a 0.16% increase in vehicle weight. These results provide empirical evidence that automakers change their vehicle weight in response to the change in the stringency of fuel economy regulation.

In Table 5, we estimate equation 18 but use $\Delta FuelEconomy_i$ as the left-hand side variable.

Table 5: Elasticity of Attributes (fuel economy) with Respect to Fuel Economy Standards

	ΔFuel Economy (km/liter)			Δln(Fuel Economy)	
	(1)	(2)	(3)	(4)	(5)
ΔFuel Economy Standard	-0.00	0.22	0.20		
	(0.06)	(0.21)	(0.21)		
Δln(Fuel Economy Standard)				0.22	0.21
				(0.16)	(0.16)
(Weight in 2008)/1000		7.80***	23.07*	0.56***	1.25
		(2.18)	(12.34)	(0.13)	(0.79)
(Weight in 2008)/1000)^2		-2.61***	-12.73	-0.19***	-0.65
		(0.69)	(7.92)	(0.04)	(0.51)
((Weight in 2008)/1000)^3			2.15		0.10
			(1.65)		(0.11)
Constant	0.26**	-5.61***	-12.95**	-0.40***	-0.73*
	(0.12)	(1.89)	(6.38)	(0.10)	(0.39)
Observations	531	531	531	531	531
Estimation	OLS	RD(2SLS)	RD(2SLS)	RD(2SLS)	RD(2SLS)

Note: This table shows the regression result in equation 18 but uses $\Delta FuelEconomy_i$ as the left-hand side variable. $\Delta FuelEconomy_i$ is the change in fuel economy (km/liter) for vehicle i from 2008 to 2013. Standard errors are in the parentheses. ***, **, and * show 1, 5, and 10% levels of statistical significance from zero.

$\Delta FuelEconomy_i$ is the change in fuel economy (km/liter) for vehicle i from 2008 to 2013. In this estimation, we test whether increases in the fuel economy standard cause an increase in the fuel economy of vehicles. The estimates indicate that more stringent fuel economy standards result in increases in fuel economy, but the estimates are not statistically significantly different from zero.

In summary, our empirical analysis provides three findings: 1) there is clear and statistically significant evidence that vehicles bunch at the attribute (weight) notch points in the Japanese fuel economy standard; 2) given the assumption (backed by theory) that bunching comes from the left (i.e., bunching comes from increases in weight), the excess bunching in our cross-sectional variation implies that the average increase in vehicle weight for each notch point is between 6 kg and 28 kg; and 3) our panel data analysis provides supporting evidence for this assumption, as we find that a one unit increase in the fuel economy standard results in an increase in vehicle weight of 19 kg and a one percent increase in fuel economy standard results in an increase in vehicle weight by 0.16%.

4.7 Welfare implications of weight manipulation

As emphasized in the theory section, the main welfare loss from the response to attribute based standards in the situation where the attribute itself has an unpriced externality will be the inefficiency from an increase in the attribute. An approximation of the distortion from the policy can therefore be calculated by multiplying the change in the attribute by the externality, as indicated in equation 9. Our empirical estimates above suggest that the increase in weight caused by the attribute-basing in the Japanese car policy is between 10 and 20 kg.

To estimate the externality associated with increased weight, we use estimates from Anderson and Auffhammer (Forthcoming), which concludes that an increase in vehicle weight of 1000 pounds (454 kg) is associated with a 0.09 percentage point increase in the probability that the vehicle is associated with a fatality, on a mean probability of 0.19%. We use a standard estimate of the value of a statistical life (VSL) of $7 million. (Both of these estimates come from the U.S., which is a potential weakness of the calculation.)

We thus calculate the welfare loss, per car sold, for a 10 kg as: 0.0009 * (10/454) * $7 million = $139 per car. An increase of 20 kg is likewise associated with a welfare loss of $278 per car. The Japanese car market sells around 5 million new cars per year, which implies an aggregate welfare distortion of between $0.7 and $1.4 billion per year.

5 Conclusion

This paper explores the economic implications of attribute-based regulation. We develop a theoretical framework that highlights conditions under which attribute basing is inefficient, and we show that, under those conditions, the use of attribute-based regulation leads to distortions that are concentrated in the provision of the attribute upon which targets are based. The model also explores cases where attribute-basing may play a role, but emphasizes that, even in those cases, the attribute-basing function deviates fundamentally from those observed in real world policies.

Empirically, the paper demonstrates that distortions in response to attribute-based fuel economy standards in Japan are clearly present. We use both established cross-sectional tools based on the notch literature, as well as novel panel techniques, to demonstrate that the Japanese car market has experienced a notable increase in weight in response to attribute-basing.

The theoretical framework makes a number of simplifying assumptions that could be relaxed in future work. In particular, it is possible that attribute-based regulation, particularly when it features notches, impacts firm pricing strategies for certain types of products. Further exploration of the cases in which attribute-basing may be justified is also warranted. On the empirical side, evidence of the responsiveness of attributes other than vehicle weight, which may be particularly easy to manipulate, would be a valuable object of study for future research.

References

Anderson, Michael and Maximilian Auffhammer. Forthcoming. "Pounds that Kill: The External Costs of Vehicle Weight." *Review of Economic Studies* .

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.

Blinder, Alan S. and Harvey S. Rosen. 1985. "Notches." *American Economic Review* 75 (4):736–747.

Borenstein, Severin. 2013. "A Microeconomic Framework for Evaluating Energy Efficiency Rebound and Some Implications." Energy Institute at Haas Working Paper 242R.

Chetty, Raj, John N. Friedman, Tore Olsen, and Luigi Pistaferri. 2011. "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records." *The Quarterly Journal of Economics* 126 (2):749–804.

Gillingham, Kenneth. 2013. "The Economics of Fuel Economy Standards versus Feebates." Manuscript: Yale University.

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

Ito, Koichiro. Forthcoming. "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing." *American Economic Review* .

Jacbosen, Mark R. 2013. "Fuel Economy and Safety: The Influences of Vehicle Class and Driver Behavior." *American Economic Journal: Applied Economics* 5 (3).

Jacobsen, Mark R. and Arthur A. van Benthem. 2013. "Vehicle Scrappage and Gasoline Policy." NBER Working Paper 19055.

Kleven, Henrik J. and Mazhar Waseem. 2013. "Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan." *The Quarterly Journal of Economics* 128 (2):669–723. URL `http://qje.oxfordjournals.org/content/128/2/669`.

Knittel, Christopher R. 2011. "Automobiles on Steroids: Product Attribute Trade-Offs and Technological Progress in the Automobile Sector." *American Economic Review* 101 (7):3368–3399.

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

Saez, E. 2003. "The effect of marginal tax rates on income: a panel study of bracket creep." *Journal of Public Economics* 87 (5-6):1231–1258.

Saez, Emmanuel. 1999. "Do Taxpayers Bunch at Kink Points?" *National Bureau of Economic Research Working Paper Series* No. 7366.

———. 2010. "Do Taxpayers Bunch at Kink Points?" *American Economic Journal: Economic Policy* 2 (3):180–212.

Saez, Emmanuel, Joel B. Slemrod, and Seth H. Giertz. 2012. "The elasticity of taxable income with respect to marginal tax rates: A critical review." *Journal of Economic Literature* 50 (1):3–50.

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

Sallee, James M. and Joel Slemrod. 2012. "Car Notches: Strategic Automaker Responses to Fuel Economy Policy." *Journal of Public Economics* 96 (11-12):981–999.

Whitefoot, Katie, Meredith Fowlie, and Steven Skerlos. 2013. "Evaluating U.S. Reformed Corporate Average Fuel Economy Standards using an Engineering Attribute Selection Model." Manuscript: University of California, Berkeley.

Whitefoot, Katie and Steven Skerlos. 2012. "Design Incentives to Increase Vehicle Size Created from the U.S. Footprint-based Fuel Economy Standards." *Energy Policy* 41:402–411.