

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

September 9, 2014

Abstract

We study “attribute-based regulations,” in which regulatory compliance depends upon some secondary attribute that is not the intended target of the regulation. We theoretically characterize how attribute-basing creates perverse incentives to distort the attribute upon which compliance depends. To test our theoretical predictions, we exploit quasi-experimental variation in Japanese fuel economy regulations, under which fuel-economy targets are downward-sloping step functions of vehicle weight. Our bunching analysis reveals large distortions to vehicle weight induced by the policy. We then leverage panel data on vehicle redesigns to empirically investigate the welfare implications of attribute-basing, including both potential benefits and likely costs.

Keywords: fuel-economy standards, energy-efficiency, corrective taxation, notches, bunching analysis

JEL: H23, Q48, L62

*Ito: ito@bu.edu. Sallee: sallee@uchicago.edu. The authors would like to thank Kunihito Sasaki for excellent research assistance. For helpful comments, we thank Hunt Allcott, Soren Anderson, Severin Borenstein, Meghan Busse, Raj Chetty, Lucas Davis, Francesco Decarolis, Meredith Fowlie, Don Fullerton, Michael Greenstone, Mark Jacobsen, Damon Jones, Hiro Kasahara, Kazunari Kainou, Ryan Kellogg, Ben Keys, Christopher Knittel, Ashley Langer, Bruce Meyer, Richard Newell, Matt Notowidigdo, Ian Parry, Mar Reguant, Nancy Rose, Mark Rysman, Jesse Shapiro, Joel Slemrod, Sarah West, Katie Whitefoot, Florian Zettelmeyer and seminar participants at the ASSA meetings, Berkeley, Boston University, Chicago, the EPA, Harvard, Michigan, Michigan State, MIT, the National Tax Association, the NBER, the RIETI, Pontifical Catholic University of Chile, Stanford, Resources for the Future, UCLA, University of Chile, and Wharton. Ito thanks the Energy Institute at Haas and the Stanford Institute for Economic Policy Research for financial support. Sallee thanks the Stigler Center at the University of Chicago for financial support. ©2014 Koichiro Ito and James M. Sallee. All rights reserved.

1 Introduction

The goal of this paper is to advance economists’ understanding of “attribute-basing”, which is a common feature of policies aimed at correcting externality-related market failures. An attribute-based regulation is a regulation that aims to change one characteristic of a product related to the externality (the “targeted characteristic”), but which takes some other characteristic (the “secondary attribute”) into consideration when determining compliance. For example, Corporate Average Fuel Economy (CAFE) standards in the United States recently adopted attribute-basing. Figure 1 shows that the new policy mandates a fuel-economy target that is a downward-sloping function of vehicle “footprint”—the square area trapped by a rectangle drawn to connect the vehicle’s tires.¹ Under this schedule, firms that make larger vehicles are allowed to have lower fuel economy. This has the potential benefit of harmonizing marginal costs of regulatory compliance across firms, but it also creates a distortionary incentive for automakers to manipulate vehicle footprint.

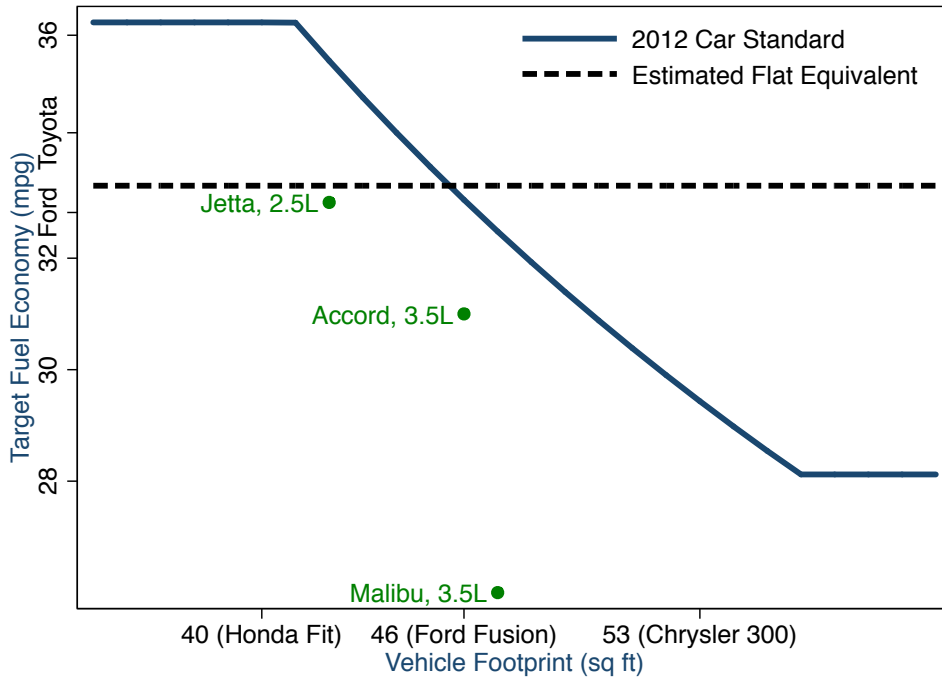
Attribute-basing is used in a variety of important economic policies. Fuel-economy regulations are attribute-based in China, Europe, Japan and the United States, which are the world’s four largest car markets. Energy efficiency standards for appliances, which allow larger products to consume more energy, are attribute-based all over the world. Regulations such as the Clean Air Act, the Family Medical Leave Act, and the Affordable Care Act are attribute-based because they exempt some firms based on size. In all of these examples, attribute-basing is designed to provide a weaker regulation for products or firms that will find compliance more difficult. This can both improve equity and enhance efficiency by equalizing marginal costs of compliance across actors, but it will generally come at the cost of distortions in the choice of the attribute upon which the regulation is based.

Despite the ubiquity of attribute-basing, its welfare implications are not well understood by economists. In this paper, we develop a theoretical framework that characterizes the welfare implications of attribute-basing.² We then test our theoretical predictions by exploiting quasi-

¹Under CAFE, firms must meet the attribute-based target on average. Not every vehicle must be in compliance.

²The real world attribute-based policies that we discuss here appear to be motivated by the goal of compliance cost equalization, and this motivation is therefore our focus. Another potential motivation for attribute-basing is to regulate secondary attributes when policymakers cannot directly target the externality. That is, if the targeted characteristic is imperfectly related to the externality, and the secondary attribute is correlated with the externality, conditional on the targeted characteristic, then optimal policy may involve attribute-basing. We discuss this alternative rationale for attribute-basing further below, though we do not model it explicitly.

Figure 1: Example of an Attribute-Based Regulation: 2012 U.S. CAFE standards



As of 2012, the U.S. Corporate Average Fuel Economy (CAFE) standards mandate a fuel-economy target that is a downward-sloping function of vehicle footprint. The flat equivalent is the EPA’s estimate based on the distribution of footprint in current fleet. Prior to 2011, the 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, as are three individual models. All data taken from the Federal Register, 40 CFR Part 85.

experimental variation in Japanese fuel economy regulations, under which fuel-economy targets are downward-sloping step functions of vehicle weight. Finally, we leverage panel data on vehicle redesigns to empirically investigate the welfare implications of attribute-basing, including both potential benefits and likely costs.

In the theory section, we develop a framework that translates attribute-based regulations (ABRs) into more familiar microeconomic incentives. We demonstrate that policies like CAFE can be thought of as creating a pair of implicit taxes in separate “markets” for the targeted characteristic and the secondary attribute. These wedges are proportional to each other and scaled by the slope of the compliance function with respect to the attribute. In many circumstances, the implicit tax or subsidy for the secondary attribute is purely distortionary. The distortionary incentive changes the level of the attribute and creates welfare losses that are directly analogous

to the deadweight loss from distortionary taxation. These distortions, which constitute the main cost of attribute-based policies, are greater as the attribute is more elastic with respect to policy incentives. This theoretical prediction motivates our empirical analysis, in which we show that attribute-based standards do indeed create large distortions in vehicle weight, which implies large welfare losses, in the Japanese automobile market.

Attribute-based regulation may also have benefits. Our model emphasizes that attribute-basing may improve efficiency by equalizing the marginal costs of regulatory compliance. To see this, consider the two cars labeled in Figure 1, the Accord and the Jetta. If both cars were forced to meet the same flat standard (the dashed line in the figure), the marginal cost of compliance would likely be higher for the Accord than the Jetta because the Accord must make a much larger change in its fuel economy. If both cars were instead required to meet the attribute-based standard, their marginal costs of compliance would likely be much more similar because each require about the same improvement in fuel economy. Equalization of marginal compliance costs is efficiency enhancing; it is equivalent to equalizing the marginal cost of providing a positive (or mitigating a negative) externality associated with the targeted characteristic. Our theory formalizes this intuition, and our empirical analysis illustrates it by simulating compliance costs under a flat standard versus an attribute-based standard in the Japanese automobile market.³

The example of the Accord and the Jetta points to *potential* benefits of ABR, but these may not be realized. For example, in practice, CAFE does not require each individual product to comply with the standard. Rather, an entire firm’s fleet must on average comply. Moreover, since 2012, firms that exceed the standard are able to sell their “excess compliance” to other firms. This implies that marginal compliance costs can be harmonized through a trading market, which obviates the need for ABR. We also show that when such a compliance trading scheme is not available, attribute-basing has a role in equalizing marginal compliance costs but is an imperfect substitute for compliance trading. Attribute-basing can only equalize compliance costs to the degree that the attribute predicts marginal costs. In addition, attribute-basing induces distortions in the attribute unless the attribute is perfectly unresponsive to policy incentives.

In the empirical analysis, we exploit quasi-experimental variation in Japanese fuel economy reg-

³Below, we also discuss additional considerations that may explain the use of attribute-basing in practice, including distributional goals, technological spillovers and additional distortions due to imperfect competition, though we do not formally model them.

ulations, where fuel economy targets are decreasing step functions of vehicle weight. The Japanese regulation offers three advantages for our empirical analysis. First, the Japanese attribute-based fuel economy regulation has been used for more than three decades, which provides us with a long analysis window. Second, there was a policy change during our sample period, which enables us to use panel variation. Third, the regulation features a “notched” compliance schedule—the requisite fuel economy for a vehicle is a decreasing *step function* of vehicle weight. Automakers therefore have a large incentive to increase vehicle weight only up to key thresholds where the mandated fuel economy drops discretely.

This notched policy schedule motivates our first empirical approach, which uses cross-sectional data to demonstrate that there is significant “bunching” (excess mass) in the distribution of vehicle weight around regulatory thresholds. Using methods developed in Saez (1999, 2010); Chetty, Friedman, Olsen, and Pistaferri (2011) and Kleven and Waseem (2013) to quantify the bunching in weight, we estimate that 10% of Japanese vehicles have had their weight increased in response to the policy. Among the affected vehicles, we estimate that weight rose by 110 kilograms on average. This not only works against the goal of petroleum conservation (because heavier cars are less fuel efficient), but also exacerbates accident-related externalities (because heavier cars are more dangerous to non-occupants). Our back-of-the-envelope estimate based on the value of a statistical life and estimates of the relationship between fatalities and vehicle weight suggests that this weight increase creates around \$1 billion of deadweight loss per year in the Japanese car market.

We also exploit panel variation created by a policy change during our sample period. In 2009, the Japanese government introduced a model-specific (rather than corporate average) subsidy for vehicles that exceeded a weight-based fuel economy threshold. This policy creates a “double notched” policy; eligibility requires that each model be above a step-function in fuel economy by weight space. Vehicles that are modified in order to become eligible for the subsidy reveal information about the relative costs of manipulating weight versus fuel economy. We use this revealed preference information and a discrete choice model to estimate the cost of modifying fuel economy and weight, leveraging quasi-experimental variation in incentives due to differences in the initial position of each model vis-à-vis the double-notched schedule.

We use these cost estimate to illustrate the welfare implications of attribute-basing by simulating three alternative policy scenarios—attribute-based fuel economy standards, a flat standard without

compliance trading, and a flat standard with compliance trading. The simulation shows key insights about potential benefits and limitations of attribute-based regulation. When compliance trading is disallowed, attribute-based standards improve efficiency as compared to a flat standard because attribute-basing helps equalize marginal compliance costs. However, this benefit is partially offset by distortions in the attributes created by the regulatory incentive. In addition, consistent with our theory, we find that attribute-basing is an imperfect substitute for compliance trading because the marginal compliance costs are not perfectly correlated with the attribute, which results in only partial equalization of the marginal compliance costs.

The findings from our theoretical and empirical analysis provide important policy implications for a variety of economic policies. In addition to the United States and Japan, a growing number of countries are adopting attribute-basing for energy efficiency policy such as fuel-economy regulation for cars and energy efficiency standards for appliances.⁴ Subsidies or taxes, such as the hybrid vehicle tax in the United States. or vehicle registration fees in Europe, are also attribute-based. Attribute-basing is also prominent in regulations outside of energy efficiency. For example, firms below a certain size (the attribute) are exempted from some policies, including provisions as diverse as the Clean Air Act, the Affordable Care Act, the Family Medical Leave Act, livestock regulations, and securities regulations. The stringency of other policies, such as liability for worker safety and unemployment insurance tax rates, also depend on firm size. Our work contributes to this literature by theoretically characterizing the potential benefits and limitations of attribute-basing and empirically testing the theoretical predictions using quasi-experimental variation in regulatory incentives.⁵

Fewer papers consider attribute-basing in energy-efficiency policies. There is a substantial literature on CAFE, including key contributions in Goldberg (1998), Kleit (2004), Gramlich (2009),

⁴For example, a refrigerator in the United States must meet a minimum efficiency that depends on its fresh food capacity and frozen food capacity, as well as its door type (French or not), the location of its freezer (top or bottom), and whether or not it has through-the-door ice.

⁵Sneeringer and Key (2011) find evidence that firms downsize livestock operations in order to stay under a size threshold that leads to milder federal regulation. The Clean Air Act, which restricts economic activity within counties that have exceeded pollution limits, exempts sufficiently small firms. Becker and Henderson (2000) find that firms locate in unregulated counties, but they do not find evidence that firms manipulate size to avoid the regulation. Gao, Wu, and Zimmerman (2009) find suggestive evidence that firms reduced size in order to avoid security regulations imposed by the Sarbanes-Oxley bill. A number of theoretical models have been developed to explain when exempting small firms for taxation or regulation might be optimal. These results are driven either by the existence of per-firm administrative costs, as in Dharmapala, Slemrod, and Wilson (2011), or by the inability of the planner to directly tax externalities, as in Brock and Evans (1985) and Kaplow (2013). These conditions differ critically from our setting, which is free of administrative costs and allows the planner to directly tax the externality.

Anderson and Sallee (2011), Jacobsen (2013a), and Whitefoot, Fowlie, and Skerlos (2013), but these papers study CAFE before the introduction of attribute-basing. Thus, we contribute to the broader CAFE by adding to the small literature that does consider attribute-basing. That literature includes Whitefoot and Skerlos (2012), which uses 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. That study concludes that an increase in footprint will likely be a major source of adjustment. Gillingham (2013) discusses the implicit incentive for the expansion of footprint in a broader discussion of CAFE policies. Jacobsen (2013b) addresses the safety impacts of footprint-based standards in the United States, which we touch on below. Our work also 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), as well as to the broader bunching literature cited above, which is surveyed in Slemrod (2010). Our panel analysis differs from existing work in this area by considering a double notch (i.e., a notch in two coordinates), which is, to the best of our knowledge, new to the literature.

2 Theory

The goal of our theoretical model is to characterize the welfare consequences of attribute-based regulation and provide theoretical predictions that motivate our empirical analysis. We define an ABR as any policy, denoted $S(a, e)$, that ostensibly targets some characteristic, e , but which takes another characteristic, a , into account when determining policy treatment.⁶ For example, CAFE targets fuel economy (e), but the CAFE fine structure ($S(a, e)$) depends on footprint (a). Our model presumes that regulation is motivated by an externality that is directly caused by the characteristic e , but we discuss below the implications of the realistic alternative in which e is imperfectly correlated with the externality in question (e.g., fuel economy itself does not cause an externality; greenhouse gas emissions do).

For energy efficiency policies, ABR mandates some minimum efficiency level that each product (or some group of products jointly) must obtain in order to be sold or to avoid fines. We write this

⁶We use the term “ostensibly” because the very nature of ABR is that it creates implicit incentives for a as well as e , so that one might say it is a regulation pertaining to both characteristics. This distinction is thus somewhat semantic, but it is nevertheless clear in the cases that we have in mind—e.g., fuel-economy regulations are intended to regulate fuel economy rather than size.

version of ABR as an inequality constraint, $e \geq \sigma(a) + \kappa$, where $\sigma(a)$ is some function of a that determines what minimum efficiency is required of a particular product and κ is a constant. We call $\sigma(a) + \kappa$ the *target function*, and we refer to its slope, denoted $\sigma'(a)$, as the policy's *attribute slope*. The solid line in Figure 1 is the target function for CAFE, and its slope is CAFE's attribute slope, which is negative, so that $\sigma'(a) < 0$. The function $\sigma(a)$ is general, so it could of course include a constant, which obviates the need to write κ separately. But, we assume that $\sigma(0) = 0$ and add the constant term κ for expositional reasons; much of our analysis concerns cases where $\sigma'(a) = 0$ for all values of a , in which case the policy reduces strictly to the choice of the constant κ .

Consumers have unit demand for a durable good with variable characteristics a and e . They receive flow utility $F_n(a_n, e_n)$ from the durable, where $n = 1, \dots, N$ indexes different types of consumers. Each type may have a distinct flow utility function F_n , which allows for taste heterogeneity, but all consumer types face the same price schedule, denoted $P(a_n, e_n)$. We assume that a and e are both goods with decreasing marginal utility; F_n is increasing and concave in each argument. For expositional ease, we assume that there are equal numbers of each type n and that the planner puts equal weight on all utilities, which obviates the need for weighting notation.

Consumers have exogenous income I_n , which they spend on the durable and a numeraire x . The numeraire has unit price, and, for simplicity, we assume that utility U is quasilinear in the durable and the numeraire and that the numeraire provides a constant marginal utility, which is normalized to 1. The consumer's payoff is written as $U_n(a_n, e_n) = F_n(a_n, e_n) + x_n$. When the consumer faces a regulatory policy that enforces a minimum, they choose a_n , e_n and x_n so as to maximize U_n subject to their budget constraint ($I_n \geq P(a_n, e_n) + x_n$) and the regulatory constraint ($e_n \geq \sigma(a_n) + \kappa$).

On the supply side, we assume perfect competition. Specifically, we assume that there is a constant marginal cost of the quantity produced of a good with attribute bundle (a, e) , which we denote as $C(a, e)$, and that there are no fixed costs or barriers to entry. This means that (in the absence of policy) a consumer will be able to purchase any bundle of attributes (a, e) and pay $P(a, e) = C(a, e)$. Furthermore, because there are no profits in equilibrium due to free entry and the lack of fixed costs, firms will make no profits; thus consumer surplus will be a sufficient welfare statistic. We assume that $C(a, e)$ is rising and convex in both attributes. The assumption of

perfect competition abstracts from a variety of important issues, and we discuss them below after establishing our baseline results.

The planner's problem differs from the individual's in summing across the different types of consumers and in recognizing the existence of an externality. We model the external benefits (or damages) as depending only on the sum total of e chosen by the consumers. This implies that benefits (or damages) depend directly on the characteristic e , and thus policy can directly target the externality. This assumption matters for policy design (it implies that the first-best is obtainable in some cases), and we discuss alternatives below. Note that our assumptions about the externality also implies that the externality does not depend on *who* chooses e .⁷ For expositional ease, we assume that total external damages are linear in aggregate e , denoted as $\phi \sum_{n=1}^N e_n$, with ϕ being the marginal external effect of e . For energy efficiency, ϕ will be positive.

Before proceeding to theoretical results, it is useful to write out the consumer's and planner's optimization problems and compare their respective first-order conditions. The planner's objective is to maximize social welfare W , which is the simple sum of private utilities over types, plus the externality:

$$\max_{a_n, e_n} W = \sum_{n=1}^N \{F_n(a_n, e_n) - C(a_n, e_n) + I_n\} + \phi \sum_{n=1}^N e_n. \quad (1)$$

The first-order conditions of equation 1, which characterize the first-best allocation of resources, require (1) that the marginal flow utility of a_n equals the marginal cost of producing a_n for each consumer type ($\frac{\partial F_n}{\partial a_n} = \frac{\partial C}{\partial a_n} \forall n = 1, \dots, N$) and (2) that the marginal flow utility of e_n is less than the marginal cost of producing e_n , by the size of external benefits ϕ , for each consumer type ($-\phi = \frac{\partial F_1}{\partial e_1} - \frac{\partial C}{\partial e_1} = \dots = \frac{\partial F_N}{\partial e_N} - \frac{\partial C}{\partial e_N}$). These are standard conditions. Because there is no externality associated with a_n , its allocation is chosen irrespective of the externality in the first-best. The first-best allocation of e_n involves a wedge, equal to marginal external benefits, which is common across all products. In the terminology of pollution mitigation, $\frac{\partial F_n}{\partial e_n} - \frac{\partial C}{\partial e_n}$ is the marginal cost of abatement for type n . At the first best, these marginal costs must be equalized across all types.

For the consumer, we substitute the budget constraint into the optimization problem to eliminate the numeraire, substitute $P(a, e) = C(a, e)$ based on our supply side assumption, and write

⁷Introducing heterogeneous damages across consumers would cause our baseline results to have elasticity-weighted second-best tax rates, as in Diamond (1973).

type n 's Lagrangean as:

$$\max_{a_n, e_n} \mathcal{L}_n = F_n(a_n, e_n) - C(a_n, e_n) + I_n + \lambda_n \times (e_n - \sigma(a_n) - \kappa), \quad (2)$$

where λ_n is the (endogenously determined) shadow price of the regulation for type n . Consumer n 's first-order conditions are:

$$\frac{\partial F_n}{\partial a_n} - \frac{\partial C}{\partial a_n} = \lambda_n \sigma'(a_n) \quad (3)$$

$$\frac{\partial F_n}{\partial e_n} - \frac{\partial C}{\partial e_n} = -\lambda_n. \quad (4)$$

Note that λ_n may be zero for some or all types if the regulation is not binding. Moreover, it will generally vary across types. These conditions bring to light the basic economics of ABR. ABR creates a pair of wedges between marginal utility and marginal cost for a and e ; it acts just like a pair of taxes (or subsidies) in the “markets” for a and e . These wedges are proportional to each other, differing by a factor of $-\sigma'(a)$.

The first-best allocation implies no wedge in the market for a . For the consumer, this can only be true if λ_n or $\sigma'(a_n)$ is zero. If $\sigma'(a_n) = 0$ (i.e., there is no attribute basing) then every consumers' first-order condition for a_n will match the planner's, but their first-order conditions for e_n will generally not meet the conditions required of efficiency. The reason is that the endogenously determined λ_n will generally vary across consumer types. For some, λ_n may be zero because their preferences lead them to choose a bundle that complies with the standard, even in the absence of policy. For others, λ_n will be positive, and it will differ across individuals. Again using the terminology of pollution mitigation, this is equivalent to the case where marginal costs of abatement differ across sources.

Under some conditions, ABR can increase a policy's efficiency by equalizing marginal costs of generating (mitigating) the externality. It does so, however, by creating a distortion in the choice of a , so that the benefits of marginal cost equalization need to be balanced against the distortions to the attribute in an optimally designed ABR. We show this result below, but before doing so, we first emphasize that ABR creates a distortion without generating any benefits in certain circumstances, namely when a regulation has a *compliance trading system*. This is relevant because real-world

attribute-based policies, like CAFE, have such a trading system.

2.1 Policies with compliance trading

Many policies, including CAFE since 2012, have a compliance trading system in which a product or firm that exceeds the standard is issued a credit for its “excess compliance”. An out-of-compliance firm can buy this credit and use it to comply with the regulation. Such a system implies that regulations bind only on the market as a whole, rather than on the individual products or firms. If the market for these credits is competitive, then the shadow price of the regulation will be uniform across firms and equal to the equilibrium price of a credit.⁸ We use λ , without a subscript, to denote the market-wide shadow price of complying with the regulation when there is trading.

When there is compliance trading, the potential benefit that ABR provides by equalizing marginal costs of compliance is obviated, which leads to Proposition 1.

Proposition 1. *The first-best allocation can be achieved for a regulation with compliance trading by setting $\sigma'(a) = 0$ and choosing κ so that $\lambda = \phi$.*

The proof of Proposition 1 (along with all other proofs) is in the appendix. Briefly, the proof shows that when $\lambda = \phi$ and $\sigma'(a) = 0$, each consumer’s optimality conditions match the planner’s exactly. This result is not surprising, but it cuts to the heart of an evaluation of ABR. It says that the marginal incentive (the shadow price) to e should equal marginal benefits, just as in a Pigouvian tax, and that the attribute-slope should be zero—i.e., there should be no attribute basing.

Proposition 1 can be recast as a tax (or subsidy, in this case) policy, rather than a regulation, which makes the link between our results and the optimal tax literature more transparent. Instead of a regulation, suppose that the attribute-based policy $S(a, e)$ is a subsidy equal to $s \times (e - \sigma(a))$ and that all needed revenue is collected lump-sum equally from each consumer in a lump-sum tax G . The consumer’s problem is then $\max_{a_n, e_n} F_n(a_n, e_n) - P(a_n, e_n) + I_n - G + s \times (e_n - \sigma(a_n) - \kappa)$, which will produce exactly the same choices as the regulation with compliance trading whenever

⁸Formally, suppose that market credits m are available at price P_m and can be used for compliance so that consumer n ’s constrained optimization problem becomes $\max_{a_n, e_n, m_n} F_n(a_n, e_n) - P(a_n, e_n) + I_n - P_m m_n + \lambda_n \times (e_n + m_n - \sigma(a_n) - \kappa)$. A consumer for whom the private cost of increasing e exceeds the market price will buy permits and reduce e , and one for whom the private cost of increasing e is less than P_m will increase their own e and sell the excess credits. In equilibrium, the shadow price λ_n for all consumers will converge to a common value, equal to P_m , which we denote as λ without a subscript. The consumer’s condition can then be rewritten to eliminate the permits (because they cancel out) and substituting a shadow price that, from the point of view of the individual, is fixed (rather than endogenous): $\max_{a_n, e_n} F_n(a_n, e_n) - P(a_n, e_n) + I_n + \lambda \times (e_n - \sigma(a_n) - \kappa)$.

$s = \lambda$ and the same target function $\sigma(a)$ is used.⁹ This leads to a corollary of Proposition 1 for a subsidy.

Corollary 1. *The first-best allocation can be achieved with a subsidy by setting $\sigma'(a) = 0$ and $s = \phi$.*

This is the same result as Proposition 1—attribute-basing is not part of the first-best solution, which features a Pigouvian subsidy—but it helps tie ABR to the literature. The “additivity property” of optimal taxation in the presence of externalities establishes that (a) the optimal tax on a commodity that produces an externality is equal to the optimal tax on that good if there were no externality, plus marginal external damages; and (b) the externality does not change the optimal tax on other goods, even if they are substitutes or complements to the externality-generating good (Kopczuk 2003). In our context, where optimal taxes are zero in the absence of the externality, this implies that the subsidy on e is just marginal external benefits and that there should be no subsidy (or tax) on a , regardless of the complementarity or substitutability of a and e .¹⁰

The subsidy version of the problem also clarifies the connection between ABR and Harberger triangles. Consider a subsidy with $s = \phi$ and $\sigma'(a) \neq 0$, which produces an undesired subsidy to a that creates deadweight loss. A “partial equilibrium” approximation (that is, one that ignores the way in which a wedge in the choice of a distorts the choice of e by shifting its costs and benefits) of this deadweight loss is the Harberger triangle, which is $1/2 \cdot \frac{\partial a}{\partial s\sigma'(a)}(s\sigma'(a))^2$. As in a standard tax setting, this approximation suggests that the deadweight loss grows with the square of the tax wedge ($s\sigma'(a)$). In addition, the derivative indicates that the deadweight loss will be larger when a is more responsive to attribute-based policy incentives. This provides an important implication for our empirical analysis—the responsiveness of a to ABR is a key parameter that determines the

⁹We assume the same normalization that $\sigma(0) = 0$, and we drop κ from our formulation of the subsidy. This is immaterial, and is done purely for exposition. For the regulation, κ is a useful parameter because raising or lowering it will change λ while holding $\sigma(a)$ fixed. In the subsidy case, where s is set directly, the κ term is superfluous as it functions as a lump-sum transfer and we abstract from income effects.

¹⁰Kopczuk (2003) shows that the additivity property holds quite generally, even when the first-best is not obtainable. The only requirement for additivity to hold is that the externality-generating good be directly taxable. Thus, our “no attribute-basing” result would hold if we introduced revenue requirements, income taxation, other goods, equity concerns, etc., so long as direct targeting of e were possible. (In such second-best settings, the optimal tax on a will not generally be zero, but this will not be because of the externality—i.e., changing the size of the externality will not change the optimal tax on a .) When the externality cannot be directly targeted, this result requires modification because, as is generally true in second-best situations, taxation of all goods will be optimal. We argue below, however, that real-world attribute-based policies cannot be justified on these grounds and are not, in practice, motivated by these concerns.

the welfare implications of ABR.

2.2 Policies with compliance trading and exogenous policy parameters

Under compliance trading, there should be no attribute-basing. But, policy makers may sometimes be constrained in their choice of policy parameters (perhaps because of political considerations). To see how constraints on the planner might influence attribute-based policy designs, we derive second-best parameter values (denoted with a superscript SB) under the assumption that the planner takes either the tax rate s or the attribute slope $\sigma'(a)$ as given.

For simplicity, we restrict attention to cases where the attribute slope is linear, so that $\sigma(a) = \hat{\sigma}a$ where $\hat{\sigma} = \sigma'(a)$ is a constant. We show results here for the subsidy rather than a regulation for ease of exposition—the subsidy rate s is a direct choice variable in the subsidy policy, whereas λ is a function of κ —but there are exact analogs for all results for a regulation whenever $\lambda = s$. The linear subsidy is written as $S(a, e) = s \times (e - \hat{\sigma}a)$.

Proposition 2 derives the second-best attribute slope ($\hat{\sigma}^{SB}$) when the subsidy rate s is taken as fixed.

Proposition 2. *For a linear subsidy, when s is fixed, the second-best attribute slope is:*

$$\hat{\sigma}^{SB} = \frac{s - \phi \left(\sum_n \frac{\partial e}{\partial \hat{\sigma}} \right) / n}{s \left(\sum_n \frac{\partial a}{\partial \hat{\sigma}} \right) / n}.$$

When the subsidy is not set at marginal benefits, the optimal attribute-slope will not be zero.¹¹ Instead it is the product of two terms, the first of which is the percentage difference between marginal benefits and the subsidy (i.e., how “wrong” is the subsidy), and the second of which is a ratio of derivatives of a and e , averaged over consumer types. The second term indicates what portion of the market’s response to a change in the attribute slope will be in e (which boosts the externality) versus in a (which creates distortion). When this is small (response is largely in a), $\hat{\sigma}^{SB}$ will be small. This is intuitive; if the attribute slope mainly induces a change in a , it will be a poor substitute for a direct tax on e . This formula motivates our empirical analysis—the market’s

¹¹Note that $\hat{\sigma}^{SB}$ does collapse to zero when $s = \phi$ because the first term will be zero.

responses of a and e to ABR are key parameters to determine the second-best attribute slope, and we estimate them below.

If attribute-basing in fuel economy policies was being used as a way to overcome standards that are too weak on average, then Proposition 2 suggests that the second-best attribute slopes would likely be *upward* sloping ($\hat{\sigma}^{SB} > 0$), whereas real-world fuel-economy ABR all have a *downward* slope ($\hat{\sigma}^{SB} < 0$). If a policy for a positive externality is too weak, then ($0 < s < \phi$), which makes the first term negative. Increasing $\hat{\sigma}$ (which means “flattening” the negative slope) makes the implicit subsidy for a smaller, which will reduce a ; thus $\frac{\partial a_n}{\partial \hat{\sigma}} < 0$. Increasing $\hat{\sigma}$ affects the choice of e indirectly by shifting a . If a and e are gross substitutes (which is uncertain, but seems likely), then $\frac{\partial e_n}{\partial \hat{\sigma}} > 0$.¹² This makes $\hat{\sigma}^{SB} > 0$, which suggests that size should be penalized, not rewarded, if attribute-basing is meant to compensate for a fuel economy policy that is too weak.

Proposition 3 describes the opposite case. It expresses the second-best subsidy rate when the attribute slope $\hat{\sigma}$ is taken as fixed.

Proposition 3. *For a linear subsidy, when $\hat{\sigma}$ is fixed, the second-best subsidy rate is:*

$$s^{SB} = \frac{\phi}{1 - \hat{\sigma} \left(\frac{(\sum_n \frac{\partial a}{\partial s})/n}{(\sum_n \frac{\partial e}{\partial s})/n} \right)} = \frac{\phi}{1 - \epsilon_a^a \epsilon_s^e},$$

where ϵ_x^y denotes the elasticity of y with respect to x for each pair of variables.

Proposition 3 shows that, in the presence of a non-zero attribute slope, the optimal corrective subsidy on e is not just marginal benefits (ϕ), as it is in the first best.¹³ Instead, it is marginal benefits divided by one minus the attribute slope times a ratio of derivatives, averaged over consumer types. These derivatives capture how a and e respond to a change in the subsidy rate. The second formulation rewrites these derivatives as elasticities.

This formula implies that the market’s responsiveness of a and e to ABR (ϵ_s^a and ϵ_s^e) are key statistics to determine the second-best subsidy rate. Our empirical analysis provides actual

¹²We think the gross substitute case is the most likely one, but it ultimately depends on the relative cross-partial derivatives of F_n and C . We expect that $\frac{\partial^2 C}{\partial a \partial e} > 0$ for energy efficiency and size—i.e., the cost of raising fuel economy for a larger vehicle is larger. For utility, $\frac{\partial^2 F}{\partial a \partial e}$ is plausibly zero, or small, because the value of energy efficiency (which saves fuel cost) is largely independent of vehicle size.

¹³Note that s^{SB} does collapse to the first best if $\hat{\sigma} = 0$ or if a is perfectly inelastic because the second term in the denominator will be zero.

estimates of ϵ_s^a and ϵ_s^e for one case, but it is worth considering whether the formula's denominator is likely to be larger or smaller than one in general. In most real-world ABRs, the attribute-slope is negative ($\epsilon_a^\sigma < 0$). A higher s increases the subsidy to both a and e , so that both sets of derivatives will generally be positive ($\epsilon_s^a > 0$ and $\epsilon_s^e > 0$). Then, the denominator of s^{SB} will be greater than one, so $s^{SB} < \phi$. This means that attribute-basing *attenuates* corrective policy and the second-best subsidy rate is smaller than the Pigouvian benchmark; the subsidy is below marginal benefits. The reason is that the subsidy induces distortions in a . The cost of increasing e by one unit is therefore not only the difference between willingness to pay and the production cost of e , but also the marginal distortion induced in the choice of a . As the market's response is tilted towards an increase in a (rather than e), the distortionary cost will be greater, the ratio of derivatives will be larger, and the optimal subsidy will be lower. In the extreme, if firms respond to a change in the subsidy exclusively by manipulating a , the subsidy will be zero.

2.3 Policies without compliance trading

We next derive results for the case where there is no compliance trading. In the absence of compliance trading, a flat standard will impose different burdens on different products and will lead to an inefficient dispersion in the marginal costs of compliance. Attribute-basing can improve the efficiency of such a policy by (partially) equalizing marginal costs of compliance across firms or products, so long as the attribute is correlated with marginal compliance costs. But, attribute-basing is not an efficient substitute for compliance trading because, unless the attribute perfectly predicts compliance costs, ABR will not fully equalize compliance costs, and also because it equalizes compliance costs by creating a distortionary incentive to change the attribute. The latter consideration causes the optimal degree of attribute-basing to stop short of harmonizing marginal costs of compliance as fully as possible, instead trading off some degree of harmonization against distortions induced in the attribute.

In the appendix, we show how to generalize the second-best formulas in Propositions 2 and 3 for the linear attribute-based policy (i.e., $\sigma'(a) = \hat{\sigma}$) when there is no compliance trading. Rather than focusing on these results, which express the optimal policies as functions of the covariance between marginal costs of compliance and the derivatives of a and e with respect to κ and $\hat{\sigma}$, we provide a different result in Proposition 4 that expresses the optimal attribute slope for the linear policy as a

function of the correlation between marginal compliance costs and the level of the attribute a . We denote the optimal attribute slope in the no compliance case with superscript NC , and we denote variable means with a bar.

Proposition 4. *When there is no compliance trading, and the constraint binds for all n , the optimal attribute slope $\hat{\sigma}^{NC}$ in the linear regulation is:*

$$\hat{\sigma}^{NC} = \frac{\frac{\sum_n (\lambda_n - \bar{\lambda})(a_n - \bar{a})}{n}}{\phi \left(\frac{\sum_n \frac{\partial a_n}{\partial \hat{\sigma}}}{n} - \bar{a}_n \frac{\sum_n \frac{\partial a_n}{\partial \kappa}}{n} \right)},$$

which is not zero unless a_n is perfectly uncorrelated with λ_n under a flat standard.

In contrast to the result under compliance trading (Proposition 1), this proposition shows that some attribute-basing is optimal ($\hat{\sigma}^{NC} \neq 0$). The exception is when the attribute is perfectly uncorrelated with marginal compliance costs (under a flat standard), in which case the numerator is zero, and attribute-basing, which cannot equalize compliance costs, is undesirable.

Consider the implications of Proposition 4 for the case of fuel economy regulation, where e is fuel economy and a is vehicle weight, as in Japan. The denominator of $\hat{\sigma}^{NC}$ represents how much a responds to changes in the regulation. Raising $\hat{\sigma}$ (flattening the negative attribute slope) lowers the implicit subsidy to weight, so we expect (and find empirically below) that $\frac{\partial a_n}{\partial \hat{\sigma}} < 0$. Raising κ (tightening the standard overall) raises the implicit subsidy to weight (by raising the shadow price), so we expect (and find empirically below) that $\frac{\partial a_n}{\partial \kappa} > 0$. Thus, the denominator will be negative. The numerator of $\hat{\sigma}^{NC}$ is the sample covariance between the marginal cost of improving fuel economy (which equals the shadow price) and vehicle weight. This covariance will be positive as long as the marginal costs of improving fuel economy are rising.¹⁴ This implies that $\hat{\sigma}^{NC} < 0$ —heavier vehicles must meet a lower standard—as is the case in real world policies.

Under a flat standard, heavier vehicles will have a higher marginal cost of compliance because they will be forced to improve fuel economy more than lighter vehicles. As the attribute-slope becomes more negative, the correlation between weight and the shadow price will get smaller. If the slope were sufficiently steep, this correlation could be driven to zero. This cannot be a solution

¹⁴Note that the formula for $\hat{\sigma}^{NC}$ is not in closed form. The shadow price λ_n is endogenous, so its correlation with a_n is also endogenous (as are the derivatives of a_n with respect to the parameters). In this example, that the correlation will be negative can be shown by contradiction.

to $\hat{\sigma}$, however, because if the denominator became zero, it would imply no attribute-basing. Rather, the planner stops short of fully equalizing marginal compliance costs, so heavier vehicles will still have higher compliance burdens under the attribute-based policy, and this will be more true as the attribute is more elastic (i.e., the denominator is larger).

In general, the proposition shows that the planner must balance two competing factors when determining the optimal degree of attribute-basing.¹⁵ The first factor is the extent to which ABR can equalize marginal compliance costs. The formula’s numerator indicates that the planner will make the slope steeper when there is a strong correlation between the attribute and marginal compliance costs. The second factor is the extent to which ABR induces distortions in the attribute. The formula’s denominator indicates that the slope should be flatter when market responses in a to the ABR are strong. An optimizing planner will thus use ABR to equalize compliance costs, but not fully; they stop short when the increases in the distortion equal the benefits of increase harmonization. Below, we estimate the market response of a and e to an ABR and relate our findings back to these theoretical results.

2.4 Discussion of modeling assumptions

Our theory considers policies where $\sigma(a)$ is smooth (differentiable everywhere), but below we empirically analyze a notched policy, in which $\sigma(a)$ is a step function. The welfare implications of attribute-basing are quite similar for notched policies, which we discuss further below were relevant as well as in an appendix. We presented the smooth case believing it to be more general and intuitive.

Our model abstracts from several factors that might make attribute-based regulation desirable, even when compliance trading is available. First, consider incidence. ABR can smooth the *average* (as opposed to marginal) compliance burden across segments of the market, or it might be useful in tilting the burden in favor of preferred producers and consumers. For example, some argue that CAFE introduced footprint-based standards in order to favor the Detroit automakers (who have an advantage in large vehicles) and their consumers relative to their Asian rivals. In Europe, the

¹⁵An interesting benchmark is the special case when the attribute a is exogenous (so that there is no distortion caused by attribute-basing and one of the competing forces is removed), and the cost of compliance for each product is equal to $(e_n - e_n^o)^2$, where e_n^o is the optimal e_n in the absence of policy. In that case, the optimal $\hat{\sigma}$ is the same as the Ordinary Least Squares solution to fitting a line between e_n and a_n . This is, in fact, how attribute slopes in both the United States and Japan are determined, which would be optimal only if there were no distortion in a .

regulation is said to have been designed to favor German automakers (whose vehicles are relatively large) vis-à-vis French and Italian producers. If distribution is the motivation for ABR, then the distortions it induces represent the costs of achieving distributional goals.

Second, our model assumes that the externality is directly related to e , which can be targeted directly. In reality, because they do not directly target externalities, it is well known that energy-efficiency policies are imperfect instruments (as are most, if not all, policies targeting externalities). Fuel-economy policies, for example, induce a rebound effect (additional travel induced by a lower cost per mile of travel), fail to create proper incentives in the used car market, and are inefficient when lifetime mileage varies across consumers.¹⁶ Attribute-based energy efficiency policies will inherit these inefficiencies.¹⁷ Truly first-best solutions will thus not be obtainable, even with compliance trading. Our focus, however, is not on determining whether or not attribute-based energy-efficiency policies are perfectly efficient, but rather on whether or not they can improve welfare as compared to energy-efficiency policies that do not employ attribute-basing.

The additional flexibility afforded by attribute-basing *could* ameliorate some of the problems created by energy-efficiency policies, but we are doubtful that this is the case for attribute-based policies as they are currently practiced. Consider the case of greenhouse gas emissions from automobile emissions, which are proportional to gallons consumed. A direct tax on gasoline would be efficient. A tax (or an implicit tax achieved through regulation) on fuel economy is inefficient if vehicles are driven different lifetimes miles, and thus create different total emissions. Heavier vehicles tend to have longer lives and thus consume more gasoline than smaller vehicles, even accounting for differences in their rates of fuel consumption (Lu 2006). Attribute-basing could offer efficiency improvements if it placed a higher tax *per unit of fuel consumption* on heavier vehicles; that is, attribute-basing would be useful in creating a tax on energy efficiency that varied with the attribute, which is not how existing policies are designed.

Fuel-economy policies are also inefficient because they induce a rebound effect. If the rebound effect is correlated with vehicle size (or another attribute), then attribute-based regulation could

¹⁶See Anderson, Parry, Sallee, and Fischer (2011) for a recent review of CAFE that describes inefficiencies of fuel-economy policy. See Borenstein (Forthcoming) for a recent discussion of the rebound effect induced by energy-efficiency policies. See Jacobsen and van Benthem (2013) for an exploration of effects on used durables markets.

¹⁷A substantial literature considers whether or not such policies can be efficiency enhancing when consumers improperly value energy efficiency. Fischer, Harrington, and Parry (2007), Allcott, Mullainathan, and Taubinsky (2014) and Heutel (2011) provide theoretical treatments, and Busse, Knittel, and Zettelmeyer (2013) and Allcott and Wozny (2012) provide key empirical tests.

improve the efficiency of a policy by shifting compliance towards vehicles that have less elastic mileage. There is evidence on heterogeneity in the rebound effect derived from micro data on how vehicle miles traveled responds to fuel prices in Gillingham (2011) and Knittel and Sandler (2012), but this data was not used in the design of policies. Instead, the attribute-slope was drawn to fit the observed correlations between fuel economy and footprint or weight.

Without direct targeting, the additivity property of corrective policies will not hold (Kopczuk 2003), and the flexibility afforded by attribute-basing will have some efficiency value for the reason that, in general, all instruments are useful in second-best settings.¹⁸ Moreover, attribute-based energy efficiency policies might be useful in fixing some of the problems associated with energy-efficiency policies, but we know of no evidence that such considerations have actually been part of the policy-design process.

Third, we assume perfect competition, which abstracts from several issues. A market with imperfect competition among differentiated products will not necessarily offer the set of goods (attribute bundles) that maximizes social welfare, nor will prices necessarily maximize social welfare conditional on the set of products on offer. Thus, a cleverly-designed ABR *might* be able to improve social welfare by changing the mix of products or shifting equilibrium prices. It is not, however, obvious under what conditions this would be true, nor do we know of any evidence that actual policies have been designed with such market inefficiencies in mind. Thus, we consider this an issue of great theoretical interest, but we are less sanguine about its practical import for explaining real-world policies.¹⁹

One way in which ABR might be helpful in practice is in fixing distortions to the “extensive

¹⁸There is a vast literature on second-best tax policies. Particularly relevant to the study of automobile emissions are Fullerton and West (2002) and Fullerton and West (2010), which derive second-best taxes on several vehicle attributes in an attempt to mimic an efficient (direct) tax on emissions. They show that counterintuitive results, such as a subsidy on vehicle size, can arise because of the interaction of multiple tax instruments on correlated attributes and multiple margins of adjustment, including not just the choice of vehicle but also the choice of miles traveled.

¹⁹To see the practical challenges involved, consider the intuition of the Spence (1975) model of quality pricing applied to automobiles, where vehicle size takes the place of quality. If the marginal consumer of a product valued size less than price, then an ABR that implicitly subsidized size might make prices more efficient. But, with many products, each good will have a different marginal consumer (or distribution of marginal consumers), so efficient wedges will differ across products. Or, note that markups are likely to be correlated with regulated attributes in many markets. Heavier vehicles, for example, are thought to have greater per unit profits than lighter vehicles (though this may itself be the outcome of fuel economy regulations). This might suggest that too few heavy vehicles are being sold, so an implicit subsidy to size through an ABR would be desirable. But, this result ultimately hinges on the form of competition, the reasons that markups differ initially (which may be hard to determine), and the endogenous response of not just prices, but attribute bundles to the policy.

margin” of a market. As Holland, Hughes, and Knittel (2009) demonstrate for the case of low-carbon fuel standards in California, an average emissions standard creates a relative subsidy for products above the standard and a relative tax for products below the standard. Depending on parameters, this could raise or lower the average price level and cause the overall product market to expand or contract. Provided that cars create negative externalities relative to the outside good, the optimal policy should shrink the car market. If a flat standard happened to create an expansion of the market, it is possible that an attribute-based standard could reverse this incentive and shrink the car market by providing a tax on all products.²⁰

Fourth and finally, advocates of attribute-basing in car markets have argued that it promotes technology adoption.²¹ Roughly speaking, automakers can comply with a flat standard by downsizing their fleet or by adopting new technologies. Attribute-based policies can be designed to limit opportunities for downsizing, which forces compliance to come from technology. If there are spillovers between companies from adopting new technologies, there might be some justification for attribute-basing. In this case, the distortions in the attribute represent a cost of spurring technology. Direct subsidies to new technology adoption may thus be more efficient.


3 Identifying Attribute Distortions through Bunching Analysis

Our theory predicts that attribute-basing will lead to distortions in the choice of the secondary attribute. In this section, we test that prediction by analyzing the distribution of the secondary attribute for the case of Japanese fuel-economy standards. The Japanese regulation has several advantages from the point of view of identification. First, the regulation features “notches”; the fuel-economy target function in Japan is a downward-sloping *step function* in vehicle weight. These notches provide substantial variation in regulatory incentives and allow us to use empirical methods developed for the study of nonlinear taxation (Saez 1999, 2010; Chetty, Friedman, Olsen, and

²⁰Note that Jacobsen and van Benthem (2013) suggest that the benefits resulting from higher new car prices will be partly offset by changes in vehicle scrappage.

²¹Advocates have also argued that attribute-basing promotes safety by promoting larger cars. This appears to be based on a misunderstanding of safety-related externalities. Larger cars are safer for the car’s occupants (which is a private benefit and should be priced into the car), but they are more dangerous to those outside the car (which is an externality). If ABR changes the distribution of sizes of cars, this could affect net safety. See Jacobsen (2013b) for a related model that concludes that footprint-based CAFE is roughly safety neutral.

Figure 2: Fuel-Economy Standards 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.

Pistaferri 2011; Kleven and Waseem 2013).²² Second, the Japanese government has been using attribute-based regulation for decades, and we have more than ten years of data available for analysis.²³ Third, our data span a policy change that enables us to use panel variation, which we pursue in section 4.

3.1 Data and Policy Background

Japanese fuel-economy standards, which were first introduced in 1979 and have changed four times since, are weight-based. Our data, which begin in 2001, span the two most recent policy regimes. The target functions for these policies are shown in Figure 2. To be in compliance with the regulation, firms must have a sales-weighted average fuel economy that exceeds the sales-weighted average target of their vehicles, given their weights, and there is no trading of compliance credits

²²Our theory concerns smooth attribute-based policies, but the welfare implications of notched ABR are the same, which we illustrate in the appendix.

²³In contrast, the United States just recently instituted attribute-based fuel economy regulation in 2012, which generates little data for analysis. In addition, its attribute-based target function is smooth, making identification more challenging.

across firms although it is one of the key issues in the ongoing policy debate.²⁴ In Japan, firms are required to meet this standard only in the “target year” of a policy. This is different from the U.S. CAFE program, which requires compliance annually.²⁵

When introducing a new standard, the Japanese government selects a set of weight categories (the widths of the steps in Figure 2).²⁶ The height of the standard is then determined by what is called the “front-runner” system. For each weight category, the new standard is set as a percentage improvement over the highest fuel economy vehicle (excluding vehicles with alternative power trains) currently sold in that segment. When the newest standard was introduced in 2009, the government also introduced a separate tax incentive that applies to each specific car model, rather than for a corporate fleet average. We make use of this policy in our panel analysis in section 4. For our bunching analysis, we simply note that both incentives are present in the latter period, and either could be motivating strategic bunching of vehicle weight at the regulatory thresholds (which are common across the two policies).

Our data, which cover all new vehicles sold in Japan from 2001 through 2013, come from the Japanese Ministry of Land, Infrastructure, Transportation, and Tourism (MLIT). The data include 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. Table 1 presents summary statistics. There are between 1,100 and 1,700 different vehicle configurations sold in the Japanese automobile market each year. This includes both domestic and imported cars. The data are not sales-weighted; we use the vehicle model as our unit of analysis throughout the paper.²⁷

²⁴Technically, this obligation extends to each weight segment separately. However, firms were allowed to apply excess credits from one weight category to offset a deficit in another. Thus, in the end, the policy is functionally equivalent to the U.S. CAFE program, where there is one firm-wide requirement (but no trading across firms).

²⁵This does not mean, however, that firms have no incentive to comply before the target year. Consumers see a car’s fuel economy relative to the standard when buying a car, so compliance may affect sales. Compliance may also be a part of long-run interactions between firms and the government. Our data show clearly that firms react to the standards even before the target year. Also, to be precise, under CAFE firms may do some limited banking and borrowing, so they must meet the standard every year, on average.

²⁶It is not transparent how these weight categories are chosen, but note that they are almost all of regular width, either 250 kilograms in the old standard or 120 kilograms in the new.

²⁷Sales-weighting might be a useful extension for some of our results, but Japanese automobile sales data suffer from a problem common to automotive sales data sets in general, which is that sales data are generally recorded at a notably higher unit of analysis and a different calendar. For example, there will be several different versions of the Toyota Camry recorded in our regulatory data, but industry sources typically record sales only for all versions of the Camry together. In addition, the relevant sales are model year totals, not calendar totals, whereas industry data typically cover calendar time and do not distinguish between, for example, a 2013 Camry and a 2014 Camry that are sold in the same month. In contrast, the dataset used in our analysis provides disaggregated data for each vehicle

Table 1: Summary Statistics

Year	N	Fuel Economy (km/liter)	Vehicle weight (kg)	Displacement (liter)	CO2 (g-CO2/km)
2001	1441	13.53 (4.58)	1241.15 (356.63)	1.84 (0.98)	195.40 (66.72)
2002	1375	13.35 (4.33)	1263.52 (347.00)	1.86 (0.97)	196.72 (66.26)
2003	1178	13.78 (4.53)	1257.15 (356.28)	1.85 (1.03)	191.88 (68.08)
2004	1558	14.20 (4.78)	1255.37 (364.69)	1.82 (1.03)	184.33 (66.67)
2005	1224	13.30 (4.66)	1324.81 (380.62)	2.00 (1.13)	198.14 (71.62)
2006	1286	13.08 (4.59)	1356.56 (391.13)	2.08 (1.17)	201.78 (72.67)
2007	1298	13.24 (4.78)	1369.41 (399.45)	2.09 (1.22)	200.35 (75.07)
2008	1169	13.38 (4.82)	1390.09 (405.77)	2.14 (1.29)	198.58 (76.27)
2009	1264	13.49 (4.93)	1396.40 (413.76)	2.15 (1.30)	197.73 (76.67)
2010	1300	13.50 (5.04)	1428.27 (438.06)	2.21 (1.30)	198.32 (77.34)
2011	1391	13.95 (5.06)	1437.21 (426.23)	2.19 (1.28)	190.15 (71.60)
2012	1541	14.50 (5.21)	1446.50 (411.87)	2.16 (1.24)	182.05 (67.26)
2013	1706	14.43 (5.40)	1476.79 (400.31)	2.24 (1.24)	183.67 (67.37)

Note: This table shows the number of observations, means and standard deviations of variables by year. Data are not sales-weighted.

3.2 Excess Bunching at Weight Notches


The notched attribute-based standards in Japan create incentives for automakers to increase vehicle weight, but only up to specific values. Increasing weight offers no regulatory benefit, unless the increase passes a vehicle over a threshold. Excess mass (“bunching”) in the weight distribution at exactly (or slightly beyond) these thresholds is thus evidence of weight manipulation. Moreover, if automakers are able to choose vehicle weight with precision, then all manipulated vehicles will have a weight exactly at a threshold. In turn, this implies that all vehicles with weights not at a regulatory threshold have not had their weight manipulated in response to the policy.

We begin by presenting histograms of vehicle weight in Figure 3. Panel A shows the histogram of cars sold between 2001 and 2008. In this period, all vehicles were subject to the old fuel economy standards, which is overlaid in the same figure. The figure reveals strong evidence of weight manipulation in response to the policy; there is visible excess mass at the notch points. The magnitude of this bunching is substantial. There are more than double the number of cars at each notch point, as compared to the surrounding weight segments. The histogram bins have a


configuration. For example, our dataset provides information about each version of the 2013 Camry as well as each version of the 2014 Camry.

Figure 3: Fuel-Economy Standard and Histogram of Vehicles

Panel A. Years 2001 to 2008 (Old Fuel Economy Standard Schedule)



Panel B. Years 2009 to 2013 (New Fuel Economy Standard Schedule)



Note: Panel A shows the histogram of vehicles from 2001 to 2008, where all vehicles had the old fuel economy standard. Panel B shows the histogram of vehicles from 2009 to 2013, in which the new fuel-economy standard was introduced.

width of only 10 kilograms (which is the finest measure available in our data), and yet the bunching appears to be isolated to the weight categories immediately at each threshold. This suggests that automakers can manipulate weight very finely.

Panel B shows the corresponding figure for data taken from years when the new standard was in effect. Between Panel A and B, the mass points shifted precisely in accordance with the change in the locations of the notch points. This shift in the location of the distributional anomalies provides further compelling evidence that firms respond to the attributed-based regulation.

In the appendix, we present analogous results for “kei-cars”. Kei-cars are very small cars with engine displacements below 0.66 liters.²⁸ The weight distribution of kei-cars also exhibits bunching at notch points, and the bunching moves over time in accordance with the policy change.

In sum, the raw data provide strong evidence that market actors responded to attribute-based fuel-economy policies in Japan by manipulating vehicle weight, as predicted by theory. In the next section, we use econometric methods to estimate the magnitude of this excess bunching.

3.3 Estimation of Excess Bunching at Notches

Econometric estimation of excess bunching in kinked or notched schedules is relatively new in the economics literature.²⁹ Saez (1999) and Saez (2010) estimate the income elasticity of taxpayers in the United States with respect to income tax rates and EITC schedules by examining excess bunching around kinks in the U.S. personal income tax schedule. Similarly, Chetty, Friedman, Olsen, and Pistaferri (2011) estimate 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) use 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.

²⁸For comparison, no car sold in the United States in 2010 would qualify as a kei-car. The two-seat Smart Car has the smallest displacement of any car in the United States that year, at 1.0 liters. Although kei-cars must comply with the same fuel economy regulations as all other cars, we present our results for them separately because they occupy a unique market segment, have different tax and insurance regulations, and are generally viewed as a distinct product category by Japanese consumers.

²⁹See Slemrod (2010) for a review of this literature.

To estimate the magnitude of the excess bunching, our first step is to estimate the counterfactual distribution as if there were no bunching at the notch points, which parallels the procedure in Chetty, Friedman, Olsen, and Pistaferri (2011). We start by grouping cars into small weight bins (10 kg bins in the application below). For bin j , we denote the number of cars in that bin by c_j and the car weight by w_j . For notches $k = 1, \dots, K$, we create dummy variables d_k that equal one if j is at notch k . (Note that there are several bins on each “step” between notches, which we can denote as $j \in (k - 1, k)$.) We then fit a polynomial of order S to the bin counts in the empirical distribution, *excluding* the data at the notches, by estimating a regression:³⁰

$$c_j = \sum_{s=0}^S \beta_s^0 \cdot (w_j)^s + \sum_{k=1}^K \gamma_k^0 \cdot d_k + \varepsilon_j, \quad (5)$$

where β_s^0 is an initial estimate for the polynomial fit, and γ_k^0 is an initial estimate for a bin fixed effect for notch k . (We refer to these as initial estimates because we will adjust them in a subsequent step.) By including a dummy for each notch, the polynomial is estimated without considering the data at the notches, defined as the 10 kg category starting at the notch. We define an initial estimate of the counterfactual distribution as the predicted values from this regression omitting the contribution of the notch dummies: $\hat{c}_j^0 = \sum_{s=0}^q \hat{\beta}_s^0 \cdot (w_j)^s$. The excess number of cars that locate at the notch relative to this counterfactual density is $\hat{B}_k^0 = c_k - \hat{c}_k^0 = \hat{d}_k^0$.

This simple calculation overestimates B_k because it does not account for the fact that the additional cars at the notch come from elsewhere in the distribution. That is, this measure does not satisfy the constraint that the area under the counterfactual distribution must equal the area under the empirical distribution. To account for this problem, we must shift the counterfactual distribution upward until it satisfies this integration constraint.

The appropriate way to shift the counterfactual distribution depends on where the excess bunching comes from. Our theory indicates that attribute-based fuel economy regulation provides incentives to *increase* car weight—that is, excess bunching should come from the “left”. We assume that this is the case.³¹ We also make the conservative assumption that the bunching observed

³⁰We use $S = 7$ for our empirical estimation below. Our estimates are not sensitive to the choice of S for the range in $S \in [3, 11]$.

³¹The regulation creates an incentive to *increase* car weight in order to bunch at a weight notch because it provides a lower fuel economy target. The regulation may also create an incentive to decrease weight if, for example, decreasing weight mechanically helps improving fuel economy. However, such an incentive is “smooth” over any weight levels in

at a given notch comes only from the adjacent step in the regulatory schedule, which limits the maximum increase in weight. That is, the bunching at notch k comes from bins $j \in (k - 1, k)$.³² In practice, automakers may increase the weight of a car so that it moves more than one weight category. In that case, our procedure will underestimate weight distortions. In this sense, our procedure provides a lower bound on weight manipulation.

In addition, estimation requires that we make some parametric assumption about the distribution of bunching. We make two such assumptions, the first of which follows Chetty, Friedman, Olsen, and Pistaferri (2011), who shift the affected part of the counterfactual distribution uniformly to satisfy the integration constraint. In this approach, we assume that the bunching comes uniformly from the range of $j \in (k - 1, k)$. We define the counterfactual distribution $\hat{c}_j = \sum_{s=0}^q \hat{\beta}_s \cdot (w_j)^s$ as the fitted values from the regression:

$$c_j + \sum_{k=1}^K \alpha_{kj} \cdot \hat{B}_k = \sum_{s=0}^S \beta_s \cdot (w_j)^s + \sum_{k=1}^K \gamma_k \cdot d_k + \varepsilon_j, \quad (6)$$

where $\hat{B}_k = c_k - \hat{c}_k = \hat{d}_k$ is the excess number of cars at the notch implied by this counterfactual. The left hand side of this equation implies that we shift c_j by $\sum_{k=1}^K \alpha_{kj} \cdot \hat{B}_k$ to satisfy the integration constraint. The uniform assumption implies that we assign $\alpha_{kj} = \frac{c_j}{\sum_{j \in (k-1, k)} c_j}$ for $j \in (k - 1, k)$ and $= 0$ for $j \notin (k - 1, k)$. Because \hat{B}_k is a function of $\tilde{\beta}_s$, the dependent variable in this regression depends on the estimates of $\tilde{\beta}_s$. We therefore estimate this regression by iteration, recomputing \hat{B}_k using the estimated $\tilde{\beta}_s$ until we reach a fixed point. The bootstrapped standard errors that we describe below adjust for this iterative estimation procedure.

The uniform assumption may underestimate or overestimate Δw if the bunching comes disproportionately from the “left” or the “right” portion of $j \in (k - 1, k)$. For example, if most of the excess mass comes from the bins near k , rather than the bins near $k - 1$, the uniform assumption will overestimate Δw . In practice, this appears to be a minor concern, because the empirical distribution in Figure 3 shows that there are no obvious holes in the distribution, which suggests

the sense that vehicles at anywhere in the weight distribution have this incentive, and therefore, the incentive does not create bunching at the notches.

³²For notch $k = 1$ (the first notch point), we use the lowest weight in the data as the minimum weight for this range. Note that this approach may underestimate the change in weight, because the minimum weight in the counterfactual distribution can be lower than the minimum weight in the observed distribution if the attribute-based regulation shifted the minimum weight upward. We want to use this approach to keep our estimate of the change in weight biased towards zero.

that the uniform assumption is reasonable. There should be clear holes if the origins of the excess bunching are substantially disproportional to the distribution. However, we prefer an approach that does not impose the uniform assumption. We propose instead an approach that defines α_j based on the empirical distribution of cars relative to the counterfactual distribution. We define the ratio between the counterfactual and observed distributions by $\theta_j = \hat{c}_j/c_j$ for $j \in (k-1, k)$ and $= 0$ for $j \notin (k-1, k)$. Then, we define $\alpha_{kj} = \frac{\theta_j}{\sum_{j \in (k-1, k)} \theta_j}$. In this approach, α_{kj} is obtained from the relative ratio between the counterfactual and observed distributions. We use this approach for our main estimate and also report estimates from the uniform assumption approach as well.

In addition to \hat{B}_k (the excess number of cars at notch k), we provide two more estimates that are relevant to our welfare calculations. The first is the excess bunching as a proportion, which is defined as $\hat{b} = c_k/\hat{c}_k$. This is the number of vehicles at a weight notch divided by the counterfactual estimate for that weight. The second estimate is the average changes in weight for cars at notch k , which is the quantity-weighted average of the estimated change in weight:

$$E[\Delta w_k] = \frac{\sum_{j \in (k-1, k)} (w_k - w_j) \cdot (\hat{c}_j - c_j)}{\sum_{j \in (k-1, k)} (\hat{c}_j - c_j)}.$$

We calculate standard errors using a parametric bootstrap procedure, which follows Chetty, Friedman, Olsen, and Pistaferri (2011) and Kleven and Waseem (2013). We draw from the estimated vector of errors ϵ_j in equation (6) with replacement to generate a new set of vehicle counts and follow the steps outlined above to calculate our estimates. We repeat this procedure and define our standard errors as the standard deviation of the distribution of these estimates. Because we observe the exact population distribution of cars in the Japanese automobile market, this standard error reflects error due to misspecification of the polynomial for the counterfactual distribution rather than sampling error.

Figure 4 depicts our procedure graphically for two notch points. In Panel A, we plot the actual distribution and estimated counterfactual distribution at 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 the excess number of cars B is 285.27 (3.75). That is, there are 285 excess cars at this notch compared to the counterfactual distribution. Bunching as a proportion b is 3.75 (0.21), which means that the observed distribution has 3.75 times more observations than the counterfactual distribution at this notch. Finally, the average weight increase $E[\Delta w]$ is 114.97 (0.22) kg for affected cars. Similarly, we illustrate our

Table 2: Excess Bunching and Weight Increases at Each Notch: Old Fuel-Economy Standard


Notch Point	Fuel Economy Standard	Main Estimates			Uniform Assumption		
	below & above the Notch (km/liter)	Excess Bunching (# of cars)	Excess Bunching (ratio)	$E[\Delta\text{weight}]$ (kg)	Excess Bunching (# of cars)	Excess Bunching (ratio)	$E[\Delta\text{weight}]$ (kg)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
830 kg	18.8	16.46	2.13	51.57	16.73	2.17	55.00
	17.9	(7.91)	(0.49)	(3.21)	(7.28)	(0.44)	N.A.
1020 kg	17.9	87.18	2.41	103.77	87.02	2.40	95.00
	16	(8.05)	(0.16)	(0.49)	(7.48)	(0.13)	N.A.
1270 kg	16	163.48	2.47	146.89	163.33	2.46	125.00
	13	(7.92)	(0.11)	(0.62)	(7.33)	(0.08)	N.A.
1520 kg	13	285.27	3.75	114.97	285.41	3.76	125.00
	10.5	(8.21)	(0.21)	(0.22)	(7.52)	(0.15)	N.A.
1770 kg	10.5	143.93	3.52	129.44	144.25	3.54	125.00
	8.9	(8.93)	(0.30)	(0.57)	(8.13)	(0.24)	N.A.
2020 kg	8.9	127.07	8.51	120.77	127.24	8.59	125.00
	7.8	(9.04)	(1.55)	(0.15)	(8.28)	(1.43)	N.A.
2270 kg	7.8	15.67	2.52	137.86	15.52	2.48	125.00
	6.4	(6.40)	(0.66)	(4.48)	(5.95)	(0.64)	N.A.
<u>Kei-Cars</u>							
710 kg	21.2	60.53	2.33	72.57	59.44	2.28	75.00
	18.8	(15.54)	(0.30)	(0.57)	(13.37)	(0.25)	N.A.
830 kg	18.8	118.36	2.06	39.79	120.15	2.09	60.00
	17.9	(15.99)	(0.14)	(0.09)	(12.77)	(0.11)	N.A.
1020 kg	17.9	21.15	2.33	92.63	19.30	2.09	95.00
	16	(9.48)	(2.30)	(1.08)	(7.51)	(2.08)	N.A.

Note: This table shows the regression result in equation 6. Bootstrapped standard errors are in the parentheses.

estimation result at the 2020 kg notch point. At this notch, $B = 127.07$ (9.04), $b = 8.51$ (1.55), and $E[\Delta w] = 120.77$ (0.15).

Table 2 presents our estimates for all notches for the data between 2001 and 2008 (the old fuel economy standard). To see the automakers' incentives at each notch, column 2 shows the stringency of the fuel economy standard (km/liter) below and above the notch (higher km/liter

Figure 4: Graphical Illustration of Estimation of Excess Bunching at Each Notch Point



Note: This figure graphically shows the estimation in equation (6). The figure also lists the estimates of B (excess bunching), b (proportional excess bunching), and $E[\Delta w]$ (the average weight increase). See the main text for details on these estimates.

Table 3: Excess Bunching and Weight Increases at Each Notch: New Fuel-Economy Standard

Notch Point	Fuel Economy	Main Estimates			Uniform Assumption		
	Standard below & above the Notch (km/liter)	Excess Bunching (# of cars)	Excess Bunching (ratio)	E[Δ weight] (kg)	Excess Bunching (# of cars)	Excess Bunching (ratio)	E[Δ weight] (kg)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
980 kg	20.8	5.37	1.60	50.13	5.35	1.60	40.00
	20.5	(5.98)	(0.50)	(1.25)	(5.17)	(0.44)	N.A.
1090 kg	20.5	14.34	2.05	51.58	14.33	2.05	55.00
	18.7	(5.47)	(0.34)	(1.31)	(4.73)	(0.31)	N.A.
1200 kg	18.7	26.56	2.44	61.40	26.58	2.44	55.00
	17.2	(4.85)	(0.24)	(0.26)	(4.19)	(0.24)	N.A.
1320 kg	17.2	20.20	1.89	64.38	20.24	1.89	60.00
	15.8	(4.37)	(0.16)	(0.45)	(3.78)	(0.16)	N.A.
1430 kg	15.8	28.54	2.12	52.18	28.57	2.12	55.00
	14.4	(4.26)	(0.15)	(0.28)	(3.68)	(0.16)	N.A.
1540 kg	14.4	44.44	2.67	61.76	44.45	2.67	55.00
	13.2	(4.45)	(0.18)	(0.10)	(3.84)	(0.21)	N.A.
1660 kg	13.2	81.08	4.13	65.17	81.07	4.13	60.00
	12.2	(4.93)	(0.33)	(0.28)	(4.24)	(0.39)	N.A.
1770 kg	12.2	43.24	2.82	53.86	43.19	2.81	55.00
	11.1	(5.44)	(0.25)	(0.58)	(4.68)	(0.26)	N.A.
1880 kg	11.1	36.60	2.79	52.56	36.52	2.78	55.00
	10.2	(5.86)	(0.29)	(0.49)	(5.04)	(0.30)	N.A.
2000 kg	10.2	33.16	3.09	68.93	33.08	3.08	60.00
	9.4	(5.99)	(0.38)	(0.26)	(5.16)	(0.39)	N.A.
2110 kg	9.4	20.62	2.81	58.08	20.56	2.80	55.00
	8.7	(5.66)	(0.44)	(0.60)	(4.89)	(0.42)	N.A.
2280 kg	8.7	6.66	2.25	91.89	6.65	2.24	88.00
	7.4	(4.09)	(1.13)	(2.21)	(3.57)	(0.87)	N.A.
<u>Kei-Cars</u>							
860 kg	21	18.72	1.66	46.93	18.66	1.66	55.00
	20.8	(4.52)	(0.14)	(0.77)	(4.04)	(0.13)	N.A.
980 kg	20.8	44.82	3.47	55.80	45.04	3.51	60.00
	20.5	(2.31)	(0.40)	(0.97)	(1.80)	(0.54)	N.A.

Note: This table shows the regression result in equation 6. Bootstrapped standard errors are in the parentheses.

numbers imply more stringent standards). Columns 3 to 5 report our main estimates based on the approach described above.

First, we find statistically significant excess bunching at all notches except for the 1020 kg notch for kei-cars, where the estimate is noisy because there are few kei-cars in this weight range. Second, we find substantial heterogeneity in the estimates across the notches. The proportional excess bunching b ranges from 2.1 to 8.5. The estimated weight increases $E[\Delta w]$ range between 40 kg to 93 kg for kei-cars and 52 kg to 147 kg for other cars. For most cars, this amounts to around a 10% increase in weight, which is substantial. Third, our two different approaches for approximating the counterfactual distribution (uniform or not) produce broadly similar results. Our estimates for B and b are not sensitive to the uniform assumption because the excess bunching is very large compared to the counterfactual distribution, so that the way that we reach the integration constraint matters little. We find slightly larger differences in $E[\Delta w]$ between our two methods. With the uniform assumption, $E[\Delta w]$ equals half the width of the regulatory weight step immediately below the notch by assumption, and therefore we have no standard errors for them. Our main estimates do not impose this assumption, but nevertheless yield similar results.

Note that the counterfactual distribution we estimate represents the distribution of vehicle weights that would exist if there was a flat (not attribute-based) fuel-economy standard with the same shadow price.³³ To see this, consider a policy with compliance trading and two weight categories, and thus one notch, at weight \tilde{a} . The shadow price term in the consumer's optimization problem is equal to $\lambda \cdot (e_n - \kappa_1)$ for $a < \tilde{a}$ and $\lambda \cdot (e_n - \kappa_2)$ for $a \geq \tilde{a}$, where $\kappa_1 > \kappa_2$. Thus, for any a , the shadow price term is equal to λ times e_n minus some constant. The marginal regulatory incentive affecting the choice of e is thus the shadow price λ , regardless of the weight category. The marginal incentive affecting the choice of a is zero because a small change in a does not affect the shadow price term, unless $a = \tilde{a}$, in which case there is a discrete jump in the regulatory incentive. Thus, the distortions in a under the notched policy comes only from the vehicles that bunch at a weight threshold, and the marginal incentives for e ($=\lambda$) and a ($=0$ away from the notches) match those in a flat standard with the same shadow price λ .³⁴

³³This is not the same counterfactual as one with no policy at all. Firms may respond to a flat policy by downsizing vehicles (lowering weight) as part of a strategy to boost fuel economy.

³⁴The Japanese policy does not have full compliance trading across firms, so the precise statement is that the counterfactual represents the distribution of weight under a flat subsidy in which each firm faces the same shadow price as in the actual policy. If there were optimization frictions, as in Chetty et al. (2011), then we might expect that some of the vehicles located discretely above the thresholds are also bunching. Our raw data suggest that automakers are able to manipulate weight precisely, because we see excess mass right at each threshold, which suggests that this is not an important issue in our context.

Table 3 presents corresponding estimates for all notches for the data between 2009 and 2013 (the new fuel-economy standard period). Note that the new fuel-economy standard has more, and narrower, notches. This mechanically lowers our estimates for $E[\Delta w]$, because our conservative approach assumes that automakers do not increase weight to move more than one step.³⁵ Results from the second policy follow the same pattern. First, we find statistical significant excess bunching at all notches except for the 980 kg notch and 2280 kg notch for normal cars. At these two notches, the estimates indicate that there is positive excess bunching, but the estimates are noisy because there are not a large number of cars in this weight range. Second, similar to the estimates in the old standard, we find substantial heterogeneity in the estimates between the notches. The proportional excess bunching b ranges between 2.1 and 4.1 for normal cars and 1.7 and 3.5 for kei-cars, depending on the notches. The estimated weight increases $E[\Delta w]$ range between 50 kg to 92 kg for normal cars and 47 kg to 56 kg for kei-cars. Finally, similar to our results for the old standard, the method with the uniform assumption provides similar estimates to our main estimates.³⁶

Overall, the results in this section provide evidence that automakers respond to the attribute-based fuel-economy regulation by changing the weight of vehicles. We find that about 10 percent of vehicles in the Japanese car market manipulated their weight to bunch at regulatory notch points. For these vehicles, the average weight increase induced by the regulation is 110 kg, which is about a 10 percent increase in vehicle weight. This weight increase has welfare implications, as described in our theory above, but it also has implications for safety-related externalities. Heavier vehicles are more dangerous to non-occupants, and when this unpriced, the weight distortions we document here exacerbate safety externalities. We briefly discuss this issue in the next subsection before moving on to our panel analysis.

³⁵Our panel data do suggest that some weight changes are large enough to cover two steps, but we have no grounds for asserting what fraction of vehicles have been thus altered in our cross-sectional analysis, leading us to prefer providing a reliable lower bound on weight changes.

³⁶As mentioned above, the new fuel-economy standards were introduced with a separate subsidy incentive that applied to each specific car model. Therefore, the bunching in the new fuel-economy schedule may come from the incentives created by either policy. The bunching in the old fuel-economy schedule comes only from the incentives created by the fuel-economy standards because there was no separate subsidy incentive. We analyze the new policy's subsidy incentive in section 4.

3.4 Safety-related welfare implications of weight manipulation

Our theory emphasized that attribute-biasing creates a welfare loss because it causes a distortion in the choice of the attribute. The welfare implications of attribute manipulation are altered when the attribute is related to another unpriced externality because the existence of a second externality implies that the privately optimal choice of the attribute is not socially optimal.³⁷

In the event of a traffic accident, heavier automobiles are safer for the occupants of the vehicle (this is a private benefit) but more dangerous for pedestrians or the occupants of other vehicles (this is an externality).³⁸ Thus, the optimal attribute-based policy should *tax* vehicle size rather than subsidize it. The implicit subsidy on vehicle weight in the Japanese fuel-economy standards therefore exacerbate accident-related externalities.

We obtain a back-of-the-envelope estimate of the magnitude of this distortion by multiplying our estimate of the average change in car weight by an estimate of the increased probability that a heavier vehicle causes a fatality during its lifetime, times an estimate of the value of a statistical life. Specifically, the weighted average increase across all cars that bunch at the notches in Table 2 is 109.62 kg. Anderson and Auffhammer (2014) estimate 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, compared to a mean probability of 0.19 percent. For the value of a statistical life, we use \$9.3 million, which comes from a study in Japan (Kniesner and Leeth 1991), and is within the range of standard estimates from the United States.

Multiplying through, we calculate the welfare loss, per car sold, for a 110 kg weight increase as: $110 * 0.0009 * (2.2/1,000) * \$9.3 \text{ million} = \$2026$ per car that changes weight in response to the policy. Our analysis suggests that the excess bunching accounts for about 10 percent of cars in the market. The Japanese car market sells around 5 million new cars per year, so we estimate our aggregate annual welfare distortion to be 10 percent of 5,000,000 times \$2026, which is \$1.0 billion. For context, automaker revenue in Japan is roughly \$150 billion per year, and Toyota's

³⁷It is straightforward to incorporate a second externality into our framework. In the simplest case when a causes a separate externality, the optimal attribute-slope will be designed to create a Pigouvian tax on a . Attribute-biasing simply provides a second policy instrument, which is necessary for dealing with a second market failure.

³⁸In principle, the externality risk may be partly priced through insurance or legal liability. White (2004), however, argues that neither tort liability nor mandatory liability insurance prices safety externalities. In brief, tort liability requires negligence, not just that one be driving a dangerous vehicle. Liability insurance generally covers the cost of damages to a vehicle, but it is but a small fraction of the value of a life. In addition, rate differences across vehicles are very coarse and reflect average driver characteristics in concert with vehicle attributes.

global revenue in 2013 was \$210 billion. Equivalently, our calculation implies that each new cohort of 5 million cars sold in Japan will be associated with an extra 103 deaths over the lifetime of those cars, which compares to an annual fatality rate of roughly 6,000. Our calculations are meant only as back-of-the-envelope estimates, but they make clear that the welfare distortions induced by the Japanese policy are economically significant.


4 A Double Notched Policy and Panel Analysis

The previous section provides evidence of attribute distortions by focusing on cross-sectional variation in regulatory incentives. In this section, we exploit panel variation created by a policy change to examine the welfare implications of attribute-biasing, including both its potential benefits and its limitations. In 2009, the Japanese government introduced new corporate average fuel-economy standards (with target year 2015) as well as a separate incentive that applies to each specific car model, rather than to an entire company’s fleet. If any individual model met the new fuel-economy standard in years before the target year, consumers purchasing that car received a direct subsidy of \$1,000 (or, \$700 for kei-cars). In addition, cars with fuel economy 10% and 20% higher than the standard received a more generous incentive, in the form of reduced registration fees that were normally levied on the consumer. This vehicle-specific subsidy created a “double notched” policy—a car had to be above a step-function in the two-dimensional space of fuel economy and weight to qualify for the subsidy. Vehicles that were redesigned to become eligible for the subsidy reveal information about the relative costs of changing weight versus fuel economy. We use this revealed preference information to estimate an adjustment cost function, which we explain further below, that enables us to simulate the welfare effects of several policy alternatives.

Our procedures require panel data, which we create by linking vehicles in 2008 (before the policy change) and 2012 (the last year of our data) using a unique product identifier (ID) that is included in the regulatory data.³⁹ We first match on product ID across years, which is often, but not always, constant over time. If automakers change the product ID between years, we match by using model name, displacement, drive type (e.g., four-wheel drive), and transmission (manual or automatic).

³⁹Product ID is narrower 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, each of which will have a unique ID.

Figure 5: Fuel Economy and Weight before and after the Policy Change



Note: This figure shows each vehicle’s fuel economy and weight before and after the introduction of the new subsidy that was applied to each vehicle individually. The scatterplot shows each car’s starting values of fuel economy and weight in 2008—the year before the policy change. For the cars that qualified for the new subsidy in 2012, we also show “arrows” connecting each car’s starting values in 2008 with its values in 2012. The figure also includes three step functions that correspond to the three tiers of the new incentive’s eligibility cutoffs.

That is, we consider two cars sold in two different years to be the same if they have matching IDs, or if they have exactly the same model name, displacement, drive type, and transmission. Entry and exit imply that not all data points are matched, and we end with 675 matched records.⁴⁰

We begin by plotting our raw data in Figure 5. The scatterplot shows each car’s starting values of fuel economy and weight in 2008—the year before the policy change. For the cars that qualified for the new subsidy in 2012, we also show “arrows” connecting each car’s starting position in 2008 with its position in 2012. The figure includes three step functions that correspond to the three

⁴⁰Our matching procedure guarantees that we match the same model name, which avoids mismatching the panel structure of the data. We take this approach because it provides transparent matching criteria. A potential drawback of this approach is that firms may change some of their model names over time, yet they are targeting similar customer segments. To address this concern, we also conduct our analysis by including unmatched cars from the first matching criteria whenever we can match them using displacement, drive type, and transmission, while ignoring model names. This procedure produces a slightly different set of matched data, but our final estimation results are very similar regardless of which matching procedure is used.

tiers of the new incentives. To qualify for the incentive, a vehicle had to be anywhere on or above these lines in the two-dimensional space. The figure indicates that 1) vehicles that gained the subsidy were usually modified so that they were just above the eligibility cutoff, 2) vehicles which started off closer to the new standard were more likely to get the incentive—the “distance” to the eligibility lines explain most of the variation in which cars get the subsidy, 3) most cars (but not kei-cars) increased both weight and fuel economy in order to become eligible—that is, they moved “northeast”.⁴¹

Kei-cars differ in a telling way. The attribute schedule is nearly flat in the weight distribution inhabited by kei-cars, so they are subject to a nearly-flat (not attribute-based) policy.⁴² Most kei-cars that obtained the subsidy increased fuel economy, while decreasing weight slightly. This accords with our theory: the steeper sloped portion of the standard induced substantial weight manipulation, while the flatter sloped part of the standard did not.⁴³

In the appendix, we also show the corresponding arrows for vehicles that did not receive a subsidy. Most of these vehicles had much smaller changes in fuel economy and weight, which suggests that secular trends were modest during this period. We also find that some of these cars increased weight to be at the weight notch underneath the subsidy cutoff lines. This is because, in addition to the model-specific subsidy incentive, vehicles were still influenced by the corporate average fuel-economy standards, which gave them an incentive to increase weight.

4.1 Discrete Choice Model of Vehicle Redesign

The descriptive panel data provide further evidence in support of our theoretical prediction that attribute-based policies cause distortions in the secondary attribute. In this section, we use this revealed preference information about the choice of fuel economy and weight to estimate a model of an adjustment cost function of these two attributes, which we then use for policy simulation.

To see the logic of our approach, suppose that we observed data from a competitive market in two time periods, the first of which has no policy, and the second of which features the model-

⁴¹To highlight this, we color vectors that show an increase in weight in red and a decrease in weight in blue.

⁴²The figure shows an approximate dividing line between kei-cars and other cars. Kei-cars are not regulated by weight, but rather by engine displacement, so this division is not strict. The flatness of the schedule is due to the front runner system and the fact that there was little variation in the top fuel efficiencies over this range of weight.

⁴³Other factors may cause the kei-cars to be different. For example, they may be targeted at consumers who value fuel economy more than average, which might make it more costly for kei-cars to increase weight than other vehicles.

specific subsidy.⁴⁴ For each vehicle, the pair of fuel economy and weight that was chosen in the first period would be welfare maximizing, given the firm's production function and the tastes of its consumers. Any deviation from that characteristic bundle would lower private welfare. Graphically, in a two-dimensional diagram like Figure 5, there would be level sets of the adjustment cost function that surround the original choice of fuel economy and weight. The adjustment cost is zero at the original bundle. It becomes larger as the new bundle gets further from the original, optimal one.

When the subsidy is introduced in the second period, there are two possibilities for a vehicle's new optimal bundle. The first case is that the vehicle changes the bundle of a and e to receive the subsidy. This happens if the minimum adjustment cost to get the subsidy is lower than the subsidy's value. The second case is that the vehicle does not change its bundle, which happens if the minimum adjustment cost to get the subsidy is higher than the subsidy's value.

Figure 5 shows which vehicles changed in order to get the subsidy, and the vectors show the path taken by those products that did change in order to gain the subsidy. If we make an assumption about the functional form of the adjustment cost function, we can use these data to estimate that function. We begin with the assumption that the loss function, denoted L is quadratic: $L = \alpha(a_n - a_n^o)^2 + \beta(e_n - e_n^o)^2 + \gamma(a_n - a_n^o)(e_n - e_n^o)$, where a_n and e_n are the second-period characteristics, a_n^o and e_n^o are the first-period characteristics (which maximize private welfare), and α , β and γ are parameters to be estimated. This functional form implies that the level sets of the adjustment cost function are ellipses centered around the private optimum (a_n^o, e_n^o) .⁴⁵ One such ellipse will be tangent to the standard at the new optimum.

We describe the choice of the new optimum (a_n, e_n) for product n as the outcome of a discrete choice over all of the possible (discretized) grid points in a by e space.⁴⁶ We denote each grid point

⁴⁴In reality, firms also had the corporate average fuel-economy standards. We begin with a simple model that ignores this regulatory incentive and incorporate it in our second model.

⁴⁵Mathematically, an ellipse can be defined in an x-y plane by $b_1(x - c_1)^2 + b_2(y - c_2)^2 + b_3(x - c_1)(y - c_2) = A$, where (c_1, c_2) is the centroid, b_1 and b_2 determine the width and height of the ellipse, b_3 determines its tilting angle, and A is the level. In our adjustment cost function, the level set has $(c_1, c_2) = (a_n^o, e_n^o)$. Note that the loss function includes the welfare losses coming from both the demand side and supply side. The deviation from the initial optimal a and e could induce additional costs for the supply side (e.g., higher production cost) and the demand side (e.g., lower consumer utility from the durable). Our model estimates the total welfare loss function; it does not separately identify demand and supply parameters.

⁴⁶While weight and fuel economy are in principle continuous measures, the regulatory data are measured in discrete units (10 kilograms for weight and tenths of a kilometer-per-liter for fuel economy), and all regulations are based on these discrete units. In our data, the maximum and minimum of fuel economy (km/liter) are 7.6 and 30.0, and those of weight (kg) are 730 and 2190. To make a choice set of the discrete choice problem, we use these maximum and minimum values to create a rectangle of grid points that include all possible bundles for our data.

as a unique value of z . The second-period optimization problem for product n is then to choose the a_n and e_n values that maximize the loss function plus the subsidy:

$$W_{nz} = \alpha(a_n - a_n^o)^2 + \beta(e_n - e_n^o)^2 + \gamma(a_n - a_n^o)(e_n - e_n^o) + \tau \cdot 1(e_n \geq \sigma(a_n)) + \varepsilon_{nz}, \quad (7)$$

where τ is also a parameter to be estimated, which represents the value of receiving the subsidy, and ε_{nz} is an error term specific to each vehicle n and each grid point z .⁴⁷ We assume that ε_{nz} is a Type-I extreme value error term and estimate equation 7 via a logit. Logit coefficients are scaled by the variance of the error term. When we interpret the parameters in terms of dollars, we rescale them by the dollar value of the subsidy by dividing by $\hat{\tau}$ (Train 2009).

Our estimation benefits from quasi-experimental variation in choice sets caused by the second-period policy. Even though all products face the same second-period policy, they each have a different set of *changes* in a and e that would gain them the subsidy. This variation comes from differences in starting points, from the introduction of new weight notches, and from the fact that the changes in the standards are different across the weight categories. As a result, some vehicles are able to make modest improvements in fuel economy to gain the subsidy, whereas others require large changes. And, some vehicles can take advantage of a weight notch with small increases in weight, but others require a large increase. These differences create a rich source of identification for our estimation. The intuition is that the relationship between each product’s “distance” to the subsidy cutoff lines and its probability of obtaining the subsidy identifies the subsidy coefficient (τ), and the “route” taken by products facing different compliance options identifies the coefficients that determine the shape of the cost function (α , β and γ).

Table 4 presents estimates from the logit for this specification, in column 1. We define Δ Weight as each vehicle’s change in weight in 100 kilogram units and Δ (Fuel consumption) as its change in liters per 100km (l/100km).⁴⁸ The coefficient on the interaction term is small and statistically insignificant, which implies that the elliptical level sets of the adjustment cost function have limited “tilt”. The first two coefficients (-1.25 and -1.15) are roughly the same, which implies that a

⁴⁷This formulation omits the payoff in the first-period, which is a constant and therefore does not influence choice.

⁴⁸We use fuel consumption (l/100km) rather than fuel economy (km/liter) here because when regulators calculate the corporate average fuel economy, they average each model’s fuel consumption rather than fuel economy (equivalently, they harmonically average fuel economy). Thus, the control variables added in later specifications need to measure fuel consumption to be correctly specified. All of our conclusions here are robust to estimating the models with fuel economy instead. (Note that the distinction is irrelevant in logs.).

change in weight by 100 kg and a change in fuel consumption by one l/100km result in approximately the same loss of private welfare.⁴⁹ In Column 2, we also estimate the incremental effects of higher subsidies by including the interactions between the subsidy dummy variable and each of the higher subsidy dummy variable. The interaction with the highest subsidy dummy is estimated with large standard errors because not many vehicles passed the highest subsidy's thresholds, which is evident in Figure 5.

Proposition 3 showed that, when the attribute slope is fixed, the relative responsiveness of a and e to a change in the subsidy would dictate the degree to which the second-best subsidy rate deviates from marginal benefits, which is the Pigouvian benchmark. The adjustment cost function estimated here allows us to calculate the relevant terms. The adjustment cost function parameters α , β and γ can be translated into the derivatives $\partial a/\partial s$ and $\partial e/\partial s$. Measuring e in terms of fuel consumption, the average attribute-slope of the policy is .44.⁵⁰ If the policy were smooth with a slope of .44, the second-best tax rate would be marginal damages (or benefits) times .83. That is, the optimal corrective policy would be attenuated by 17%.

This simple specification provides transparent interpretations on the estimates of the adjustment cost function. However, it considers only the product-specific subsidy and ignores the fleet-average regulation. This is potentially important because the model-specific subsidy policy was introduced at the same time that the new fleet-average regulation was announced. The effect of the new fleet-average regulation might be minimal because it is legally binding only in the target year, which is 2015. However, the new fleet-average regulation might affect firm decisions before 2015, as automobiles are redesigned on a multi-year cycle and are likely redesigned before the target year. Moreover, the bundles of a and e before the policy change may not be privately optimal if they were distorted by the old fleet-average regulation. Even in the absence of the product-specific subsidy, therefore, the optimal choice of a and e could shift between the two periods; i.e., the level sets of the loss function could be centered around some alternative point, rather than around the original

⁴⁹For a car at the average fuel economy in our sample, this implies that a change in weight of 200 lb and a change in fuel consumption of 4.5 miles per gallon result in approximately the same loss. In dollars, this means that, for example, a change in weight by 100 kg (220 lb) induces a loss of $\$1621 = 1.24 \cdot \$1000/0.77$ and a change in fuel consumption by l/100km (equivalent to 4.5 miles per gallon at the fleet average fuel economy) costs $\$1494 = 1.15 \cdot \$1000/0.77$. This is a reasonable order of magnitude. If gasoline costs \$2.50 per gallon, a vehicle of average fuel economy driven 12,000 miles per year for 13 years at a 5% discount rate would save around \$2,100 in fuel costs from a 4.5 mpg improvement.

⁵⁰We estimate this as the best linear fit of the end point of the vehicle weight bins and fuel consumption standards.

Table 4: Estimates of the Adjustment Cost Function

	(1)	(2)	(3)	(4)
$(\Delta\text{Weight})^2$	-1.24 (0.08)	-1.25 (0.08)	-1.25 (0.08)	-1.25 (0.08)
$(\Delta\text{Fuel consumption})^2$	-1.15 (0.07)	-1.19 (0.08)	-1.21 (0.08)	-1.21 (0.09)
$\Delta\text{Weight} \times \Delta\text{Fuel consumption}$	0.13 (0.11)	0.15 (0.11)	0.16 (0.11)	0.16 (0.11)
$1\{\text{Subsidy}\}$	0.77 (0.15)	0.69 (0.16)	0.37 (0.18)	0.32 (0.19)
$1\{\text{Subsidy}\} \times 1\{\text{Higher Subsidy}\}$		0.39 (0.21)		0.26 (0.21)
$1\{\text{Subsidy}\} \times 1\{\text{Highest Subsidy}\}$		0.19 (0.33)		-0.08 (0.34)
Shadow price for new policy (λ)			0.45 (0.14)	0.44 (0.15)
– Shadow price for old policy ($-\dot{\lambda}$)			-0.11 (0.12)	-0.11 (0.12)

Note: This table shows the estimation results of the discrete choice models in equation (7) in columns 1 and 2 and equation (8) in columns 3 and 4. The dependent variable is a set of the possible bundles of fuel consumption and weight. The variable equals 1 if the bundle was chosen by product n and 0 otherwise. The weight and fuel consumption variables are rescaled in the estimation—we define ΔWeight as each vehicle’s change in weight in 100 kilogram units and $\Delta(\text{Fuel consumption})$ as its change in liters per 100km (l/100km). The estimation is based on the panel data of 675 matched vehicle models.

choice of (a_n^o, e_n^o) .

We can account for the effects of the regulations by augmenting our estimating equation. The derivation is straightforward, but it is algebraically involved, so we relegate it to the appendix. There, we show that the following estimating equation provides estimates of the parameters necessary to identify the adjustment cost function, taking into account the effects of the regulations:

$$V_{nz} = \alpha(a_n - a_n^o)^2 + \beta(e_n - e_n^o)^2 + \gamma(a_n - a_n^o)(e_n - e_n^o) + \lambda(e_n - \sigma(a_n)) + \tau 1(e_n \geq \sigma(a_n)) - \lambda^o(e_n - \sigma^o(a_n)) + \varepsilon_{nz}, \quad (8)$$

where λ^o is the shadow price of the original regulation and λ is the shadow price of the new regulation. By including $(e_n - \sigma(a_n))$ (the position of the vehicle vis-à-vis the new policy) and $(e_n - \sigma^o(a_n))$ (the position of the vehicle vis-à-vis the old policy evaluated at the *new* choice of a and e) in our estimating equation, we obtain estimates of these shadow prices from our coefficients (along with α , β , γ and τ). Note that equation (8) is identical to equation (7) except that equation (8) includes two additional variables that estimate the incentives created by the fleet-average regulation. Furthermore, in the appendix, we show that the adjustment cost function excluding the effect of the fleet-level regulation effects will be $\alpha(a_n - a_n^o)^2 + \beta(e_n - e_n^o)^2 + \gamma(a_n - a_n^o)(e_n - e_n^o) + \lambda^o[e_n^o - \sigma^o(a_n^o) - (e_n - \sigma^o(a_n))]$. We can estimate that function from our coefficients after estimating equation (8).

Columns 3 and 4 of Table 4 present the estimation results with the regulatory control variables. The addition of the regulatory control variables have very little impact on the coefficients on the adjustment function. The subsidy coefficient, however, becomes smaller in this specification. This is because we are likely to overestimate the subsidy effect in columns 1 and 2 since we do not control for the corporate compliance regulation. The changes in e and a are driven not only from the car-specific incentive but also from the firm-level regulatory incentive. Columns 3 and 4 explicitly control for both of those incentives. Importantly, while the coefficient on the subsidy affects the dollar values of the adjustment cost, it does not have major impacts on the *relative* costs of changing e versus a because the estimates on $(\Delta a)^2$ and $(\Delta e)^2$ are robust between the two approaches. In the analysis of welfare implicates, we report our results based on both of these approaches and

also provide the implications in terms of relative costs.⁵¹

Our estimation in this section is based on strong assumptions. It assumes that the functions determining taste and price are unchanged between the two periods. It is based on only our matched observations, and does not model entry and exit. It also implicitly assumes that markups on each vehicle do not change between the two periods. Relaxing this assumption would be an interesting and important direction for future research, but it is well beyond the scope of this paper.⁵² Because our estimation is based on these assumptions, we do not intend to claim that our procedure provides a complete welfare analysis of the Japanese fuel economy regulation. Our procedure does, however, (a) establish a method for analyzing double notched policies, which have not been studied in the economics literature, and (b) enable us to use real data on vehicle redesign to illustrate the basic economic implications of attribute-basing, which is what we pursue in the next section via policy simulation.

4.2 Policy Simulations and Welfare Implications

We use our estimates of the adjustment cost function to simulate three policies: attribute-based fuel-economy standards, a flat fuel-economy standard with no compliance trading, and a fully efficient policy that is equivalent to a flat standard with compliance trading. Our aim is to highlight the basic economics of attribute-basing with an example based on real data and policy.

The first policy schedule we consider is the actual new attribute-based fuel economy standard schedule in Japan—the bottom schedule in Figure 5. The benefits of attribute-basing are limited as there is more averaging (or trading) across products. Thus, to best highlight the benefits of attribute-basing, we suppose that the schedule represents a mandatory minimum *for each vehicle*,

⁵¹Though it is not our focus, we do note that this procedure also delivers an estimate of the shadow price of fuel-economy regulations, which are of much interest to the literature. Our procedure differs from the existing literature in leveraging panel data around a policy *change* to identify the shadow price, whereas existing work either examines a specific policy loophole (Anderson and Sallee 2011) or uses static structural models (Goldberg 1998; Gramlich 2009; Jacobsen 2013a). The estimate in column 3 implies a shadow price of \$1,162 ($= .45/.37 * \1000) per unit of 1/100km car per. This translates into \$258 per mpg per car at the average fuel economy in our sample. Our simple approach does not account for imperfect competition, so we do not stress these results, but simply note that they are the same order of magnitude as results found in Gramlich (2009) and Jacobsen (2013a) for the United States.

⁵²For example, Berry, Levinsohn, and Pakes (1995) and the subsequent literature provide methods to estimate markups in differentiated product markets in a static setting. The estimation requires that all non-price attributes are fixed and not endogenously determined by firms. That is, the method can be applied to a static setting in which vehicle attributes are fixed and exogenous to firm decisions, but it cannot be applied to our context, which is concerned explicitly with how multiple product characteristics are altered in a dynamic setting. Relaxing this assumption and incorporating the endogenous choices of product characteristics into complete structural models is one of the key ongoing research topics in the literature (see, for example, Fan (2013)).

rather than for the fleet average. Appliance standards in many countries, including the United States, are structured this way, and the Chinese vehicle standards initially imposed per-product minima. Graphically, this policy requires all vehicles in Figure 5 to move above the lowest line. Given each vehicle’s initial point and our estimates of the loss function L , we find the a and e that achieves compliance at the lowest possible cost. We then calculate the resulting Δa , Δe , and ΔL for each car’s new optimal point.⁵³

Table 5 reports results for this attribute-based regulation in the first row. Panel A is based on the adjustment cost function estimated in the simple quadratic specification (column 2 of Table 4). The ABR lowers fuel consumption e by 0.76 l/100km on average, which is approximately a 10 percent improvement in fuel economy.⁵⁴ While this improvement produces benefits from the externality, it comes at a private cost (a loss in L) of \$1843 per unit sold, averaged across all model types. This cost can be decomposed into a cost from Δe and from Δa , which is shown in the table. Consistent with the vectors in Figure 5, the ABR causes a change in a even though the policy’s target is e . The simulated ABR causes a 33 kilogram (73.37 lb) average weight increase, which is about a 3 percent increase in weight. The cost associated with this weight increase accounts for about 30% of the total cost of the regulation. This cost represents the Harberger triangle associated with Proposition 1 in our theory. The table also reports the standard deviation in the marginal cost of increasing e at the optimal choice across models, which is calculable from our loss function. This is an important statistic for measuring the efficiency of the regulation because the benefit of the ABR as compared to a flat standard is greater equalization of marginal costs of compliance.

Our second simulated policy is a flat fuel economy standard with no compliance trading. That is, all cars have a common fuel economy standard regardless of their weight, and each of them has to comply with the standard. To compare this policy with our first policy (the ABR), we find a flat policy that generates exactly the same average improvement in e as the one in the first policy. We find the flat policy that satisfies this condition numerically. The resulting standard is 8.04 l/100km (12.4 km/liter and 29.2 mpg) for Panel A.⁵⁵

⁵³In our simulation, we assume that all vehicles stay in the market after the introduction of this simulated policy, although in reality there can be entry and exit.

⁵⁴In our data, the average fuel consumption in 2008 is approximately 7.6 l/100km, which is 13 km/liter and 30.5 miles per gallon. For a car with the average fuel consumption in 2008, an improvement of fuel consumption by 0.76 l/100km implies an improvement of fuel economy by 1.43 km/liter or by 3.35 mpg.

⁵⁵For Panel B, the resulting flat standard is 8.09 l/100km (12.3 km/liter and 29.0 mpg). Again, we assume that all products must comply and that there is no exit. In reality, such mandates would likely cause some products, which

Table 5: Welfare Implications of Attribute-Based Regulation and Alternative Policies

	Δe : Fuel consumption (liter/100km)	Δa : Weight (kg)	Cost from Δe (\$/car)	Cost from Δa (\$/car)	Welfare cost (\$/car)	Cost relative to ABR	S.D. of MC (\$/car)
Panel A) Based on the Loss Function without Controls for Compliance Regulation							
ABR	-0.76	33.28	-1319	-524	-1843	1.00	1805
Flat	-0.76	0.00	-3590	0	-3590	1.95	4044
Efficient	-0.76	0.00	-731	0	-731	0.40	0
Panel B) Based on the Loss Function with Controls for Compliance Regulation							
ABR	-0.74	35.43	-2584	-1391	-3975	1.00	3955
Flat	-0.74	0.00	-7272	0	-7272	1.83	8614
Efficient	-0.74	0.00	-1309	0	-1309	0.33	0

Note: This table shows the results of our three policy simulations: 1) attribute-based fuel-economy standards (ABR), a flat fuel-economy standard with no compliance trading (Flat), and a fully efficient policy that is equivalent to a flat standard with compliance trading (Efficient).


The flat standard has the benefit of not distorting weight.⁵⁶ The flat standard, however, creates many infra-marginal vehicles; that is, vehicles that are in compliance with the standard without any change in fuel economy. These infra-marginal vehicles have a zero marginal cost of increasing e , but do not change e at all, because there is no regulatory incentive. Other vehicles have very large marginal costs of increasing e because they have to improve fuel economy by a large amount in order to comply with the policy. This dispersion in marginal costs is inefficient, and it results in welfare costs that are, on average, 1.95 times larger than the welfare costs of the ABR. To illustrate this, Figure 6 plots the distribution of marginal costs of compliance under the ABR and the flat policy. Under the flat policy, there are many more infra-marginal (zero marginal cost) observations, and the remaining distribution is also more diffuse. The benefit of attribute-basing is the (partial) harmonization of these marginal costs.

This harmonization, however, is incomplete and that makes attribute-basing an inefficient substitute for a compliance trading system. Our third policy, which is a flat policy with compliance

are particularly far away from the standard, to exit the market, in both the flat and that attribute-based policies.

⁵⁶The fact that weight does not change by a detectable amount is due to the fact that we estimate a very small magnitude on the interaction term in our cost function. If that estimate were strongly positive or negative, the flat standard could induce a decrease or increase in weight, but this would not represent a distortion.

Figure 6: Distributions of the Marginal Costs of Compliance



Note: This figure shows the distributions of the marginal cost of compliance for the three simulated policies: 1) the attribute-based regulation (left), 2) the counterfactual flat policy (right), and 3) the counterfactual efficient policy (dashed line). The distribution of the efficient policy is the vertical line because the standard deviation is zero.

trading that generates the same Δe as the ones in the first and second policy, demonstrates this. As shown in the theory section, this policy is equivalent to the efficient Pigouvian subsidy that generates the same average Δe . Under this policy, all cars improve fuel economy by the same amount (because we assume a common cost function). This policy completely harmonizes marginal costs, no products are infra-marginal, and therefore the total welfare loss is minimized. Because the standard deviation of the marginal compliance cost is zero, its distribution collapses to a constant, which we label in Figure 6. Finally, we calculate the efficient policy’s welfare loss relative to the first policy (ABR). The last column in Panel A shows that the efficient policy would lower total costs by 60 percent, while achieving the same improvement in fuel economy.

Panel B calculates the same statistics using the loss function that includes control variables for corporate compliance regulation (column 4 of Table 4). In relative terms, all results in Panel A and B come to essentially the same conclusion. Compared to the ABR, the efficient policy would lower

total welfare loss by 67 percent, and the flat policy would increase total welfare loss by 83 percent, for the same improvement in e . The difference between Panels A and B is their dollar values. The estimates from the discrete choice model in Table 4 imply that we are likely to overestimate the effect of the subsidy if we do not control for the corporate average regulation.⁵⁷ Overestimating the subsidy effect underestimates the dollar values of the adjustment cost function. This is why we find lower dollar values for Panel B. However, importantly, this scaling issue has only a slight effect on our illustration of the *differences* between the three policies because the scaling affects all policies in the same way.

In sum, our simulation highlights three key policy implications consistent with our theoretical predictions. First, ABR creates substantial distortions in the attribute, which accounts for 30% of the total regulatory cost in our simulation. Second, a benefit of ABR is the *partial* harmonization of the marginal costs of compliance. Relative to a counterfactual non-ABR flat policy, ABR produces smaller dispersions of the marginal compliance costs between products. This smaller dispersion produces a smaller total regulatory cost compared to a flat policy. Third, ABR is nevertheless an inefficient substitute for a fully efficient policy such as a non-ABR policy with compliance trading because the efficient policy does not create attribute distortions and fully equalizes the marginal costs of compliance.

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 improve efficiency by equalizing marginal costs of regulatory compliance, but we emphasize that, even in those cases, the optimal attribute-basing function deviates fundamentally from those observed in real world policies because it must balance efficiency gains from cost equalization with the distortion it induces in the attribute.

Empirically, the paper demonstrates that distortions in response to attribute-based fuel-economy

⁵⁷This is because the changes in e and a are driven not only from the car-specific subsidy incentive but also from the firm-level regulatory incentive. Our second specification explicitly controls for both of the incentives.

standards in Japan are clearly present. We use both established cross-sectional tools based on the notch literature, as well as novel panel techniques that take advantage of a “double notched” policy, to demonstrate that the Japanese car market has experienced a notable increase in weight in response to attribute-based regulation.

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. 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

- Allcott, Hunt, Sendhil Mullainathan, and Dmitry Taubinsky. 2014. “Energy Policy with Externalities and Internalities.” *Journal of Public Economics* 112:72–88.
- Allcott, Hunt and Nathan Wozny. 2012. “Gasoline Prices, Fuel Economy, and the Energy Paradox.” *Review of Economics and Statistics* .
- Anderson, Michael and Maximilian Auffhammer. 2014. “Pounds that Kill: The External Costs of Vehicle Weight.” *Review of Economic Studies* 82 (2):535–571.
- Anderson, Soren T., Ian W.H. Parry, James M. Sallee, and Carolyn Fischer. 2011. “Automobile Fuel Economy Standards: Impacts, Efficiency and Alternatives.” *Review of Environmental Economics and Policy* 5 (1):89–108.
- Anderson, Soren T. and James M. Sallee. 2011. “Using Loopholes to Reveal the Marginal Cost of Regulation: The Case of Fuel-Economy Standards.” *American Economic Review* 101 (4):1375–1409.
- Becker, Randy and Vernon Henderson. 2000. “Effects of Air Quality Regulations on Polluting Industries.” *Journal of Political Economy* 108 (2):379–421.
- Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. “Automobile Prices in Market Equilibrium.” *Econometrica* 64 (4):841–890.
- Blinder, Alan S. and Harvey S. Rosen. 1985. “Notches.” *American Economic Review* 75 (4):736–747.
- Borenstein, Severin. Forthcoming. “A Microeconomic Framework for Evaluating Energy Efficiency Rebound and Some Implications.” *Energy Journal* .
- Brock, William A. and David S. Evans. 1985. “The Economics of Regulatory Tiering.” *The Rand Journal of Economics* 16 (3):398–409.
- Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer. 2013. “Are Consumers Myopic? Evidence from New and Used Car Purchases.” *American Economic Review* 103 (1):220–256.

- 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.
- Dharmapala, Dhammika, Joel Slemrod, and John Douglas Wilson. 2011. "Tax Policy and the Missing Middle: Optimal Tax Remittance with Firm-Level Administrative Costs." *Journal of Public Economics* 95 (9-10):1036–1047.
- Diamond, Peter A. 1973. "Consumption Externalities and Imperfect Corrective Pricing." *The Bell Journal of Economics* 4 (2):526–538.
- Fan, Ying. 2013. "Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market." *American Economic Review* 103 (5):1598–1628.
- Fischer, Carolyn, Winston Harrington, and Ian W.H. Parry. 2007. "Should Automobile Fuel Economy Standards Be Tightened?" *The Energy Journal* 28 (4):1–29.
- Fullerton, Don and Sarah E. West. 2002. "Can Taxes on Cars and Gasoline Mimic an Unavailable Tax on Emissions?" *Journal of Environmental Economics and Management* 42 (1):135–157.
- . 2010. "Tax and Subsidy Combinations for the Control of Car Pollution." *The B.E. Journal of Economic Analysis & Policy (Advances)* 10 (1).
- Gao, Feng, Joanna Shuang Wu, and Jerold Zimmerman. 2009. "Unintended Consequences of Granting Small Firms Exemptions from Securities Regulation: Evidence from the Sarbanes-Oxley Act." *Journal of Accounting Research* 47 (2):459–506.
- Gillingham, Kenneth. 2011. "How Do Consumers Respond to Gasoline Price Shocks? Heterogeneity in Vehicle Choice and Driving Behavior." Manuscript: Yale University.
- . 2013. "The Economics of Fuel Economy Standards versus Feebates." Manuscript: Yale University.
- Goldberg, Pinelopi Koujianou. 1998. "The Effects of the Corporate Average Fuel Efficiency Standards in the U.S." *The Journal of Industrial Economics* 46 (1):1–33.
- Gramlich, Jacob. 2009. "Gas Prices, Fuel Efficiency, and Endogenous Product Choice in the U.S. Automobile Industry." [Http://faculty.msb.edu/jpg72/Autos.GramlichJacob.pdf](http://faculty.msb.edu/jpg72/Autos.GramlichJacob.pdf).
- Heutel, Garth. 2011. "Optimal Policy Instruments for externality-producing Durable Goods Under Time Inconsistency." NBER Working Paper 17083.
- Holland, Stephen P., Jonathan E. Hughes, and Christopher R. Knittel. 2009. "Greenhouse Gas Reductions Under Low Carbon Fuel Standards?" *American Economic Journal: Economic Policy* 1 (1):106–146.
- Jacobsen, Mark R. 2013a. "Evaluating U.S. Fuel Economy Standards In a Model with Producer and Household Heterogeneity." *American Economic Journal: Economic Policy* 5 (2):148–187.
- . 2013b. "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.

- Kaplow, Louis. 2013. "Optimal Regulation with Exemptions and Corrective Taxes." Working Paper: Harvard University.
- Kleit, Andrew N. 2004. "Impacts of Long-Range Increases in the Fuel Economy (CAFE) Standard." *Economic Inquiry* 42 (2):279–294.
- 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.
- Kniesner, Thomas J. and John D. Leeth. 1991. "Compensating wage differentials for fatal injury risk in Australia, Japan, and the United States." *Journal of Risk and Uncertainty* 4 (1):75–90.
- Knittel, Christopher R. and Ryan Sandler. 2012. "Carbon Prices and Automobile Greenhouse Gas Emissions: The Extensive and Intensive Margins." In *The Design and Implementation of U.S. Climate Policy*, edited by Don Fullerton and Catherine Wolfram. NBER, 287–299.
- Kopczuk, Wojciech. 2003. "A Note on Optimal Taxation in the Presence of Externalities." *Economics Letters* 80:81–86.
- Lu, S. 2006. "Vehicle Survivability and Travel Mileage Schedules." Technical Report DOT HS 809 952, National Highway Traffic Safety Administration.
- 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.
- 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.
- Slemrod, Joel. 2010. "Buenas Notches: Lines and Notches in Tax System Design." Manuscript: University of Michigan.
- Sneeringer, Stacy and Nigel Key. 2011. "Effects of Size-Based Environmental Regulations: Evidence of Regulatory Avoidance." *American Journal of Agricultural Economics* 93 (4):1189–1211.
- Spence, Michael A. 1975. "Monopoly, Quality, and Regulation." *The Bell Journal of Economics* 6 (2):417–429.
- Train, Kenneth. 2009. *Discrete choice methods with simulation*. Cambridge university press.
- White, Michelle J. 2004. "The "Arms Race" on American Roads: The Effect of Sport Utility Vehicles and Pickup Trucks on Traffic Safety." *Journal of Law and Economics* 47 (2):333–355.
- 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.

Online Appendix (Not for Publication)

Appendix A Proofs of propositions

Proposition 1. *The first-best allocation can be achieved for a regulation with compliance trading by setting $\sigma'(a) = 0$ and choosing κ so that $\lambda = \phi$.*

Under compliance trading, a single shadow price, denoted λ will prevail. As shown in footnote 8, consumer n 's problem can be written as:

$$\max_{a_n, e_n} U_n = F_n(a_n, e_n) + I_n - P(a_n, e_n) + \lambda \times (e_n - \sigma(a_n) - \kappa).$$

The first-order conditions are:

$$\frac{\partial U_n}{\partial a_n} = \frac{\partial F_n}{\partial a_n} - \frac{\partial P}{\partial a_n} - \lambda \sigma'(a_n) = 0 \quad (9)$$

$$\frac{\partial U_n}{\partial e_n} = \frac{\partial F_n}{\partial e_n} - \frac{\partial P}{\partial e_n} + \lambda = 0. \quad (10)$$

The planner's direct allocation problem (equation 1) is:

$$\max_{a_n, e_n} W = \sum_{n=1}^N \{F_n(a_n, e_n) - C(a_n, e_n) + I_n\} + \phi \sum_{n=1}^N e_n.$$

The first-best optimization conditions are found by differentiation:

$$\frac{\partial W}{\partial a_n} = \frac{\partial F_n}{\partial a_n} - \frac{\partial C}{\partial a_n} = 0 \quad (11)$$

$$\frac{\partial W}{\partial e_n} = \frac{\partial F_n}{\partial e_n} - \frac{\partial C}{\partial e_n} + \phi = 0. \quad (12)$$

Under perfect competition, prices equal marginal costs. Then, it is apparent that the planner's and consumers' first-order conditions are identical if and only if $\sigma'(a_n) = 0$ for all n and $\lambda = \phi$. The endogenous λ will be an increasing function of κ because of the convexity of the cost function. Thus, some value of κ exists for which $\lambda = \phi$. Choosing that value of κ and $\sigma' = 0$ makes the consumers' first-order conditions identical to the planner's first-best.

Corollary 1. *The first-best allocation can be achieved with a subsidy by setting $\sigma'(a) = 0$ and*

$s = \phi$.

The planner's conditions are the same as above. The consumer's problem is now:

$$\max_{a_n, e_n} U_n = F_n(a_n, e_n) + I_n - P(a_n, e_n) + s \times (e_n - \sigma(a_n)).$$

The first-order conditions are:

$$\frac{\partial U_n}{\partial a_n} = \frac{\partial F_n}{\partial a_n} - \frac{\partial P}{\partial a_n} - s\sigma'(a_n) = 0 \quad (13)$$

$$\frac{\partial U_n}{\partial e_n} = \frac{\partial F_n}{\partial e_n} - \frac{\partial P}{\partial e_n} + s = 0. \quad (14)$$

The consumer's conditions match the first-best if and only if $\sigma'(a_n) = 0$ for all n and $s = \phi$.

Proposition 2. *For a linear subsidy, when s is fixed, the second-best attribute slope is:*

$$\hat{\sigma}^{SB} = \frac{s - \phi \left(\sum_n \frac{\partial e}{\partial \hat{\sigma}} \right) / n}{s \left(\sum_n \frac{\partial a}{\partial \hat{\sigma}} \right) / n}.$$

Taking s as fixed, the planner's second-best choice of $\hat{\sigma}$ solves:

$$\max_{\hat{\sigma}} W = \sum_{n=1}^N \{F_n(a_n, e_n) - C(a_n, e_n) + I_n\} + \phi \sum_{n=1}^N e_n.$$

The first-order condition of this problem is:

$$\sum_n \left(\frac{\partial F_n}{\partial e_n} - \frac{\partial C}{\partial e_n} + \phi \right) \frac{\partial e_n}{\partial \hat{\sigma}} + \left(\frac{\partial F_n}{\partial a_n} - \frac{\partial C}{\partial a_n} \right) \frac{\partial a_n}{\partial \hat{\sigma}} = 0. \quad (15)$$

We substitute the optimality conditions from the consumer's problem and pull constants out of the summations:

$$(-s + \phi) \sum_n \frac{\partial e_n}{\partial \hat{\sigma}} + (\hat{\sigma}s) \sum_n \frac{\partial a_n}{\partial \hat{\sigma}} = 0.$$

Dividing by n (to express the summations as averages) and rearranging yields the result.

Proposition 3. For a linear subsidy, when $\hat{\sigma}$ is fixed, the second-best subsidy rate is:

$$s^{SB} = \frac{\phi}{1 - \hat{\sigma} \left(\frac{(\sum_n \frac{\partial a}{\partial s})/n}{(\sum_n \frac{\partial e}{\partial s})/n} \right)} = \frac{\phi}{1 - \epsilon_a^\sigma \frac{\epsilon_s^a}{\epsilon_s^e}},$$

where ϵ_x^y denotes the elasticity of y with respect to x for each pair of variables.

Taking $\hat{\sigma}$ as fixed, the planner's second-best choice of s solves:

$$\max_s W = \sum_{n=1}^N \{F_n(a_n, e_n) - C(a_n, e_n) + I_n\} + \phi \sum_{n=1}^N e_n.$$

The first-order condition is:

$$\sum_n \left(\frac{\partial F_n}{\partial e_n} - \frac{\partial C}{\partial e_n} + \phi \right) \frac{\partial e_n}{\partial s} + \left(\frac{\partial F_n}{\partial a_n} - \frac{\partial C}{\partial a_n} \right) \frac{\partial a_n}{\partial s} = 0.$$

We substitute the optimality conditions from the consumer's problem, just as above, and rearrange.

To get the second equivalence, we substitute the definition of an elasticity and use the fact that, when the standard binds, $e_n = \hat{\sigma} a_n + \kappa$.

Proposition 4. When there is no compliance trading, and the constraint binds for all n , the optimal attribute slope $\hat{\sigma}^{NC}$ in the linear regulation is:

$$\hat{\sigma}^{NC} = \frac{\frac{\sum_n (\lambda_n - \bar{\lambda})(a_n - \bar{a})}{n}}{\phi \left(\frac{\sum_n \frac{\partial a_n}{\partial \hat{\sigma}}}{n} - \bar{a}_n \frac{\sum_n \frac{\partial a_n}{\partial \kappa}}{n} \right)},$$

which is not zero unless a_n is perfectly uncorrelated with λ_n under a flat standard.

The planner solves:

$$\max_{\hat{\sigma}, \kappa} W = \sum_{n=1}^N \{F_n(a_n, e_n) - C(a_n, e_n) + I_n\} + \phi \sum_{n=1}^N e_n.$$

The first-order condition with respect to κ is:

$$\sum_n \left(\frac{\partial F_n}{\partial a_n} - \frac{\partial C}{\partial a_n} \right) \frac{\partial a_n}{\partial \kappa} + \left(\frac{\partial F_n}{\partial e_n} - \frac{\partial C}{\partial e_n} + \phi \right) \frac{\partial e_n}{\partial \kappa} = 0.$$

Using the optimality conditions from the consumer's problem, this can be rewritten as:

$$\sum_n \lambda_n \hat{\sigma} \frac{\partial a_n}{\partial \kappa} + (\phi - \lambda_n) \frac{\partial e_n}{\partial \kappa} = 0. \quad (16)$$

When the constraint is binding, $e_n = \hat{\sigma} a_n + \kappa$. Total differentiation of this constraint yields a relationship between $\partial a_n / \partial \kappa$ and $\partial e_n / \partial \kappa$, namely that $\partial e_n / \partial \kappa = \hat{\sigma} \cdot \partial a_n / \partial \kappa + 1$. Using this substitution and rearranging equation 16 yields:

$$\frac{\sum_n \lambda_n}{n} = \phi \left(1 + \underbrace{\hat{\sigma} \frac{\sum_n \frac{\partial a_n}{\partial \kappa}}{n}}_+ \right), \quad (17)$$

The first-order condition for $\hat{\sigma}$ is:

$$\sum_n \left(\frac{\partial F_n}{\partial a_n} - \frac{\partial C}{\partial a_n} \right) \frac{\partial a_n}{\partial \hat{\sigma}} + \left(\frac{\partial F_n}{\partial e_n} - \frac{\partial C}{\partial e_n} + \phi \right) \frac{\partial e_n}{\partial \hat{\sigma}} = 0.$$

Substituting the consumer's optimality conditions yields:

$$\sum_n (\hat{\sigma} \lambda_n) \frac{\partial a_n}{\partial \hat{\sigma}} + (\phi - \lambda_n) \frac{\partial e_n}{\partial \hat{\sigma}} = 0.$$

Total differentiation of the constraint yields $\partial e / \partial \hat{\sigma} = \hat{\sigma} \partial a / \partial \hat{\sigma} + a$. Substitution yields:

$$\sum_n (\hat{\sigma} \lambda_n) \frac{\partial a_n}{\partial \hat{\sigma}} + (-\lambda_n + \phi) \left(\hat{\sigma} \frac{\partial a_n}{\partial \hat{\sigma}} + a_n \right) = 0.$$

Canceling terms yields:

$$\sum_n -\lambda_n a_n + \phi \hat{\sigma} \frac{\partial a_n}{\partial \hat{\sigma}} + \phi a_n = 0.$$

We then use the definition of the sample covariance of λ_n and a_n to rewrite $\sum_n \lambda_n a_n$ as $\sum_n (\lambda_n - \bar{\lambda})(a_n - \bar{a}) + n^{-1} \sum_n a_n \sum_n \lambda_n$, substitute equation 17 to rewrite the average shadow price, and rearrange. This yields the result.

It is apparent that $\hat{\sigma}^{NC}$ must be non-zero, unless the covariance between the shadow price and the attribute is zero under a flat standard. Otherwise, there is a contradiction.

A.1 Supplementary proofs

We provide here two additional proofs that describe the optimal $\hat{\sigma}$ and λ for the case without compliance trading in a form that shows the relationship to Propositions 2 and 3 (which assume compliance trading) clearly.

Proposition 5. *When there is no compliance trading, and the constraint binds for all n , the optimal average shadow price (denoted $\bar{\lambda}^{NC}$) is:*

$$\bar{\lambda}^{NC} = \frac{\phi}{1 - \hat{\sigma} \left(\frac{(\sum_n \frac{\partial a}{\partial \kappa})/n}{(\sum_n \frac{\partial e}{\partial \kappa})/n} \right)} + \frac{\frac{\sum_n (\lambda_n - \bar{\lambda}_n) (\frac{\partial e_n}{\partial \kappa} - \frac{\partial \bar{e}_n}{\partial \kappa})}{n} - \hat{\sigma} \frac{\sum_n (\lambda_n - \bar{\lambda}_n) (\frac{\partial a_n}{\partial \kappa} - \frac{\partial \bar{a}_n}{\partial \kappa})}{n}}{\hat{\sigma} \frac{\partial \bar{a}}{\partial \kappa} - \frac{\partial \bar{e}}{\partial \kappa}}.$$

The first-order condition (with optimal consumer conditions substituted in) is the same as above, in equation 16. Rearranging yields:

$$\hat{\sigma} \sum_n \lambda_n \frac{\partial a_n}{\partial \kappa} = -\phi \sum_n \frac{\partial e_n}{\partial \kappa} + \sum_n \lambda_n \frac{\partial e_n}{\partial \kappa}$$

Then, we substitute for the $\lambda_n \frac{\partial a_n}{\partial \kappa}$ and $\lambda_n \frac{\partial e_n}{\partial \kappa}$ terms, using the definition of a covariance.

$$\begin{aligned} \hat{\sigma} \left[\sum_n (\lambda_n - \bar{\lambda}_n) \left(\frac{\partial a_n}{\partial \kappa} - \frac{\partial \bar{a}_n}{\partial \kappa} \right) + \bar{\lambda} \sum_n \frac{\partial a_n}{\partial \kappa} \right] &= -\phi \sum_n \frac{\partial e_n}{\partial \kappa} \\ &+ \left[\sum_n (\lambda_n - \bar{\lambda}_n) \left(\frac{\partial e_n}{\partial \kappa} - \frac{\partial \bar{e}_n}{\partial \kappa} \right) + \bar{\lambda} \sum_n \frac{\partial e_n}{\partial \kappa} \right] \end{aligned}$$

Solving for the $\bar{\lambda}$ terms that are outside of the covariance terms yields the following expression for the average value of the shadow price at the optimum:

$$\bar{\lambda} = \frac{\phi}{1 - \hat{\sigma} \left(\frac{(\sum_n \frac{\partial a}{\partial \kappa})/n}{(\sum_n \frac{\partial e}{\partial \kappa})/n} \right)} + \frac{\frac{\sum_n (\lambda_n - \bar{\lambda}_n) (\frac{\partial e_n}{\partial \kappa} - \frac{\partial \bar{e}_n}{\partial \kappa})}{n} - \hat{\sigma} \frac{\sum_n (\lambda_n - \bar{\lambda}_n) (\frac{\partial a_n}{\partial \kappa} - \frac{\partial \bar{a}_n}{\partial \kappa})}{n}}{\hat{\sigma} \frac{\partial \bar{a}}{\partial \kappa} - \frac{\partial \bar{e}}{\partial \kappa}}$$

The average shadow price is monotonically rising in κ , so there will be a unique κ^{NC} that yields the desired average shadow price.

As is the case with Proposition 4, the result is not in closed form. But, the merit of writing

the result in this way is that it is clear when the non-compliance trading case collapses to the same result as the compliance trading solution. If there is no correlation between the shadow price and the *derivatives* of e_n and a_n with respect to the policy’s stringency, then the second term will be zero and the non-trading result collapses to Proposition 2, where the only difference is that the shadow price is an average, not a constant across products.

Proposition 6. *When there is no compliance trading, and the constraint binds for all n , the optimal attribute slope is:*

$$\hat{\sigma}^{NC} = \frac{(\bar{\lambda} - \phi) \frac{\partial \bar{e}_n}{\partial \hat{\sigma}} + \frac{\sum_n (\lambda_n - \bar{\lambda}_n) (\frac{\partial e_n}{\partial \hat{\sigma}} - \frac{\partial \bar{e}_n}{\partial \hat{\sigma}})}{n}}{\bar{\lambda} \frac{\partial \bar{a}_n}{\partial \hat{\sigma}} + \frac{\sum_n (\lambda_n - \bar{\lambda}_n) (\frac{\partial a_n}{\partial \hat{\sigma}} - \frac{\partial \bar{a}_n}{\partial \hat{\sigma}})}{n}}.$$


This is derived in the same way as the prior result. Start with the same first-order condition for $\hat{\sigma}$, which is equation 15. Substitute in the optimality conditions for the consumer. Substitute in the definition of covariance to eliminate the products of the shadow price and the derivatives of a_n and e_n with respect to $\hat{\sigma}$. Solving for $\hat{\sigma}$ then yields the result.

Again, this result is not in closed form. It is intended to highlight the relationship between the non-trading case with Proposition 3. When there is no covariance between shadow prices and the derivatives of the characteristics with respect to the attribute slope, the covariance terms drop out and the formula collapses to the equation in Proposition 3 (the case with trading). As long as there is a correlation between the shadow price and responses to the policy, the optimal policy will involve some degree of attribute-basing.

Appendix B Welfare implications of notched attribute-based policies

Our theory models smooth attribute-based functions, that is, cases where the target function $\sigma(a)$ is everywhere differentiable in a . The Japanese policy that we analyze empirically has notches, so that $\sigma(a)$ is a step-function. Here, we briefly argue that the welfare implications from our theory carry over to notched policies. We first consider “single notched” systems, like the Japanese corporate

Figure 7: Isocost Curve with a Notched Attribute-Based Subsidy



average fuel-economy standard, and then discuss “double notched” systems like the model-specific subsidy in Japan. In both cases, we discuss a subsidy policy rather than a regulation for notational ease.

How does a single notched policy, where $\sigma(a)$ is a step function but the marginal incentive for e is smooth, affect choice? We provide initial intuition graphically. Figure 7 shows an isocost curve, that is, the set of values of a and e for which a consumer spends a constant amount on the durable net of the subsidy, $P(a, e) - s \times (e - \sigma(a))$. The figure is drawn with several notches, at a' , a'' and a''' . 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 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 subsidy (i.e, if $S(a, e) = se$). 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 only the indirect change due to a —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 changed at all by the attribute-basing 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' , a'' or a''' . That is, all of the distortion is evident from those who “bunch” at the notch points.

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 subsidy jumps by amount $\tau > 0$. Then, the tax function can be written as:

$$S(a, e) = \begin{cases} s \cdot e & \text{if } a < a' \\ s \cdot e + \tau & \text{if } a \geq a'. \end{cases} \quad (18)$$

Denote by (a^*, e^*) the bundle chosen by a consumer facing a Pigouvian tax of $s \cdot e$. 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 $\mathcal{L} = F(a, e) - P(a, e) + I - G + se + \tau + \mu[a - a']$, 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:

$$\max_e = F(a', e) - P(a', e) + I - G + se + \tau, \quad (19)$$

which has the same first order condition for e as the case without τ . Just as in the smooth case, any distortion to the choice of e comes through “general equilibrium” effects, through which a distortion in a shifts the marginal costs and benefits of e , which might result in a change in e .

The distortion in a will be analogous to a traditional Harberger triangle and thus rising in τ^2 . The consumer will choose $\tilde{a} = a'$ if and only if:

$$-\tau > P(a', \tilde{e}) - P(a^*, e^*) - (F(a', \tilde{e}) - F(a^*, e^*)), \quad (20)$$


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 has the same intuition as a traditional Harberger triangle, just as in the smooth case.

For our purposes, the 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.

B.1 Double notched policies

We next briefly describe the incentives created by a double notched policy, where the subsidy is not everywhere differentiable in e or a . 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 subsidy for such a system can

Figure 8: Isocost curve with notches for both a and e



be described algebraically as:

$$S(a, e) = \begin{cases} s_1 & \text{if } e > e'' \\ s_2 & \text{if } e'' > e > e' \text{ and } a > a' \\ 0 & \text{otherwise.} \end{cases} \quad (21)$$

An isocost curve for this case is shown in Figure 8. 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 s_1 for any allocation above e'' . Note that there are large regions 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 s_1 and s_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 $s_1 \neq s_2$. In many policy examples, $s_1 = s_2$, 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 8, 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.

Appendix C Adjustment Cost Functions in Section 4

This section provides a detailed description of how we derive the adjustment cost functions in Section 4. We analyze data before and after the policy change. For vehicle n , we denote a_n and e_n as the second-period characteristics and a_n^o and e_n^o as the first-period characteristics. We make the assumption of perfect competition. Then, using the notations in Section 2, the welfare for vehicle n , omitting regulatory incentives and dropping the numeraire, can be written as $F_n(a_n, e_n) - C(a_n, e_n)$.

First, consider a simple case, in which there is no fleet-average compliance regulation. Before the policy change, there is no regulation. After the policy change, there is a car-specific subsidy for cars that meet the standard ($e_n \geq \sigma(a_n)$). Because the first-period characteristics a_n^o and e_n^o are at the private optimum, any deviation from that point creates a loss, which is $L_n = F_n(a_n, e_n) - C(a_n, e_n) - [F_n(a_n^o, e_n^o) - C(a_n^o, e_n^o)]$. The second-period optimization problem for product n is then to choose the a_n and e_n values that maximize the loss plus the subsidy:

$$W_n = F_n(a_n, e_n) - C(a_n, e_n) - [F_n(a_n^o, e_n^o) - C(a_n^o, e_n^o)] + \tau \cdot 1(e_n \geq \sigma(a_n)) + \varepsilon_{nz},$$

where $L_n = F_n(a_n, e_n) - C(a_n, e_n) - [F_n(a_n^o, e_n^o) - C(a_n^o, e_n^o)] \leq 0$ and L_n is peaked at $(a_n, e_n) = (a_n^o, e_n^o)$. That is, no changes in a and e would produce the lowest possible loss (defined to be zero). This motivates us to begin with a quadratic functional form for L_n in our estimation. In the first specification in Section 4, we use a quadratic function: $L_n = \alpha(a_n - a_n^o)^2 + \beta(e_n - e_n^o)^2 + \gamma(a_n - a_n^o)(e_n - e_n^o)$ and estimate:

$$W_n = \alpha(a_n - a_n^o)^2 + \beta(e_n - e_n^o)^2 + \gamma(a_n - a_n^o)(e_n - e_n^o) + \tau \cdot 1(e_n \geq \sigma(a_n)) + \varepsilon_{nz},$$

which is equation 7 in Section 4.

Second, consider the case with fleet-average compliance regulation. We denote λ and λ^o as the shadow prices of the fleet-average regulation at the second period and first period. Then, the payoff to first period choices is $F_n(a_n^o, e_n^o) - C(a_n^o, e_n^o) + \lambda^o \times (e_n^o - \sigma^o(a_n^o))$. In the second period, there is a new fleet-average regulation and vehicle-specific subsidy policy. Then, the payoff to second-period choices is $F_n(a_n, e_n) - C(a_n, e_n) + \lambda \times (e_n - \sigma(a_n)) + \tau \cdot 1(e_n \geq \sigma(a_n))$. The second-period optimization problem for product n is then to choose the a_n and e_n values that maximize the objective function:

$$\begin{aligned} V_n &= F_n(a_n, e_n) - C(a_n, e_n) - [F_n(a_n^o, e_n^o) - C(a_n^o, e_n^o)] \\ &\quad + \lambda(e_n - \sigma(a_n)) - \lambda^o(e_n^o - \sigma^o(a_n^o)) + \tau \cdot 1(e_n \geq \sigma(a_n)) + \varepsilon_{nz} \\ &= f_n + \lambda(e_n - \sigma(a_n)) - \lambda_n^o(e_n^o - \sigma^o(a_n^o)) + \tau \cdot 1(e_n \geq \sigma(a_n)) + \varepsilon_{nz}, \end{aligned}$$

where $f_n \equiv F_n(a_n, e_n) - C(a_n, e_n) - [F_n(a_n^o, e_n^o) - C(a_n^o, e_n^o)]$. The problem with estimating this equation directly is that f_n is not peaked at $(a_n, e_n) = (a_n^o, e_n^o)$, and we wish to use a functional form approximation that requires knowing the location of the peak. That is, no changes in a and e would not necessarily produce the lowest possible loss for f_n . This is because (a_n^o, e_n^o) is the optimal in the presence of the first-period regulation, rather than the private optimal in the absence of regulation.

To address this problem, we rewrite V_n by adding and subtracting $\lambda^o(e_n - \sigma^o(a_n))$. Note that this is a mixed object—it is the second-period choice of a and e put into the first-period policy function and shadow price:

$$\begin{aligned} V_n &= F_n(a_n, e_n) - C(a_n, e_n) - [F_n(a_n^o, e_n^o) - C(a_n^o, e_n^o)] \\ &\quad + \lambda(e_n - \sigma(a_n)) - \lambda^o(e_n^o - \sigma^o(a_n^o)) + \tau \cdot 1(e_n \geq \sigma(a_n)) + \varepsilon_{nz} \\ &\quad + \lambda^o(e_n - \sigma^o(a_n)) - \lambda^o(e_n - \sigma^o(a_n)) \\ &= [F_n(a_n, e_n) - C(a_n, e_n) + \lambda^o(e_n - \sigma^o(a_n)) - F_n(a_n^o, e_n^o) + C(a_n^o, e_n^o) - \lambda^o(e_n^o - \sigma^o(a_n^o))] \\ &\quad + \tau \cdot 1(e_n \geq \sigma(a_n)) + \lambda(e_n - \sigma(a_n)) - \lambda^o(e_n - \sigma^o(a_n)) + \varepsilon_{nz} \\ &= g_n + \tau \cdot 1(e_n \geq \sigma(a_n)) + \lambda(e_n - \sigma(a_n)) - \lambda^o(e_n - \sigma^o(a_n)) + \varepsilon_{nz} \end{aligned}$$

where $g_n \equiv [F_n(a_n, e_n) - C(a_n, e_n) + \lambda^o(e_n - \sigma^o(a_n)) - F_n(a_n^o, e_n^o) + C(a_n^o, e_n^o) - \lambda^o(e_n^o - \sigma^o(a_n^o))]$.

Importantly, g_n is peaked at $(a_n, e_n) = (a_n^o, e_n^o)$. It is the adjustment cost function **holding constant the old policy**. Because (a_n^o, e_n^o) is the optimum in the presence of the old policy, $g_n \leq 0$ and g_n is peaked at $(a_n, e_n) = (a_n^o, e_n^o)$. That is, no changes in a and e would produce the lowest possible loss. This motivates us to have a quadratic functional form for g_n in our second specification in Section 4. We use a quadratic function: $g_n = \alpha(a_n - a_n^o)^2 + \beta(e_n - e_n^o)^2 + \gamma(a_n - a_n^o)(e_n - e_n^o)$ and estimate:


$$V_n = \alpha(a_n - a_n^o)^2 + \beta(e_n - e_n^o)^2 + \gamma(a_n - a_n^o)(e_n - e_n^o) \\ + \tau \cdot 1(e_n \geq \sigma(a_n)) + \lambda(e_n - \sigma(a_n)) - \lambda^o(e_n - \sigma^o(a_n)) + \varepsilon_{nz},$$

which is equation (8) in Section 4. Finally, we can recover f_n from g_n . Note that $g_n = f_n - [\lambda^o(e_n^o - \sigma^o(a_n^o)) - \lambda^o(e_n - \sigma^o(a_n))]$. Therefore, once we have parameter estimates for g_n and λ^o , we can recover $f_n = g_n + [\lambda^o(e_n^o - \sigma^o(a_n^o)) - \lambda^o(e_n - \sigma^o(a_n))]$. This f_n provides the adjustment cost function for a_n and e_n excluding the effects of the old regulation.


Appendix D Additional Figures and Tables

Figure A.1: Fuel-Economy Standard and Histogram of Vehicles: Kei-Cars (small cars)

Panel A. Years 2001 to 2008 (Old Fuel Economy Standard Schedule)




Panel B. Years 2009 to 2013 (New Fuel Economy Standard Schedule)




Note: “Kei-car” is a Japanese category of small vehicles; the displacement of kei-cars have to be less than 660 cc. Most kei-cars are not exported to other countries. Panel A shows the histogram of vehicles from 2001 to 2008, where all vehicles had the old fuel economy standard. Panel B shows the histogram of vehicles from 2009 to 2013, in which the new fuel economy standard was introduced.

Figure A.2: Fuel Economy and Weight before and after the Policy Change for Vehicles that Did Not Receive a Subsidy

Panel A. Vehicles that did not receive a subsidy but bunched at weight notches



Panel B. Vehicles that did not receive a subsidy and did not bunch at weight notches



Note: This figure shows each vehicle's fuel economy and weight before and after the introduction of the new subsidy that was applied to each vehicle individually. The scatterplot shows each car's starting values of fuel economy and weight in 2008—the year before the policy change. We also show “arrows” connecting each car's starting values in 2008 with its values in 2012. The figure also includes three step functions that correspond to the three tiers of the new incentive's eligibility cutoffs.