

**Distorted Technical Change?
Evidence from New Vehicle Launches in the Japanese Automobile Industry¹**

Yoshifumi Konishi[†]
Department of Policy and Planning Sciences
Faculty of Engineering, Information and Systems
University of Tsukuba

Kong Joo Shin
Department of Urban and Environmental Engineering
School of Engineering, Kyushu University

Shunsuke Managi
Department of Urban and Environmental Engineering
School of Engineering, Kyushu University

Preliminary Version: May 18, 2018

¹We thank Corbett Grainger, Sumeet Gulati, and other seminar participants at the University of British Columbia, the National Institute of Environmental Studies (Japan), and the Research Institute of Economy, Trade, and Industry (Japan) for their helpful comments.

[†]Corresponding author. Mailing address: 1-1-1 Tennodai, Tsukuba, Ibaragi, 305-8573, Japan. E-mail: ykonishi@sk.tsukuba.ac.jp. Phone: +81-29-853-5072. Fax: +81-29-853-5070.

Abstract: We empirically examine the distortionary impact of an attribute-based regulation on technical change, in the context of new product launches in the Japanese automobile industry. Under the Japanese regulation, fuel economy standards are a step function of curb weight, and their stringency levels vary substantially over time across weight bins of different sizes. We explicitly exploit these quasi-experimental variations in the difference-in-difference-in-differences framework to control for confounders that may be correlated with regulatory assignment. We find strong evidence in support of our theoretical prediction: An attribute-based regulation distorts technical change when it creates trade-offs between the targeted and secondary attributes that differ from technically feasible trade-offs. We also demonstrate that bunching behavior reported elsewhere was evident only in reporting to the government, but not in product offerings in the market.

JEL Codes: D22, H23, K32, L62, Q48, Q55

Key Words: Attribute-based regulation, automobile, difference-in-difference-in-differences, energy policy, fuel economy regulation, technical change, technology policy

1. Introduction

Since Porter (1991), economists have long been interested in empirically examining the effect of environmental regulation on technical change [e.g., Newell, Jaffe, and Stavins (1999); Popp (2002); and Calel and Dechezleprêtre (2016)]. The literature to date, however, has primarily focused on the *direct* (or *intended*) effect of environmental regulation: i.e., the effect of a regulation-induced increase in the implicit price of pollution on technical change in sectors that use pollution as a factor of production [Copeland and Taylor (1994)]. Real-world environmental regulation, however, often entails design features that offer ‘loopholes’ that may not be necessarily ideal in first-best settings [Anderson and Sallee (2011), Sallee and Slemrod (2012), and Ito and Sallee (forthcoming)]. These design features may create misguided incentives for innovation, the effects of which can potentially persist over time via the market size effect of technical progress [Acemoglu (2002), Acemoglu and Linn (2004), Acemoglu *et al.* (2012), and Aghion *et al.* (2016)].

We empirically examine ‘attribute-basing’ in fuel economy regulation for automobiles as a primary example of such (unintended) effect of regulatory loopholes on technical change. Fuel economy regulation is often attribute-based (Ito and Sallee, forthcoming): i.e., it relies on a secondary attribute that is not the direct target of the regulation. For example, fuel economy standards are a function of vehicle footprint (the U.S.) or curb weight (Japan and the EU).¹ In theory, an attribute-based regulation can distort firm’s incentives in product offerings, for it may create trade-offs between the attributes that differ from firm’s technically feasible trade-offs. This distortion in product offerings can also translate into distorted technical change because firms may need to choose the level of technical upgrade in either of these attributes when the attributes of the products chosen to be offered lie outside the currently feasible technology set.²

To test this economic prediction, we exploit the unique quasi-experimental setup created due to the Japanese weight-based fuel economy regulation. Under the regulation, the fuel economy standards are a step function (or notched function) of curb weight: i.e., vehicles are classified into discrete weight bins with varying levels of fuel economy standards. Importantly, when revising the standards in 2007, the regulatory authority chose narrower weight bins, creating two or more weight bins within each of the old weight segments. Consequently, we have substantial variation, in terms of stringency and width, across weight bins. We translate this variation into the variation in attribute trade-offs by computing the

¹To be precise, the EU standards target carbon emissions per kilometer of driving distance rather than fuel economy per se, but the former is a direct function of the former.

²By design, fuel economy regulation requires firms to offer vehicles outside their currently feasible technology set.

‘slope’ of each weight bin, defined as a decrease in the fuel economy standard per unit of weight increase. These slopes of the regulation are a convenient measure of the attributes trade-offs imposed by the fuel economy regulation. In principle then, we should be able to test our hypothesis by comparing the outcomes of car models assigned to different regulatory slopes. The key here is how to control for vehicle-level unobservables as well as market-level confounders.

To that end, we employ a difference-in-difference-in-differences (DDD) research design, exploiting the variation in regulatory slope and stringency across weight bins. Our control structures are three-fold: (a) cross-sectional with models assigned to low-slope weight bins as a control group, (b) temporal with the pre-2007 period as a control period, and (c) within-group cross-sectional with models assigned to non-stringent segments as an additional control. This last control serves two purposes. First, our estimate would be unbiased under a milder assumption than the conventional common-trend assumption. Second, it allows us to compare outcomes of comparable groups that faced roughly the same regulatory stringency, but different regulatory trade-offs. This way, we are able to attribute the difference in outcomes solely to the difference in *attribute trade-offs* (rather than in stringency levels per se) created under the fuel economy regulation. We implement this DDD strategy using vehicle characteristics data for all domestic passenger vehicles offered between 2004 and 2012, excluding electric, diesel, and hybrid cars as well as those launched in the interim regulatory period 2007-2009.

Our DDD estimate indicates a sizable, statistically significant distortionary effect on technical progress. The TPF for those assigned to high-slope weight bins would have lied strictly above the observed TPF if they had not been assigned to these weight bins. With the Cobb-Douglas specification for the TPF, this translates to a slowdown of fuel economy improvements by roughly 20 percentage points. The economic significance of this impact can be cast in light of the work by Knittel (2011). Using variant-level data from the U.S. automobile industry, Knittel estimates that U.S. passenger cars could have improved fuel economy by roughly 60% over the 25-year period between 1980 and 2006 if their curb weights (and other attributes) had stayed at the 1980 level. Employing a similar exercise along with our DDD estimate suggests that the Japanese passenger cars would have improved fuel economy by roughly 60% just over the 8-year period between 2004 and 2012 if they had not been assigned to the high-slope bins and their vehicle weight had stayed the same at the 2004 level — the size of technical progress comparable to that over the 25-year period in the U.S.

Our results have important implications for attribute-basing in technology regulation as well as for other institutional arrangements that utilize several attributes for standard-

setting. Energy efficiency regulations around the world are often attribute-based. For example, fuel economy regulations in China, EU, and Japan as well as energy efficiency labels and standards for buildings, consumer electronics, and home appliances have similar features. Attribute-based regulations are often preferred over uniform regulations in the regulatory arena for efficiency as well as equity concerns.³ Our analysis demonstrates that design features of such technology regulations matter, not only for product differentiation but also for technical progress. Our policy advice in this case is clear. On one hand, to assure no bias in technical progress, regulators need to make regulatory trade-offs between targeted and secondary attributes as close as technically feasible trade-offs. On the other hand, regulators could manipulate the regulatory trade-offs in a way to favor or disfavor a particular direction for technical progress. This advice is useful *both* when credit trading is available *and* when it is not. Furthermore, our results can be cast more broadly in light of other institutional arrangements that use two or more attributes for standard-setting. For example, admissions screening for secondary and higher education often entails trade-offs between test scores on multiple subjects. In some cases, only the overall test scores matter so that a test score on one subject (say, English) can be traded one-to-one with a test score on another subject (say, Mathematics). In other cases, the passing test score for one subject may be a function of the test scores on other subjects. The admissions standards then effectively create trade-offs for test takers. By manipulating these trade-offs, institutions can favor or disfavor the level of students' efforts toward one subject over others. Other examples include multiple-attribute considerations in auctions, hiring, and other business contracts.

This manuscript nicely complements the large literature investigating the effects of regulatory 'loopholes' (e.g., Anderson and Sallee, 2011; Sallee and Slemrod, 2012; and Ito and Sallee, forthcoming). While they differ in their research focus and design, they all provide evidence that firms do exploit the loopholes when they reduce (marginal) costs of compliance. The regulatory loopholes can take a variety of forms, including notches (e.g., Sallee and Slemrod, 2012), a special exemption (e.g., Anderson and Sallee, 2011), and a secondary attribute (e.g., Ito and Sallee, forthcoming), of which taxes/subsidies or technology standards are scheduled. Sallee and Slemrod (2012) examine a notched schedule in the U.S. Gas Guzzler Tax and show that automakers manipulated fuel economy ratings so as to qualify for lower tax rates. Our empirical strategy does exploit a similar behavioral response we observe for the notched schedule of the Japanese weight-based fuel economy standards. Our point,

³For example, engineering estimates suggest that fuel economy ratings of vehicles decline with vehicle footprints. The U.S. CAFE standards are calibrated based upon this property. As a result, the marginal costs of compliance with the standards are likely to be close to each other over all vehicle footprint levels. Furthermore, precisely because the costs of compliance are similar across segments, the footprint-based standards are intended to avoid unequal distribution of costs on firms and consumers.

however, is that such an incentive to exploit the notches is constrained by trade-offs that are technically feasible, and hence, the variation in the trade-offs can be used to identify the distortionary impact of the regulatory loophole. Anderson and Sallee (2011) estimate the marginal cost of compliance with the U.S. CAFE regulation, exploiting the fact that firms can effectively relax the CAFE constraint by equipping a vehicle with flexible-fuel capacity, and applying insight that firms equalize their marginal costs of compliance across different compliance strategies. Our results add to their discussion another important dimension: Such distortion may have an second-order impact on technical change.

Our work is most closely related to the work by Ito and Sallee (forthcoming). They examine attribute-based regulations both theoretically and empirically, using model-level data on Japanese passenger cars. They demonstrate that significant bunching occurred at the weight cutoffs and that the bunching resulted in a large welfare loss. Our work offers three important contributions related to their work. First, while Ito and Sallee estimate the incidence of bunching in the vehicle weight distribution, we focus on its distortionary impact on the technical progress in trade-offs between fuel economy and vehicle weight. This difference is important since ‘bunching’ in the weight distribution can occur without any distortion in technical frontiers. In this sense, their estimated welfare loss due to the regulation is only of the first order impact; Ours offer evidence on the second-order impact on technical change.

Second, we construct a new data set using the web catalogue of one of the largest car dealers, Carsensor, in Japan. The catalogue data are reported at the variant level far finer than those used in Ito and Sallee. They used the fuel economy data published each year by the Ministry of Land, Infrastructure, and Transportation (MLIT). The MLIT data report vehicle attributes only at the model/configuration level, with only the range (i.e., minimum and maximum) of curb weights for most of the vehicle configurations. The weight range can be as large as 200 kg, averaging at around 35 kg. In contrast, the web catalogue contains all vehicle offerings at the variant level *and* each vehicle variant is reported with an exact and unique curb weight. We demonstrate that the incidence of bunching largely disappears once we construct the weight distribution using these actual vehicle offerings. Our analysis suggests that what Ito and Sallee find is most likely the evidence of manipulation by automakers in *reporting* their vehicle weight ranges, but not necessarily the evidence of weight manipulation in their actual vehicle offerings. We do, however, report on the evidence in support of Ito and Sallee — our results indicate that vehicle models assigned to high-slope weight bins increased their weights more than those assigned to low-slope bins.

Third, our three-fold control structures enable us to isolate the impact of the fuel economy regulation from that of the tax incentives offered during the same study period. During the

study period (2009-2012), the tax incentives were, in fact, based on the old 2001 standards, not the new 2007 standards. Hence, the automakers faced double standards during this period in their product offerings: the 2001 standards for eco-car subsidy/tax credits and the 2007 standards for the fuel economy regulation. We demonstrate that the incidence of bunching in the MLIT reporting data indeed correspond to the weight cutoffs of *the 2001 standards*, instead of the 2007 standards. Unlike Ito and Sallee, we use this regulatory setup explicitly in our identification — by constructing comparable subsamples within each of the 2001 weight categories, we get around the confounding effect of tax incentives, which should affect these subsamples the same way.

The paper proceeds as follows. The next section revisits the theory of attribute-based regulation, and sets up our empirical research, introducing the concept of technology possibility frontiers. Section 3 describes the regulatory background, introduces our data set, and highlights the distinction between ours and the data used in Ito and Sallee (forthcoming). Section 4 discusses our identification and estimation strategy. The results are discussed in Section 5. Section 6 discusses the welfare implications of our empirical findings, highlighting important differences from the previous studies. The last section concludes.

2. Empirical Framework

2.A. Revisiting the Theory of Attribute-based Regulation

Using a static model of attribute-based regulation, Ito and Sallee (forthcoming) demonstrate that (1) in the presence of (efficient) credit trading, no attribute-basing (i.e., a flat standard) is optimal, but (2) some attribute-basing (i.e., a sloped standard) is optimal in its absence. Most importantly, their model clarifies that it is *not* optimal to perfectly equalize the marginal costs of compliance, highlighting the importance of striking a balance between marginal cost harmonization versus bias minimization in firm’s attribute choice.

We introduce the concept of a technology possibility frontier (TPF) into the analytical framework. The concept is somewhat implicit in Ito and Sallee and other related literature, yet has not been explicitly addressed. The concept originates from Knittel’s work (2011), which shows that technical trade-offs exist between fuel economy and other vehicle attributes for automobiles in the U.S. market and that the technical trade-offs change over time as firms’ technologies improve over time. Indeed, we see a similar, remarkable shift in the technical trade-offs in the Japanese automobile industry over the last 25 years. In 1990, the (unweighted) average fuel economy of all Japanese passenger cars was roughly 13.1 km/L. In 2015, that number increased by more than 70% to 22.3 km/L. This improvement in fuel

economy did not come from downsizing vehicle weight. Indeed, the average curb weight increased by roughly 10% from 1,169 kg in 1990 to 1,293 kg in 2015. Following Knittel (2011), **Figure 1** exhibits technical trade-offs between fuel economy and curb weight for Toyota’s passenger vehicles offered between 1990 and 2015, demonstrating its substantial technical change over the last 25 years.

From this empirical regularity, it seems natural to define the technology possibility set (TPS) as the set of technically feasible trade-offs between product attributes, and each firm can choose its vehicle offerings only from this set in the short run. The TPF is then the upper envelope of the TPS. Note that the TPF does not describe technical trade-offs that induce the same production cost or profit level.⁴ The TPF, instead, gets at attributes trade-offs that are technically feasible when technical inputs are used most efficiently given the technology capital. Hence, the TPF is analogous in spirit to the conventional production possibility frontier (PPF). While the PPF is a set of (most efficient) input combinations that are feasible to produce one unit of a product, the TPF is a set of (most efficient) attribute combinations that are feasible to design a product. Of course, firms may design a product that lies strictly below this frontier, just like firms may produce a product below the PPF. But under some regularity condition, no profit-maximizing firms would design a product that lies in the interior of the technology possibility set. Therefore, the TPF in any given period can be estimated using its vehicle offerings because demand-side factors (e.g., fuel prices, preferences, and tax incentives) can only affect the locations of these offerings along the TPF in the short run. Technical change can then be identified by changes in either the shape or the level of the TPF. In what follows, we demonstrate the importance of the TPF in understanding the distortionary impact of attribute-based regulation.

Our model follows the empirical industrial organization literature, which substantiates the importance of imperfect competition, not only for markup pricing, but also for endogenous product choice and technology upgrade. Consider an automobile industry consisting of M firms. Each firm m produces a unique product, which we treat as fixed. We assume away the multi-product nature of the automobile manufacturer to maintain our focus. Each product can be fully described by two-dimensional product attributes (f, w) , where f represents ‘fuel economy’ and w represents ‘vehicle weight’ for ease of interpretation. Consider the two-period decision of firm m choosing the next-period product attributes (f^1, w^1) conditional on the current-period product attributes (f^0, w^0) . All of the economic rents that result from the current-period choice (f^0, w^0) are treated as ‘sunk’ at the time of choosing next-period

⁴Knittel (2011) defines technically feasible trade-offs as the iso-cost curve in the attribute space. However, marginal cost of production (or profit per unit) should, in principle, differ substantially along this technology frontier, as evident in studies that estimate product-level marginal costs in Japan or the U.S. [see Konishi and Zhao (2017) on the former and Berry, Levinsohn, and Pakes (1995) on the latter].

attributes (f^1, w^1) .

Suppose that firms compete in three steps. In the first stage, each firm chooses the level of investment in technical capital $s \geq 0$, which shifts up the technology possibility frontier:

$$f = T(w, s).$$

In the second stage (at the beginning of the next period), they compete in product choice (f, w) . Then in the third stage, firms compete in prices in the Bertrand manner given the consumer demand and the second-stage product profiles $\{f_m, w_m\}_{m=1, \dots, M}$. Let us simplify our analysis by assuming that a unique (pure-strategy) Bertrand-Nash equilibrium of the third-stage price competition exists.⁵ We describe the third-stage product-specific profit for each firm m by $\pi(f, w; \Omega)$ given Ω , where Ω denotes a collection of other firms' product strategies. Note that the cost of producing a product with attributes (f, w) is already part of π . Then in the second stage, each firm m chooses (f, w) so as to maximize:

$$\pi(f, w; \Omega) - c(s), \quad \text{subject to } f \leq T(w, s),$$

where c is the fixed cost of investment, which is sunk at the time of choosing product attributes. The current-period technology capital is normalized to zero, so the level of investment is conveniently identified with the next-period technology capital s . The regulator sets an attribute-based regulation R , which mandates $f \geq R(w)$, in stage 'zero' before firms engage in this three-stage competition. Our interest lies in understanding the distortionary impact of attribute-based regulation on firm's choice on (f, w) and s .

To that end, we make further simplifying assumptions on π , c , T and R . (A1) The technology possibility function T is (locally) linear and strictly decreasing in w , with $\partial^2 T / \partial w \partial s = 0$. That is, technical upgrade does not affect the shape of the TPF. (A2) The profit function π is increasing in both f and w at the decreasing rate, and the iso-profit curve $\{(f, w) : \pi(f, w; \Omega) = a\}$ is strictly convex in (f, w) . (A3) The cost of technical upgrade c is increasing in s at the increasing rate. (A4) R is (locally) linear and decreasing in w .

The linearity of the TPF and the regulatory constraint in (A1) and (A4) is not as restrictive as it may appear. As discussed in the subsequent sections [and in Knittel (2011)], linear regression is surprisingly well fit to observed attributes in logged values. Moreover, firms are likely to face approximately linear technical trade-offs in the neighborhood of the pre-existing car model. In theory, it is known that (A2) may not hold in general. Without it, however, one cannot solve for or characterize the equilibrium. Assumption (A3) is

⁵We are fully aware that, unfortunately, such a unique Nash equilibrium may not exist. Nonetheless, we make this assumption to focus on the essence of our analysis.

a standard regularity condition. Under these assumptions, we obtain the proposition that characterizes the distortionary nature of attribute-based regulation (its proof is available in the appendix).

Proposition: Suppose (A1)-(A4) hold. Then an attribute-based regulation does not distort firms' incentives neither on technical upgrade s nor attribute choice (f, w) if and only if the slope of the regulatory constraint is the same as that of the technology possibility frontier. The firms have incentives to increase vehicle weight w if the slope of the regulatory constraint is larger than that of the technology possibility frontier and to decrease vehicle weight w if the reverse holds. The firms have incentives to invest less in technology capital s in either case.

The proposition states that (1) not every attribute-basing is distortionary, but (2) the distortionary incentives depend on the slope of the regulation relative to that of TPF, and (3) the distortion translates into the distortion in technical change. **Figure 2** helps us illustrate the idea behind this proposition. The current product offering, denoted O , is on the solid black line, which represents the current-period TPF. The blue dashed lines represent firm's third-stage iso-profit curves, holding Ω . Without the regulation, any profit-maximizing firm should choose (f^1, w^1) , labeled A , at the tangency between an iso-profit curve and the next period's TPF (the solid green line). Hence, (f^1, w^1) is uniquely pinned down given s and Ω . The firm should choose the level of technical upgrade s such that the marginal cost of doing so equals the marginal increase in profits. Now, let us see the impact of the regulation. Let the solid red line represent the attribute-based regulation. We draw the case where R is steeper than T , and hence, by regulatory design, R must cut through T from the above. The firm must choose the product attributes (f^1, w^1) that lie on R . (This assumes the regulation is binding, which would not be the case if the intersection B lies to the left of A , in which case the firm would continue to choose A). Let pick a point B that lies on the intersection of R and the TPF that would be realized under no regulation. It is easy to see that the firm would never pick this attribute bundle B under the regulation. To see this, let us draw another iso-profit curve that goes through B . Then by construction, this iso-profit curve lies below the iso-profit that goes through A , and hence. This means that bundle B achieves lower profits at the same cost of technical upgrade. In other words, the constrained profit is always lower at every technology level s . This lowers the incentive to invest in s . The firm must achieve the optimal bundle with lower investment in s while also meeting the regulation (i.e., along R). It is clear then that such a bundle lie to the right of A , like C . The reverse holds when R is flatter than the TPF.

2.B. From Theory to Empirics

Our approach to testing the theory presented in the previous subsection is primarily data-driven. Japan’s fuel economy standards are set at a *variant level* whereas their enforcement is based on sales-weighted averages at a *manufacturer level* (see **Subsection 3.A.** for a more detailed account of the regulation). Since the regulation is enforced at the manufacturer level, we would ideally model manufacturers’ strategic incentives to offer different variants of different car models in different years explicitly, fully endogenizing both pricing and product choice (e.g., Seim, 2006; Hitsch, 2006, Fan, 2013, Crawford *et al.*, 2015, Wollmann, forthcoming). However, such structural modeling of endogenous product choice requires demand-side information that is far more detailed than we have at hand. Since we are interested in the effect of the fuel economy standards that are imposed at a variant level, we need demand-side information that can vary at a variant level. With more than 1,000 variants offered each year, we lack enough sources of variation to separately identify the influences of variant-level demand factors from those of the regulation in the structural framework.

We take a simpler approach, and focus on the reduced-form estimate of the impact of the standards on firms’ technology possibility frontiers, exploiting policy-induced variations across weight segments over time in the difference-in-difference-in-differences (DDD) research design. Following Knittel (2011) (in spirit), we define TPF as follows. Fuel economy f of vehicle variant i of manufacturer m of vehicle model j in time t is a function of vehicle weight w , a vector of other observable product attributes \mathbf{x} , and a variable s that expresses the level of technology capital:

$$f_{imt} = T(w_{imt}, \mathbf{x}_{imt}; s_{jmt}). \quad (1)$$

The function is allowed to vary by firm and year to incorporate differences in technical efficiency across firms and over time.

This empirical model relies on three identifying assumptions. First, we are assuming that in the short run (i.e., in each year), firms can only choose variants of their car models on their technology possibility curves given their technology capital and engine type. This assumption enables us to recover the TPF for each manufacturer in each year from the variant-level vehicle characteristics data without the need to explicitly model firm’s or consumer’s choices. Empirical regularities found in Knittel (2011) and our data seem to support the validity of this assumption. Given a combustion engine type (i.e., diesel, electric, fuel, and hybrid), the technical attributes trade-offs seem rather stable over time — the curves that represent the technical trade-offs show persistent patterns over time, with changes in the level of the curves over time.

Our second assumption is that technology capital exists at the car-model level. Firms often delegate development of a car model to a specific group of engineers in the form of a division or a team, and the group of engineers apply and accumulate the knowledge/technology capital in designing the car. Hence, firm’s technology frontier can vary at the model level, at least in the short run, though the technology capital acquired through developing a model will be shared across car models that share the same platform in the intermediate term, and eventually across all car models within the firm in the long run. Toyota, for example, developed a hybrid system through development of its famous Prius (debut in 1997), yet it was only four years later Toyota used that system in another car model, Estima, in 2001. For the same token, we presume that it generally takes a few years for a superior fuel combustion system to be applied in other car models. We emphasize here this assumption neither imply nor require that technology capital does not exist at the firm level or segment level. All we require is the existence of some technology capital at the model level.⁶ Given this nature of technical progress, we posit that firms choose the level of model-specific technology capital in response to model-level regulatory assignment. For example, if a firm sees that many variants of a car model fall in a very tight fuel economy standard, then it makes variant-level choices for that car model in the subsequent period. Such a firm may decide to eliminate all grades for the car model entirely, or change the characteristics of car grades, or re-design the platform, or offer a completely new model under a different name.

Third, the model presented in **Subsection 2.A.** tells us that attribute-basing leads to distortion in product choice and technical change, as long as the regulatory trade-offs between attributes differ from technically feasible trade-offs firms face. In other words, the distortion occurs *both* when the regulatory slope is higher *and* when it is lower than the TPF slope. This poses a challenge in identifying the regulatory impact because we do not observe the TPF in the absence of regulation, and hence, we cannot directly compare the regulatory slope with the TPF slope. In the case of automobiles, however, it is known to be extremely costly for the firm to decrease weight given the vehicle’s platform design. Hence, in the present context, we assume that the distortionary incentives are unidirectional. Hence, the hypothesis to be tested in our empirical context is, *The weight-based fuel economy regulation distorts technical change if the regulatory slope is higher than the slope of the (average) firm’s TPF.* This unidirectional nature helps us use policy-induced variations for identifying the distortionary impact of the regulation.

Under these assumptions, we should be able to identify the impact of the fuel economy

⁶Nowadays, it is very common for automakers to share technologies and platform designs across different models. Hence, technology capital does exist at a higher level than the model level. However, there is still likely to be a difference between the level of technology capital at the model level versus that at the firm or the shared-model level. That difference is all that is required for our empirical strategy.

regulation, in principle, by comparing differences in the observed TPFs across different vehicle models assigned to different weight segments. The challenge, of course, is how to control for other confounders that might have affected the TPFs. The theory presented in the previous section indeed motivates this strategy. We discuss our identification and estimation strategies in more detail in **Section 4**.

3. Background and Data

3.A. Regulatory Background

The Japanese fuel economy regulation is based on what is known as the *Top-runner* system. The system was first introduced under the 1999 Amendments to the Energy Conservation Act for all manufacturing products that consume energy in utilization. Under the *Top-runner* system, the government first classifies all vehicles according to their curb weights, and then chooses the highest observed energy efficiency level as the standard for that product category. This results in the fuel economy standards that are a step function of curb weights. The first weight-based fuel economy standards under this system were adopted in 2001 with a target year 2010. Since then, the standards were revised twice, in 2007 and 2013. Like the Corporate Average Fuel Economy (CAFE) standard in U.S., the Japanese fuel economy standards are enforced only at the firm level, based on the sales-weighted corporate average. **Figure 3** depicts the 2001 standards and the 2007 standards. The Ministry of Land, Infrastructure, Transport and Tourism (MLIT) adopted a new fuel economy rating method, known as *JC08 Mode*, for the new standards. Hence, the figure reports the new standards in both the old measure (known as *10.15 Mode*) and the new measure (*JC08 Mode*). The method of conversion between the two measures is described in **Subsection 3.B.** in more detail.

There are three important differences between the Japanese fuel economy standards and the U.S. CAFE. First, no credit trading is allowed, either across firms or segments. Hence, the marginal costs of compliance are not equalized across weight segments and firms. Second, the Japanese fuel economy regulation is somewhat close to voluntary regulation. Fines for non-compliance are only 1 million JPY (\approx \$10,000) per firm, much smaller than fines under the U.S. CAFE.⁷ Moreover, the Japanese standards are not enforced every year, and instead, firms are expected to meet the standards only by the (respective) target years. Despite the voluntary nature of the standards before the target year, Japanese firms take them seriously. All firms met the 2001 standards by 2005 well ahead of its target year 2010. Hence, the new

⁷Under the U.S. CAFE, fines are \$55 per vehicle. A recent study by the NHTSA shows the U.S. automobile industry has been paying roughly \$20 million annually since 2010.

standards were adopted in July 2007. The firms again met the 2007 standards by 2012 before its target year 2015. Hence, the Japanese government again adopted the latest standards in March 2013 with a target year 2020.⁸ **Figure A1 in the Appendix** demonstrates that at the beginning of the new standards, all domestic car makers were behind the required fuel economy standards, and hence, are likely to have made some efforts to meet the standards during the post-2007 period. Ito and Sallee (forthcoming) provide more direct evidence of the binding nature of the Japanese fuel economy standards.⁹ Third, the Japanese government also introduced a series of tax/subsidy incentives since 2009. Interestingly, these incentives were tied to the 2001 standards, rather than the 2007 standards, until 2012 [for details, refer to Konishi and Meng (2017)]. Hence, firms faced the same tax incentives within each old weight segment until 2012. For this reason, we constrain our main empirical analysis up to year 2012. This allows us to isolate the confounding impact of the tax/subsidy incentives offered discontinuously since 2009.

The new 2007 standards created an interesting regulatory setup, and thus, is a focus of our study. The government chose a smaller bin width to define each product category. As a result, each old weight segment was divided into two or more segments, resulting effectively in 16 weight bins in total under the new standards in contrast to 9 under the old standards. For some reason (not transparent in regulatory documents), the bin width differed substantially across weight bins. Furthermore, because the fuel economy performance of the top-runner relative to the peers in the same weight bin differed substantially across different weight bins, the required fuel economy improvement relative to the old standard also differed substantially across these bins. Consequently, there are bins that are relatively steeper in slope than others relative to the old standards ('slope' as in a decrease in the fuel economy standard per unit of weight increase [see formula (2) in **Subsection 3.D.**]). We expect that this variability in slope and stringency levels across weight bins distorts economic incentives for firms' product offerings.

To illustrate our point, let us take an old weight segment 970-1,265 kg and focus on three new weight bins on this segment as an example for our exposition. In **Figure 4**, both old and new fuel economy standards are drawn (red and blue lines, respectively). The first point to note is that under the Top-runner system, the government chooses the highest fuel economy rating that was achieved for each weight bin as the standard for that bin. This means that the

⁸In this paper, we refer to the old standards as the "2001 standards" and the new standards as the "2007 standards" both for clarity and for economizing space, although they are often referred to as the 2010 and the 2015 standards, respectively, in the Japanese regulatory context.

⁹In addition, a recent scandal on Mitsubishi Motors may suggest that the fuel economy regulation is indeed binding for some brands. In April 2016, MLIT found that Mitsubishi Motors manipulated fuel economy data for nearly 20 models over the last 10 years (Japan Times, Jun 17, 2016).

standards approximately trace out the technology frontier of ‘the most fuel efficient’ vehicles that were available as of 2007, which would presumably be substantially different from that of a ‘typical’ or average vehicle for each bin. Consider a line connecting the two endpoints A and D of this segment. For the moment, this line represents the technical frontier of a ‘typical’ or average firm. Then the slope of the regulatory standard is clearly steeper than this technical frontier for some bins, and flatter for other bins. In **Figure 4**, the lightest bin turns out to be the high-slope bin and the heaviest bin turns out to be the low-slope bin. Then by virtue of our **Proposition**, we should expect the firm to *increase* curb weight for vehicles that lie in the high-slope weight bin and to *decrease* curb weight for vehicles that lie in the low-slope bin, for then the required increase in fuel economy would be smaller. In reality, however, it is generally much easier and less costly to increase than to decrease curb weight given the platform design of a vehicle (Ito and Sallee, forthcoming). Provided that this is true for all firms, we expect that firms respond to the new standards by increasing curb weights and then investing in technical upgrade to improve fuel economy for vehicles in high-slope weight bins, while investing in technical upgrade without decreasing curb weight for vehicles in slow-slope weight bins. Consequently, vehicles that were assigned to the high-slope bins are likely to lie on a lower technical frontier than those that were assigned to the low-slope bins in the future model changes. That the standards are enforced on sales-weighted averages is likely to simply accelerate this incentive to increase curb weight for vehicles in the high-slope bins because it is easier for the firms to meet the overall standards if the firms have more car variants in low-slope weight bins.

Note that our argument does not quite depend on the assumption that the line connecting endpoints A and D of the old segment represents the technical frontier of a ‘typical’ firm. What matters for our empirical analysis is that different bins with different slopes relative to this line are likely to give different economic incentives for technical progress. For that, we simply assume that vehicle models that lie on the same old segment are likely similar to each other. Given this assumption, the likelihood in which vehicle models have flatter technical frontiers than this line is higher for the high-slope bins than for the low-slope bins. This creates differences in average incentives across weight bins.

3.B. Data

Our data come from Carsensor.com, one of the largest online car retailers in Japan. The compiled data set contains variant-level information on observable attributes of virtually all vehicles sold since 1991: e.g., model year/month, curb weight, displacement level, fuel economy rating, horsepower, list price, size, torque, transmission and other available options.

Importantly, because we have information on grade year/month at the variant level, we can identify the year in which each vehicle variant was first offered to the market. Our main analysis covers a subsample vehicles launched during the 2004-2012 excluding observations in 2007-2009 because the new standards are implemented in July 2007 and we anticipate that it takes at least a few years before the regulation influences firm’s technical capital. Hence, we use 2004-2006 as the pre-treatment control period and 2010-2012 as the treatment period. More detailed justifications for this choice follow below.¹⁰ In our robustness analysis, we also use observations from 2001 to 2003 and from 2013 to 2015.

We drop diesel, electric, and hybrid vehicles as well as commercial vehicles since they are not subject to the same fuel economy regulation as outlined in **Subsection 3.A**.¹¹ We also drop observations on imported brands because foreign manufacturers can always choose to sell a subset of their models to Japan, and thus, their TPFs are unlikely to fully respond to the incentives created through the Japanese regulations. We also exclude vehicles produced by Mitsubishi Motors because it might severely contaminate our results if included, since the recent scandal revealed that their reported fuel economy ratings during our study period do not meet the same regulatory guidelines as others (see **footnote 8**).

A complication arises in compiling fuel economy data. The Ministry of Land, Infrastructure, Transport and Tourism (MLIT) changed the method to measure fuel economy as an effort to align reported fuel economy with actual on-road fuel economy. As a result, all new vehicles offered after October 2010 must report fuel economy in a new measure, known as *JC08 Mode*, while all vehicles offered before October 2010 report in an old measure, known as *10.15 Mode*. These two measures are not directly comparable. Fortunately, however, the MLIT also mandated that all old vehicles must also record fuel economy in *JC08 Mode* if they are still sold in the market. Hence, the Japanese manufactures tend to report fuel economy in both measures in our study period. We fit a regression of 10.15-mode fuel economy on JC08-mode fuel economy on these observations, and then use the predicted fuel economy in case of vehicles missing fuel economy data in *10.15 Mode*.¹² From here on, all fuel economy data are reported in the *10.15 mode*.

3.C. Reporting versus Product Offerings

¹⁰The statistical significance and direction of the regulatory impact are largely intact, though the magnitude of the impact does change, if we also include observations from 2008 and 2009.

¹¹There is a separate weight-based fuel efficiency regulation on diesel cars. The sales of diesel cars accounts for a tiny portion of the overall sales in the Japanese market. Hence, to avoid unduly complications, we drop diesel cars from our analysis. Hybrid vehicles are subject to the same regulation, but their fuel economy ratings are well above the fuel economy standards, and therefore, weight category assignment should not influence their technical progress.

¹²The regression is surprisingly well fit with $R^2 \approx 0.99$.

We clarify important differences between our data and the data used in Ito and Sallee (forthcoming). Their data come from the list of new cars reported at the end of each fiscal year by the MLIT. The list contains all cars sold as ‘new cars’ as of the end of each fiscal year. This means that some car models are reported in multiple years in the MLIT data. For example, Toyota Vitz 2010-model, which was sold as a new car between December 2010 and April 2012, are reported twice in fiscal years 2010 and 2011. Our data do not suffer from this double counting because we have information on model years and we count each observation only once for the year it was first launched. Furthermore, the MLIT data are reported at the car configuration (or ‘*Katashiki*’) level, which is finer than the car model level, but is coarser than the car variant level like in our data. As a result, the MLIT data contain a smaller number of observations in each year than our Carsensor data (despite their possible double counting). For example, in 2010, approximately 1,200 car variants are available in our Carsensor data whereas only 700 car configurations are reported in the MLIT data. The most unappealing aspect of the MLIT data is that it reports only the *range* of curb weights for each car configuration (see **Table A1 in the Appendix** for the raw image of the MLIT table). Of 2,012 observations in the MLIT data between 2010 and 2012, only 25% are reported with exact weights. The remaining 75% of observations are reported only with ranges (i.e., minimum and maximum weights).

These important differences naturally lead to the question: Do the weight distributions differ between the two data sets? **Figure 5** demonstrates that they are indeed substantially different. The **top panel (a) of Figure 5** displays three vehicle weight distributions for all car configurations reported between 2010 and 2012 in the MLIT data: (i) observations reported with exact weights, (ii) minimum weights using observations reported with weight ranges, and (iii) maximum weights using observations reported with weight ranges. The figure confirms Ito and Sallee’s assertion — significant bunching occurs at the weight cutoffs for fuel standards. The figure, however, shows a few other points not discussed in their paper.

First, the incidence of bunching is primarily driven by the observations reported with ranges. Interestingly, the minimum weights are clustered at the right of the weight cutoffs while the maximum weights are clustered at the left of the weight cutoffs. That is, automakers report the maximum weights so as *not* to cross over to the *heavier* weight category while they report the minimum weights so as *not* to cross over to the *lighter* weight category. This mechanism is substantially different from that discussed in Ito and Sallee. This behavior occurs because the regulatory agency assigns the car models to the lightest weight bin when their weights range over two or more weight categories. Therefore, automakers have very strong incentives not to cross over to the lighter weight bins.

Second, the incidence of bunching mostly corresponds to the 2001 standards, not the 2007 standards. This can be most clearly seen in the weight cutoffs around 1,500 kg. The 2001 weight cutoff around this segment was 1,515 kg whereas the 2007 weight cutoff was 1,530 kg. The bunching is occurring at 1,520 kg, i.e., to the right of the 2001 standard’s cutoff and to the left of the 2007 standard’s cutoff. This behavior is consistent with the regulatory background discussed in **Subsection 3.A**. During the 2010-2012 period, the tax incentives were based on model’s fuel economy performance relative to the old 2001 standards, and the government used the MLIT list in determining the amount of tax incentives for each car model.

Lastly and most importantly, the incidence of bunching at weight cutoffs largely disappears once we use actual product offerings at the variant level. This is evident in the **top panel (b) of Figure 5**. Spikes in counts of car variants do occur, but none of these spikes corresponds to the weight cutoffs for either the 2001 or the 2007 standards. We believe this weight distribution is more consistent with findings in the empirical industrial organization literature. For automakers, how best to serve consumer demand and to strategically position and price their products against their market competitors in markets is of the first-order importance. It would probably not be ideal for automakers to bunch up so many of their vehicles at the weight cutoffs even when they can reduce costs of compliance by doing so. This is particularly so when they know that they don’t need to meet the standards until much later, year 2015.

The question arises then: What explains the behavior in the MLIT data? Our explanation is as follows. Automakers offer many different variants of the same car model/configuration at different weights within the weight range reported in the MLIT data. The automakers do not know how well the new model performs in the markets, and thus, how many variants of the model they wish to offer over the course of the model year, at the time of reporting the new model data to the MLIT. Hence, they would like to keep the weight range as large as possible while they would also like to avoid assignment of their models to the lighter weight bin. The best strategy then is to report the minimum possible weight of their car models at or above the lower weight cutoffs.

Table 1 reports weight distributions in the two data sets, and confirms these points more forcefully. The MLIT data set reports vehicle weights with range for 75% of the observations whereas our data set reports exact weights for all observations. The range in the MLIT data can be as large as 200 kg, averaging at around 35 kg. When we use observations reported without range, we see, in both data sets, that vehicles are roughly equally distributed to the right and to the left of the 2001 standards’ cutoffs, but reported more frequently to the *left* than to the *right* of the 2007 standards’ cutoffs. The latter contradicts Ito and

Sallee’s argument that firms increase weights to move to lower standards. More importantly, when we use observations reported with range (in the MLIT data), *minimum* weights are reported more frequently to the *right* of the 2001 standards’ cutoffs for the 2001 standards, yet *maximum* weights are reported more frequently to the *left* of the cutoffs. Interestingly, at the 2007 standards’ cutoffs, the frequencies stay the same between minimum and maximum.

We take these as suggesting that the MLIT data offers the evidence of bunching in ‘reporting’ to the MLIT rather than actual ‘product offerings’ in the market. Our empirical analysis delivers more convincing evidence on the existence of an incentive to increase weights in actual vehicle offerings. Unfortunately, however, the effect of this incentive is obscured by the other incentive to diversify product offerings, and hence, does not show up as vividly as we wish as bunching at weight cutoffs.

3.D. Constructing Treatment Variables

As discussed in **Subection 3.A.**, some weight segments under the 2007 standard are more stringent than others in terms of required improvements relative to the old standards. Roughly, this difference in stringency levels represents the difference in costs of compliance across weight segments. At the same time, some segments are narrower than others, resulting in the difference in the ‘slopes’ of weight segments: i.e., the required increase in fuel economy per unit of decrease in curb weight differs substantially across segments. Combined, these two variations are likely to create different marginal incentives across segments for firms’ compliance strategies.

Table 2 reports the number of vehicle variants as well as the mean and standard deviation of fuel economy ratings for each weight segment. Each row represents a new weight bin, which we define as the intersection of the old and the new weight segments. The solid lines represent weight segments under the 2001 standards, and the dashed lines represent those of the 2007 standards. For each bin, the table also reports the required fuel economy improvement relative to the old standard, and the required fuel economy improvement per unit of decrease in weight (i.e., slope) for each weight bin, which is calculated as follows:

$$\Delta = \left| \frac{h_{b+1} - h_b}{w_{b+1} - w_b} \right|, \quad (2)$$

where w_b and h_b are, respectively, the weight cutoff and the fuel economy standard for b th weight bin under the new 2007 standards. From here on, we call a weight bin that requires a large fuel economy improvement relative to the old standard as "high-(compliance) cost"

bin and a weight bin with a large slope as a "high-slope" bin. A detailed discussion on how to define "high" vs. "low" follows below.

The table indicates substantial variation in both compliance costs and slopes across weight bins. To get a sense of variation in compliance costs, we classify weight bins into quartiles with 1 denoting bins that fall in the lowest 25th percentile and 4 that fall in the highest 25th percentiles. We also classify weight bins into high- versus low-slope bins according to whether their slopes are steeper than the average slope of the joint segments connecting all bins within each old segment as illustrated in **Figure 4**. By this, we are assuming that vehicles within each old weight band faced roughly the same technical frontier and that a new segment steeper than this average slope give more incentives to manipulate on curb weight. The high slope, by definition, represents a larger decrease in the required fuel economy improvement per unit of weight increase, but firms would not be able to capture that gain unless the slope is steeper than the technical frontier they face. An additional benefit of defining this way is that it allows us to have both treatment and control bins within each old weight segment. We discuss this point in more detail in **Section 4**.

The table also shows substantial variation in the number of new offerings across weight bins. There is some indication that firms are avoiding new offerings in the high-cost/high-slope weight bins. This is really an analogue of the ‘bunching’ effect Ito and Sallee (forthcoming) point out.¹³ However, the tendency is not necessarily clear — there are high-cost/high-slope weight bins that received roughly the same number of new offerings between the pre-2007 and the post-2007 periods. This occurs presumably because firms may strategically offer models in the stringent weight segments as a way to avoid tough competition in less stringent segments. This is one reason why we think our reduced-form approach to identify only the TPFs is more viable than a structural approach. Another take-away message from the table is that there is no high-slope weight bin that falls in either the 1st quartile or the 3rd quartile of compliance costs. Ideally, we would like to compare outcomes of vehicles assigned to the high slope bin with those of the low slope, conditioning on the same compliance costs. Thus, for cleaner results, we drop vehicle models that fall in the 1st and the 3rd compliance cost quartiles during the pre-2007 period. This also comes with an added benefit of being able to classify weight bins into high cost (4th quartile) and low cost (2nd quartile) categories.

¹³In **Section 7**, we demonstrate that in our data, bunching does not occur at the weight cutoffs, and explain why that’s the case as well as why ours is likely to be a more accurate account of behavioral responses.

Given the above, we define two regulatory treatment variables H_j and T_j as follows:

$$H_j = \begin{cases} 1 & \text{if model } j \text{ (or its ancestor) was} \\ & \text{assigned to a high-compliance-cost segment in 2007,} \\ 0 & \text{otherwise.} \end{cases}$$

$$T_j = \begin{cases} 1 & \text{if model } j \text{ (or its ancestor) was} \\ & \text{assigned to a high-slope segment in 2007,} \\ 0 & \text{otherwise.} \end{cases}$$

Provided that there exists model-specific technology capital, these regulatory treatments should influence the level of technology capital s in the future period.

Since there are many variants of each vehicle model, all of the variants of a model may not necessarily fall in a single weight bin. We calculate the unweighted mean of curb weights of all variants for each vehicle model during the pre-2007 period, and then classify the vehicle model according to that mean.¹⁴ By this, we are implicitly assuming that firms' model-specific technical investment depends on their model-level average characteristics. We also experimented several model-level moments. Although some variations do occur, the results are largely intact. **Figure A2 in the Appendix** plots box diagrams describing the distribution of variant-level curb weights for all vehicle models that were assigned into the high-slope weight bins. The graph demonstrates that for virtually all models, the mean and the median values lie within a single weight bin.

Lastly, we need to trace out model histories for these vehicles, so that all vehicle variants offered during the post-2007 period can be associated with our treatment variables. For models that continue to exist, this is easy because they can be easily matched by model name. For discontinued models, we search through publicly available articles and company reports to see if there is any successor model for each retired model. **Table A1 in the Appendix** reports the summary of model history for car models assigned to the high-slope weight bins. Of the 30 models, 11 models did not introduce any new variants between 2010 and 2012, and thus, are classified as 'discontinued'. Of these 11 models, only 2 models had clear successor models. Others either had no clear successor model or were merged to another existing model.

3.E. Graphical Evidence

¹⁴Assignment based only on a single year, say, 2006 or 2007, is problematic in our setup because each vehicle observation is recorded with the year in which that vehicle was first offered. Because Japanese car models typically run on a 3-4 year cycle, including all the three-year observations likely cover all variants of models that are still produced as of 2006.

Before moving to our main analysis, we take a glance at graphical evidence. **Panel (a) of Figure 6** displays an unconditional scatter plot of logged fuel economy ratings against logged curb weights for vehicle grades introduced before the 2007 standards. As explained in the previous sections, the figure excludes imported cars, commercial vans and trucks, diesel, electric, and hybrid cars as well as vehicles that fall in the 1st and 3rd quartiles of compliance costs during the pre-2007 period. Variants of vehicle models that were assigned to the high-slope weight bins are marked with circle; those assigned to the low-slope bins are marked with \times . The figure indicates no sign of a significant difference in the technical trade-offs between fuel economy and weight prior to the new standards. This supports our identifying assumption that the regulatory assignment is not causally related to technical efficiency levels s prior to the standards — vehicle models assigned to the high-slope bins are not those with lower technical capital to start with. In fact, the trend line for those assigned to the low-slope bins lie slightly below, rather than above, those assigned to the high-slope bins.

Panel (b) of Figure 6 displays the same scatter plot for those introduced between 2010 and 2012 under the new standards. In this figure, variants of the successor models of those assigned to the high-slope bins are also marked with circle. We now see a substantial difference in technical trade-offs between those assigned to the high-slope bins versus the low-slope bins. Those assigned to the high-slope weight bands lie roughly the same technical frontier as that prior to the new standards *despite the fact that these are the new vehicles introduced after the new standards*. In contrast, new variants of those assigned to the low-slope bands lie on a *higher* and *steeper* technical frontier. Our empirical strategy is designed to answer two questions that arise naturally from this graph. The first question is, Can this difference be attributed to the regulatory assignment after controlling for all potential confounders? The second question is, Why do we observe differences not only in the level but also in the slope of technical frontiers between the two groups, which may seem somewhat inconsistent with our theory? Our analysis shows that these differences arise due to differences in regulatory assignment, and thus, are indeed consistent with our theory.

4. Identification and Estimation

4.A. Identification

The goal of our empirical study is to quantify the impact of model-level assignment to high-slope weight bins on the rate of technical progress for fuel economy improvements. Identifying this impact requires us to control for any confounders that might have systematically

influenced the technical progress of vehicles assigned to the high-slope weight segments. Even after controlling for observable vehicle characteristics such as weight, horsepower, and manufacturer fixed effects, some unobservable factors may still remain that are correlated with the regulatory assignment. For example, vehicle models assigned to the high slope segments may be those that had attained high fuel economy ratings (due to the *Top-runner* system), and therefore, might have experienced a systematically different technical progress anyway even in the absence of the assignment. The direction of the bias is unknown a priori since those vehicles may exhibit either faster or slower technical progress than their peers. Our strategy is to employ a DDD research design, exploiting temporal as well as cross-sectional variation in fuel economy standards across weight segments. Our three-fold control structure is as follows:

- (a) Cross-sectional between high-slope vs. low-slope segments (with low-slope segments as control)
- (b) Temporal over years (with years 2004-2006 as control)
- (c) Cross-sectional within groups (with low-compliance-cost segments as control)

By using temporal variation with years 2004-2006 as an additional control, we are able to control for any stationary differences between high-slope and low-slope groups as well as time-varying factors that are common to both groups. However, this difference-in-differences (DD) structure is not sufficient to control for unobservables that affect the two groups differently over time. To take care of this concern, we use another within-group variation. By construction, some of the high-slope segments also tend to be the high-compliance-cost segments: i.e., they also require substantial improvements in fuel economy relative to the old standards. Firms may face substantially higher incentives to manipulate vehicle characteristics in these high-cost segments. This reasoning suggests that we can potentially use cross-sectional variation in compliance costs within the same-slope segments as an additional control. That is, we compare the fuel economy ratings of the treatment vehicles assigned to the segments with high slopes *and* high compliance costs to the control vehicles assigned to the same high-slope segments with low compliance costs. The resulting DDD estimate is consistent under a weaker identifying assumption: i.e., *unobservables that affect the rate of technical progress differently between the high-slope and the low-slope segments do not systematically differ between the high-cost and low-cost segments.*

Besides the weaker condition for identification, this DDD structure comes with an additional benefit. That is, any pairwise DD estimate, in addition to the DDD estimate, is

also consistent if any pair of treatment/control groups satisfies the standard common-trend assumption. For example, if the contemporaneous shocks that affected the high-slope and low-cost groups have the same trend over time, then the DD estimate on a subsample consisting only of the high-cost segments is also consistent. Of course, this DDD estimate may not get at the pure impact of the high-slope assignment if the high-slope assignment also induced firms to manipulate vehicles in the low-cost segments. We offer evidence in **Section 5** that this is unlikely the case: in a DD regression on a subsample consisting only of the low-cost segments, the impact of the high-slope is positive and statistically significant. This means that those assigned to the high-slope segments were associated with a *faster*, not slower, fuel economy improvement in the low-cost segments. It follows then that the DD estimate would be biased toward zero, and controlling for this bias is important.

Our control structure also helps us purge out the confounding effect of tax incentives offered after 2009. As discussed in **Subsection 3.A.**, the government offered eco-car subsidy and tax credits based on fuel economy improvements relative to the old 2001 standards, despite that the new 2007 standards were already in effect. In **Subsection 3.C.**, we discuss how firms might have manipulated in reporting their car model weights to the government, and the reported weights clearly responded to the 2001 standards, not the 2007 standards, during the 2010-2012 period. It then follows that these tax incentives create the same incentive for fuel economy improvements for car models that lie within the same old weight category. Hence, any difference in behavioral response across new weight bins within each old weight category should be attributed to differences in fuel economy standards, not to the tax incentives.

To offer support for our identifying assumption, we plot **(a)** the means of fuel economy ratings by year and by treatment (i.e., high-slope vs. low-slope groups) and **(b)** the ratios of the mean fuel economy ratings of the high-cost groups relative to the low-cost groups by year by treatment. **Figure 7-(a)** demonstrates that both groups showed a steady increase in average fuel economy, yet the low-cost group increased fuel economy more sharply after 2009. The figure does seem to refute the concern that those assigned to the high-slope segments tend to be those that attained high rates of technical progress prior to the assignment. However, the temporal patterns between the two groups before 2007 do not appear quite identical, suggesting there might be other confounders that affect the two groups differently over time. In contrast, **Figure 7-(b)** demonstrates that the ratios of relative average fuel economy between the two groups have roughly identical temporal patterns between the high-slope and the low-slope groups. This boosts our confidence in our DDD estimates. The figures also point to another complication we might take into account. They show that changes in responses to the regulatory assignment are more discernible after 2009, rather

than immediately after the regulatory change in 2007. This may be attributed to the fact that it takes generally a few years for firms to introduce new vehicle variants to fully respond to the regulatory change or that firms' incentives to respond to the new standards became stronger after the Japanese government mandated reporting of fuel economy for all vehicles in *JC08 mode*, with which the new standards are enforced, in April 2009. For this reason, our main analysis uses only the observations between 2010 and 2012.

4.B. Raw DDD Estimate

Table 3 implements the unconditional DDD estimation of the effect of weight assignment on fuel economy ratings. Each cell reports the mean and standard errors of fuel economy in km/L for the indicated period-segment group as well as the number of observations (= vehicle variants) and the number of variants per model. The top panel A displays the statistics during the pre-2007 control period (2004-2006) whereas the bottom panel B reports the post-2007 treatment period (2010-2012). Each panel is further divided into the two layers of treatment by weight assignment. The left panel 1 (the right panel 2) concerns vehicle models assigned to high (low) compliance cost segments. The left panel 1 indicates that in the high compliance cost segments, firms improved fuel economy of vehicles assigned to the low slope segments by 5.1 km/L, but those assigned to the high slope segments only by 1.4 km/L. These result in the DD estimate of the effect of high-slope assignment by -3.7 km/L (statistically significant at 0.001). The right panel 2 repeats the same procedure, but for the low compliance segments. The DD estimate is -0.06 km/L, but is statistically highly insignificant. This is consistent with our argument that firms may not face much incentive to manipulate weight when it is relatively easy to comply with the new standards.

Although the results thus far are consistent with our theoretical prediction, the DD estimate would be biased if there were some unobservable confounders that affected the treatment and control segments differently over time. To take care of this concern, we obtain the DDD estimate by taking the difference between the two DD estimates in the left and right panels. The DDD estimate is -3.6 km/L and is again highly statistically significant (at 0.001). Both DD and DDD estimates are of the same order of magnitude. This boosts our confidence in our identification strategy. Moreover, because the DD estimate in the left panel compares the outcomes between the high-slope vs. the low-slope segments in the same high compliance cost segments, the results also support our claim that it is the slope, a measure of ease with which to manipulate the second attribute, not the high compliance costs per se, that induce firms to manipulate on the second attribute. The DDD estimate may be, however, imprecise (with a standard error of 0.7) since it fails to capture important

vehicle-level variation within each weight segment. To improve the precision of the estimate, we employ a regression framework below.

4.C. Estimation Strategies

As discussed in **Section 2**, a vehicle’s fuel economy is a function of its vehicle characteristics such as weight and horsepower as well as model-level or firm-level technical capital. Presumably, different vehicles produced by different manufacturers with different design features would respond differently to regulatory assignment. For example, vehicles with relatively superior fuel economy ratings, say, due to their design, size/weight, or technology features, would respond differently from those with relatively inferior fuel economy ratings.

To further exploit this important source of variation, we employ parametric DDD regression in a manner analogous to Gruber (1994), augmenting the Cobb-Dougllass specification of the technical possibility frontier à la Knittel (2011). Thus our regression equation takes the following form:

$$\begin{aligned} \ln f_{ijt} = & \alpha + \beta_1 H_j + \beta_2 R_t + \beta_3 T_j \\ & \dots + \beta_4 (H_j \times R_t) + \beta_5 (H_j \times T_j) + \beta_6 (R_t \times T_j) \\ & \dots + \beta_7 (H_j \times R_t \times T_j) + X'_{ijt} \gamma + \epsilon_{ijst}, \end{aligned} \tag{3}$$

where $\ln f_{ijt}$ is a logged fuel economy of vehicle i of model j assigned to segment s in 2007 in period t , H_j is a dummy, which equals 1 if j was assigned to a stringent (i.e., "high cost") segment, R_t indexes a regulatory period and equals 1 during the post-2007 period, T_j is our treatment variable and equals 1 if j was assigned to a high-slope segment, and X_{ijt} is a vector of observable vehicle characteristics including weight (w), horsepower (hp), size ($size$), torque (tq) (all in logged values), transmission type as well as brand and year dummies.

The standard argument shows that the OLS estimate of β_7 captures the causal impact of the regulatory assignment under the assumption that differences in unobservable time trends between the treatment and the control groups are the same across the high-cost and low-cost segments. One advantage of this DDD regression is that it nests the DD structure. That is, the DDD estimates are also consistent under the standard common-trend assumption that makes the DD estimates consistent. There are, however, two disadvantages with the DDD regression. First, it assumes that the steepness of the slopes does not affect the TPF for vehicles assigned to the low-cost segments: i.e., $\beta_7(H_j \times R_t \times T_j) = 0$ if $H_j = 0$. If the steepness of the slopes matters for the low-cost segments, the estimated impact is likely biased

downward. To gauge the extent of the bias, we also run a DD regression on a subsample where observations are restricted to those assigned to the low cost segments. Second, it assumes that the slopes of the technical frontier stays constant over time. As Knittel (2011) points out, the estimates of technical progress (and, hence, the DDD estimate) may be biased downward if the technical trade-offs between fuel economy and other attributes are not as large in later years. Of course, one could always allow slope coefficients to vary, say, by interacting them with our treatment variables. However, a priori, there is no obvious reason for one set of coefficients to vary while others do not.

An alternative strategy would be to use propensity score matching to control for the effects of these observable covariates. A disadvantage of the PSM estimator is that it requires a stronger identifying assumption than the DDD regression. That is, the DDD regression only requires that conditional on a set of covariates, differences in trend for unobservables between the treatment and the control groups stay the same between the high-cost and the low-cost segments while the PSM requires that the unobservables have zero means conditional on the set of covariates (i.e., conditional independence assumption). Because the PSM does not control for differences in unobservable time trend between the treatment and the control groups, the PSM estimates may be biased upward if the control groups exhibit a larger change in unobservable factors. Because the PSM is likely biased upward and the DDD is likely biased downward, the true impact of the regulation is likely to fall somewhere inbetween.¹⁵

5. DDD Regression Estimates

5.A. Main Results

Table 4 reports the results of four regression models. The first model estimates DD regressions on the pooled sample, with high-slope bins against low-slope bins as the primary treatment. The coefficients of this model would be biased downward if vehicles assigned to the low-cost weight bins also respond to the high slopes, even if the common-trend assumption between the treated and the control groups is satisfied. The second and the third models estimate the same regressions, but on subsamples consisting only of those of high-cost bins and low-cost bins, respectively. The last model estimates full DDD regressions on the pooled sample. Each of these models is estimated with or without brand dummies. All specifications include logged values of weight (w), horsepower (hp), size ($size$), torque (tq) as well as AT/CVT dummy as covariates.

¹⁵Our earlier attempt to employ PSM estimator confirms this prediction.

The DD estimate of the impact of the high slope on the pooled sample is negative and statistically significant. The magnitude of the estimate gets much larger when the same regression is run on a subsample consisting only of those assigned to the high-cost weight bins. In contrast, the DD estimate turns positive on a subsample consisting only of those assigned to the low-cost bins. Though the magnitude is quite small, the estimate is statistically significant. This implies, given our identifying assumption, that vehicles assigned to the high-slope bins might have been those with a higher rate of technical progress in fuel economy improvements. Hence, controlling for this confounding effect is important for an unbiased estimate of the regulatory impact, and gives support for our DDD design. The DDD estimate $\hat{\beta}_7$ is negative and highly statistically significant. Moreover, the magnitude is quite large: With the Cobb-Douglas specification, the estimate implies that the assignment to high-slope weight bins slowed down fuel economy improvements by roughly 20%. Because we control for all relevant covariates, this also implies that the TPF for those assigned to high-slope weight bins would have lied strictly above the observed TPF if they had been assigned to low-slope weight bins instead. Both the magnitude and statistical significance of these estimates are of the same order, with or without brand dummies.¹⁶

These results also explain why the observed TPFs seem flatter for those assigned to high-slope bins than those assigned to low-slope bins after than before the 2007 in **Figure 6-(b)**. As shown in **Table 2**, heavier weight bins tend to have less stringent standards (i.e., lower compliance costs). The assignment to high-slope bins in these heavier weight bins does not induce quantitatively large impacts on technical progress, whereas it has large negative impacts in lighter weight bins. Consequently, the observed TPF should look flatter. A scatter plot of predicted fuel economy against weight (in logged values) confirms this observation.

5.B. Robustness Check

We verify robustness of our regression results in two ways per conventional wisdom. First, we arbitrarily perturb our weight bin assignment and see if our results continue to

¹⁶We also estimated the same regressions using translog specification as well as year fixed effects. The results are virtually identical. The translog specification improves the fit by a small margin, but it also seems over-parameterized as in Knittel (2011). Since the Cobb-Douglas specification is already well fit, we do not report these alternative specification. Moreover, because parametric DD regression is known for its dependence to functional form assumptions (in particular with respect to changes in slope coefficients), we also tested propensity score matching estimators. We implemented the PSM estimation in the following manner. We take variants offered in 2004-2006 as a control group, and variants offered in 2010-2012 as a treatment groups. Then the treatment sample is further divided into two subgroups: variants of models assigned to high-slope bins versus low-slope bins. We then run the PSM estimator to get the average treatment effect on the treated (ATT) separately for each subgroup against the control group. Then compare the difference between the two ATT estimates to obtain the impact of regulatory assignment. The results were qualitatively similar to those reported here. These additional results are available upon request.

hold. Specifically, we shift weight cutoffs w_b in (2) by an arbitrary number k and run the same regressions as above. Because our estimate must capture the effect of the factual weight bin assignment, this fictitious assignment should qualitatively alter the results.¹⁷ **Top panel A of Table 5** below reports the DD and DDD estimates when $k = 25$ in kg and all the same covariates as in Table 3 (incl. brand dummies) are used.¹⁸ Both DD and DDD estimates on the pooled sample (columns 1 and 4) are still negative and statistically significant, yet their magnitudes become much smaller. Furthermore, the DD estimate on a subsample consisting only of high-cost bins turns positive and statistically highly insignificant. We take these results as support for our main results — the placebo experiments make the impact of the high-slope assignment smaller on the pooled sample and go away on the subsample in which the distortionary effect is expected to be strong. We suspect that the statistically significant estimates on the pooled sample are probably capturing the influence of some models that fall in both actual and fictitious high-slope weight bins.

Next, we perturb on temporal dimension, holding weight bin assignment. We use 2001-2003 as the control period and our original control period (2004-2006) as the fictitious treatment period. Because our estimate must capture the impact of regulatory change that took place in 2007, the DD or DDD estimate on this placebo treatment should be statistical insignificant. **Bottom panel B of Table 5** reports the results of the DD and DDD estimates on this placebo treatment. Unfortunately, both DD and DDD estimates are still negative and statistically significant. But their magnitudes are substantially smaller than those reported in **Table 4**. In particular, the DD estimate on a subsample consisting only of those assigned to high-cost bins is -0.03, compared to -0.15 in **Table 4**. Similarly, the DDD estimate on the pooled sample is -0.03 on this placebo treatment, compared to -0.23 on the actual treatment. These results suggest that vehicle models assigned to high-slope, high-cost weight bins may be those that showed slow technical progress anyway, but the bias that it causes to our DDD estimate is likely to be small.

5.C. Economic Mechanism

Our results so far confirm a statistically and qualitatively large impact of regulatory assignment to high-slope weight bins. A question remains as to exactly what economic

¹⁷It may not necessarily make the statistical significance of estimates completely disappear since in our regulatory setup, we expect influences of weight bin assignment in almost every weight bin.

¹⁸Note that we cannot choose k to be too small or too large. Because we average weights over all variants of each model, virtually all models would be assigned to the same weight bins if we choose k to be too small. In the meantime, choosing too large a number is problematic because it would end by shifting virtually all models to the next weight bins. The average bin size is roughly 75 kg. Hence, we end up choosing a number between 20 and 30.

mechanism caused that effect. The economic mechanism outlined in **Section 2** is that weight segments that have steeper slopes relative to the pre-existing TPFs would induce firm to increase curb weights for vehicle models and save costs of compliance to meet the fuel economy standards in these segments. If this is indeed the economic mechanism, we should also observe an increase in average curb weight for vehicle models assigned to the high-slope weight bins. Identifying this effect is, however, more challenging than identifying the effect on fuel economy for several reasons.

First, this logic suggests that vehicles in such weight bins should increase weights only up to the next weight cutoffs. This means that the anticipated weight increase should be bound, in principle, by bin size (measured as $|w_{b+1} - w_b|$ in kg). This is in contrast to fuel economy improvements, for which there is no apparent bound because how much to improve fuel economy given other product attributes (incl. weight) should only depend on the net marginal benefits of doing so. Hence, from the outset, the expected impact on curb weight may not be large enough compared to the variance of curb weight for each weight bin. This issue is further complicated by the fact that there is large variation in bin size. Larger (i.e., longer) weight bins may exhibit two counteractive effects. First, because firms have incentives to increase vehicle weight only to the next weight cutoffs, we might expect a larger weight increase in larger weight bins. However, larger weight bins also mean that it takes a more weight increase to cross the next weight cutoff. Given the design and size of a vehicle, it may be easy to increase weight by, say, 20 kg, but may be hard to increase weight by, say, 100 kg. A priori, there is no clear reason to expect which effect is stronger.

These reasonings suggest that for cleaner results, we might control for bin size. To do so, we first calculate bin sizes of all segments (excluding the lightest and the heaviest weight bins), and classify them into quartiles of bin sizes. By tabulating our main sample by these quartiles, we find that the 1st bin size quartile (i.e., the smallest bins) contains observations in all compliance cost \times slope subsamples. Hence, we run DD and DDD regressions of logged curb weight on the same set of covariates as in **Table 4** (excluding logged weight, of course). **Table 6** reports the results of these regressions. The DD estimate on the pooled sample (i.e., consisting of both high-compliance-cost and low-compliance-cost groups) is positive and statistically highly significant. The DD estimate is also positive and statistically highly significant on a separate subsample, either the high-cost group or the low-cost group. The magnitude of the estimate is roughly of the same size between the two samples. Consequently, the DDD estimate is statistically insignificant. These results seem to indicate that in both high-cost and low-cost segments, high-slope weight bins do create incentives to increase curb weight, and, hence, the results increase credence in our main results.

Unfortunately, the fact that the regulatory impact also exists in the low-cost segments

violates the condition for the unbiasedness of the DDD estimates — the high slope does not create incentives in the low-cost segments. Note, however, that this result does not necessarily nullify the usefulness of our three-hold control structures. To see this, recall our result in **Subsection 5.A.** that assignment to the high-slope in the low-cost segments was associated with a *higher* rate of fuel economy improvements. Combined with the result that the high slope also creates incentives to increase vehicle weights (and slow down technical progress) in these low-cost segments, the counterfactual fuel economy improvement in the absence of high-slope assignment should be even greater there. But if the same trend is occurring in the high-cost segments, then the DD estimate would be biased there too, and our DDD estimate should still be unbiased in that case. On the other hand, if the same trend is not occurring in the high-cost segments, then our DDD estimate is biased toward zero, and the true impact could be larger or smaller.

6. Welfare Implications

In this section, we discuss the welfare implications of our empirical findings, and substantiate the important differences from those of other studies that investigate the distortionary impacts of regulatory ‘loopholes’ in environmental regulation [Anderson and Sallee (2011), Sallee and Slemrod (2012), and Ito and Sallee (forthcoming)]. Regulatory loopholes distort firm’s or consumer’s choice, and thereby, result in welfare losses relative to the counterfactual in which there are no such loopholes. Without accounting for effects on technical change, however, the magnitude of the losses may be substantially understated.¹⁹ For example, Anderson and Sallee (2011) write, in their study on the flexible-fuel credits under the CAFE regulation, "the flexible-fuel loophole may actually increase welfare by allowing firms to relax an inefficient (fuel-economy standards) constraint (p. 106, parenthesis added). Ito and Sallee (forthcoming) also demonstrate that attribute-basing in the fuel economy regulation is welfare-increasing relative to no attribute-basing without an efficient credit trading mechanism. All these studies, however, assume the distortion shows up in firm’s or consumer’s choice *conditional on firm’s technology choice*, assuming away the distortion in technical change. For example, Sallee and Slemrod (2012) examine the distortionary impact of the notched schedule of the U.S. Gas Guzzler Tax, and estimate the welfare effects of marginally adjusting fuel economy ratings *given the first stage choices* "regarding engine size, body style and vehicle features that cannot be changed quickly and have large impacts on fuel economy" (p. 991). The same principle is also applied in Ito and Sallee (forthcoming), who estimate

¹⁹This may be the reason why regulators opt for such loopholes in the first place — regulators may be aware of the suboptimal nature of the regulation, but still use them in light of other merits.

the discrete choice model of firms’ new product offerings without accounting for consumer demand. In their model, firms simply weigh the distance in the attribute space between the existing offerings and the new offerings against the level of (one-time) tax incentives associated with that distance. By construction, therefore, these welfare estimates do not account for consumer surplus resulting from the change in technical change. Furthermore, the recent IO literature substantiates the importance of accounting for the gain in consumer surplus from new or improved products [e.g., Nevo (2000), Petrin (2002), Goolsbee and Petrin (2004)]. Failure to account for consumer loss from the distortion in technical change, therefore, is likely to understate the welfare effects of regulatory loopholes.

To demonstrate our point, we compute a ‘wild’ estimate of the welfare loss from the distortion in technical change. In doing so, we construct a counterfactual in which the fuel economy ratings of vehicles assigned to the high-slope segments would improve the same way as those assigned to the low-slope segments, yet all other vehicle attributes would stay the same as observed. This allows us to isolate the impact of the change in TPF from the change in other product attributes. We call it a ‘wild’ estimate because we do not incorporate strategic responses by firms in pricing or product choice under the counterfactual scenario, despite that firms would, in general, adjust all attributes of vehicles fully mindful of their competitors when they can offer vehicles on a higher TPF. A structural approach would give us a more realistic estimate, but we defer it to a future research. The approach we take here, however, is still consistent with the convention in applied welfare economics and the non-market valuation literature in environmental economics.

Specifically, we proceed as follows. We have annual vehicle sales data for our policy period (2010-2012). However, the data are only reported at the model level. Hence, we transform the grade-level attributes to the model-level attributes by calculating model-level averages. We then borrow the estimates of mean marginal utility parameters with respect to income λ and kilometer per yen θ (i.e., fuel economy divided by gasoline price p) from the random-coefficient logit model of consumer demand that is estimated for the same study period in Konishi and Zhao (2017). These estimates can give us an estimate of consumer’s (marginal) willingness to pay for fuel economy improvements, $-\theta/\lambda p$. Assuming that prices and other attributes stay the same, we can also assume that consumer demand for each car model stays the same. We can then approximate the change in consumer welfare purely due to the change in fuel economy technology by multiplying the model-level sales with the marginal WTP for the counterfactual fuel economy improvements and summing them over all affected vehicles. Our estimate comes at an annual welfare loss of roughly \$1.8 million or \$23,183 per vehicle model (**with standard errors to be calculated**). To place this number in context, Ito and Sallee estimate that the Japanese weight-based regulation has a

welfare *gain* of \$1,766-\$2,133 per vehicle model relative to no attribute-basing (but welfare loss of \$1,226-\$1,788 relative to an efficient credit trading). Thus, a failure to account for technical change (and associated welfare loss in consumer surplus) may lead to misguided policy advice.

Figure 8 visualizes what is driving our welfare estimate. **Top panel (a) of Figure 8** displays the sales distribution in the attribute space of vehicles introduced during the 2010-2012 period. Each circle represents a new vehicle model, and the size of the circle represents the annual sales of that model for only the year in which it is introduced. The figure shows how disperse the sales are over the attribute space. **Bottom panel (b)** plots the counterfactual in which the fuel economy ratings of vehicles assigned to the high-slope segments improve the same way as those assigned to the low-slope segments, yet all other vehicle attributes and sales amounts stay the same as observed. The figure also display the same data as in **Panel (b)** by x's without bubbles representing sales. As we assume a constant marginal WTP and ignore strategic interactions or substitution between vehicles, the welfare effects are simply larger for vehicle models with larger sales.

7. Concluding Remarks

We examine the distortionary impact of the Japanese weight-based fuel economy regulation on technical progress. We first set up a simple model of firm's choice over technology upgrade and product attributes to obtain a clear-cut economic prediction: An attribute-based regulation distorts technical change when it creates trade-offs between the targeted and secondary attributes that differ from technically feasible trade-offs. We use the variant-level vehicle characteristics data for new vehicle models launched between 2004 and 2012 in Japan to estimate the distortionary impact on firm's technology possibility frontier (TPF). To control for confounders, we employ a difference-in-difference-in-differences strategy, exploiting the quasi-experimental variations created due to changes in weight segmentation under the 2007 fuel economy standards. Our results indicate the stark impact of the regulation: Assignment to high-slope weight bins slowed down the rate of fuel economy improvements by roughly 20 percentage points. Hence, we conclude that the attribute-based regulation has significantly distorted technological change in the Japanese automobile industry.

The findings of the paper deliver four important messages. First, it is not just the notched schedule but rather the slope of the attribute-based regulation that induces distortion in product offerings. Second, the distortion in product offerings do translate into distortion in technical progress. Third, bunching behavior found in Ito and Sallee (forthcoming) is most likely the evidence of manipulation in reporting to the government rather than distortion

in actual product offerings in the market. Our paper, however, gets at the latter. Lastly, the paper points to an important policy advice: To remove such a bias, the regulator needs to make the slope of the attribute-based regulation as close as that of firm's TPF *both* when credit trading is in place *and* when it is absent. Credit trading is known to equalize the marginal cost of compliance. A conventional wisdom is that a flat standard would be innocuous in the presence of credit trading (Ito and Sallee, forthcoming). However, a flat standard would be too flat compared to the firm's TPF. It might still distort firm's incentives by offering low-cost compliance strategies (lower than buying credits). Credit trading lessens, but does not necessarily eliminate, the distortionary incentives. This last point is also important for other types of attribute-based regulation (e.g., product labels) as well as other institutional arrangements that involve multiple attributes in evaluation such as auctions, hiring, and other business contracts.

References

- [1] Acemoglu, D. (2002). Directed Technical Change. *Review of Economic Studies* 69(4): 781–809.
- [2] Acemoglu, D., Aghion, P., Bursztyn, L., and Hemous, D. (2012). The Environment and Directed Technical Change. *American Economic Review* 102 (1):131–166.
- [3] Acemoglu, D. and J. Linn (2004). Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry. *The Quarterly Journal of Economics* 119 (3): 1049-1090.
- [4] Aghion, P., A. Dechezleprêtre, D. Hemous, R. Martin, and J. V. Reenen. (2016). Carbon Taxes, Path Dependency and Directed Technical Change: Evidence from the Auto Industry. *Journal of Political Economy* 124(1): 1-51.
- [5] Anderson, S. T., and J. 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.
- [6] Berry, Steven, James Levinsohn, and Ariel Pakes. 1995 Automobile Prices in Market Equilibrium. *Econometrica* 63 (4): 841-890.
- [7] Calel, Raphael and Antoine Dechezleprêtre. (2016). Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market. *The Review of Economics and Statistics* 98(1): 173–191.

- [8] Copeland, Brian R. and M. Scott Taylor. (1994). North-South Trade and the Environment. *The Quarterly Journal of Economics* 109 (3): 755-787.
- [9] Goolsbee, A. and A. Petrin. (2004). The Consumer Gains from Direct Broadcast Satellites and the Competition with Cable TV. 72 (2): 351-381.
- [10] Gruber, J. (1994) The Incidence of Mandated Maternity Benefits. *American Economic Review* 84(3): 622-641.
- [11] Fan, Y. (2013). Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market. *American Economic Review* 103 (5): 1598-1628.
- [12] Hitsch, G. J. (2006). An Empirical Model of Optimal Dynamic Product Launch and Exit Under Demand Uncertainty. *Marketing Science* 25 (1): 25–50.
- [13] Ito, K. and J. Sallee. *Forthcoming*. The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel-Economy Standards. *The Review of Economics and Statistics*.
- [14] Knittel, C. R. (2011). Automobiles on Steroids: Product Attribute Trade-Offs and Technological Progress in the Automobile Sector. *American Economic Review* 101 (7): 3368-3399.
- [15] Konishi, Yoshifumi and Meng Zhao. (2017). Can Green Car Taxes Restore Efficiency? Evidence from the Japanese New Car Market. *Journal of the Association of Environmental and Resource Economists* 4(1): 51-87.
- [16] Nevo, A. (2000). Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry. *The RAND Journal of Economics*. 31 (3): 395-421.
- [17] Newell, R. G., Jaffe, A. B., and Stavins, R. N. (1999). The Induced Innovation Hypothesis and Energy-Saving Technological Change. *The Quarterly Journal of Economics* 114(3): 941–975.
- [18] Petrin, A. (2002). Quantifying the Benefits of New Products: The Case of the Minivan. 110 (4): 705-729
- [19] Popp, D. (2002). Induced Innovation and Energy Prices. *The American Economic Review* 92 (1): 160–180.
- [20] Porter, M. (1991). America’s Green Strategy. *Scientific American* 264(4), 168.

- [21] Sallee, J. M., and J. Slemrod. (2012) Car Notches: Strategic Automaker Responses to Fuel Economy Policy. *Journal of Public Economics* 96(11–12): 981-999,
- [22] Seim, K. (2006). An Empirical Model of Firm Entry with Endogenous Product-type Choices. *RAND Journal of Economics* 37 (3): 619-640.
- [23] Wollmann, T. *Forthcoming*. Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicle. *American Economic Review*.

Appendix. Proof of Proposition

Let other firms' product portfolios Ω be given. Then the Lagrangian of the firm's second-stage optimization program under no regulation is:

$$\mathcal{L}^N = \pi(f, w; \Omega) - c(s) + \lambda[T(w, s) - f],$$

where λ is the shadow value of the technology constraint. The first-order condition can be rearranged to yield an optimality condition:

$$\frac{\partial \pi}{\partial w} \bigg/ \frac{\partial \pi}{\partial f} = \rho.$$

Given (A1) and (A2), this optimality condition is necessary and sufficient. Under (A1), the technology constraint is binding: $f = T(w; s)$. Hence, the optimal fuel economy f^N is uniquely pinned down by $f^N = T(w^N; s)$ once w^N is pinned down. Along with the tangency condition above, this gives us a unique solution to the optimization program. Let $(f^N(s), w^N(s))$ denote the optimal solution given s .

On the other hand, under the attribute-based regulation, the Lagrangian can be written as

$$\mathcal{L}^R = \pi(f, w; \Omega) - c(s) + \lambda[T(w, s) - f] + \mu[f - R(w)],$$

where μ is the shadow value of the regulatory constraint.

Combining (A1) and (A4), we can write

$$\sigma = \rho + \alpha, \quad \text{for some } \alpha \in \mathbb{R}, \tag{4}$$

where $\sigma = -dR/dw$ is the slope of the regulatory constraint and $\rho = -\partial T/\partial w$ is the slope of the TPF. Using (4), the first-order condition of the Lagrangian under the regulation can then be rearranged to yield:

$$\frac{\partial \pi}{\partial w} \bigg/ \frac{\partial \pi}{\partial f} + \frac{\mu \alpha}{\partial \pi / \partial f} = \rho.$$

Again, let $w^R(s)$ and $f^R(s) = T(w^R(s); s)$ denote the optimal solution given s under the regulation. By (A1), the second term of the LHS is strictly positive as long as the regulation is binding. This means that the optimal attributes occur at the tangency between the iso-profit curve and a *flatter* TPF curve, instead of the true TPF. It follows then that the optimal attributes $(f^R(s), w^R(s))$ under the regulation lie to the right of the optimal attributes under no regulation $(f^N(s), w^N(s))$ if $\alpha > 0$., and to the left if $\alpha < 0$.

Now, let us consider the first-stage decision on technology capital s . We can solve for the optimum by maximizing the following objective function:

$$\pi(T(w^r(s); s), w^r(s); \Omega) - c(s),$$

taking the second-stage solution $w^N(s)$ and $w^R(s)$ as given, for $r = N$ under no regulation and $r = R$ under the regulation. Then under no regulation, the optimality condition is given by:


$$\left. \frac{\partial \pi}{\partial f} \frac{\partial T}{\partial s} \right|_{w^N(s)} = \frac{dc}{ds},$$

whereas that under the regulation is:

$$\left. \frac{\partial \pi}{\partial f} \frac{\partial T}{\partial s} \right|_{w^R(s)} - \mu \alpha \frac{dw^R(s)}{ds} = \frac{dc}{ds}.$$

Given (A3), these conditions are necessary and sufficient. The second term of the LHS is positive as $dw^R(s)/ds \geq 0$ for $\alpha > 0$. This means that the firm under the regulation values the marginal increase in profits less than under no regulation. Hence, the firm invests less in technology capital under the regulation. ■

Figure 1. Changes in Technology Trade-offs for Toyota's Passenger Cars between 1991 and 2015



Note: The figure excludes commercial vans and trucks, imported brands, diesel, hybrid, and electric cars.

Figure 2. Impact of Attribute-based Regulation on TPF




Figure 3. The Old and New Fuel Economy Standards

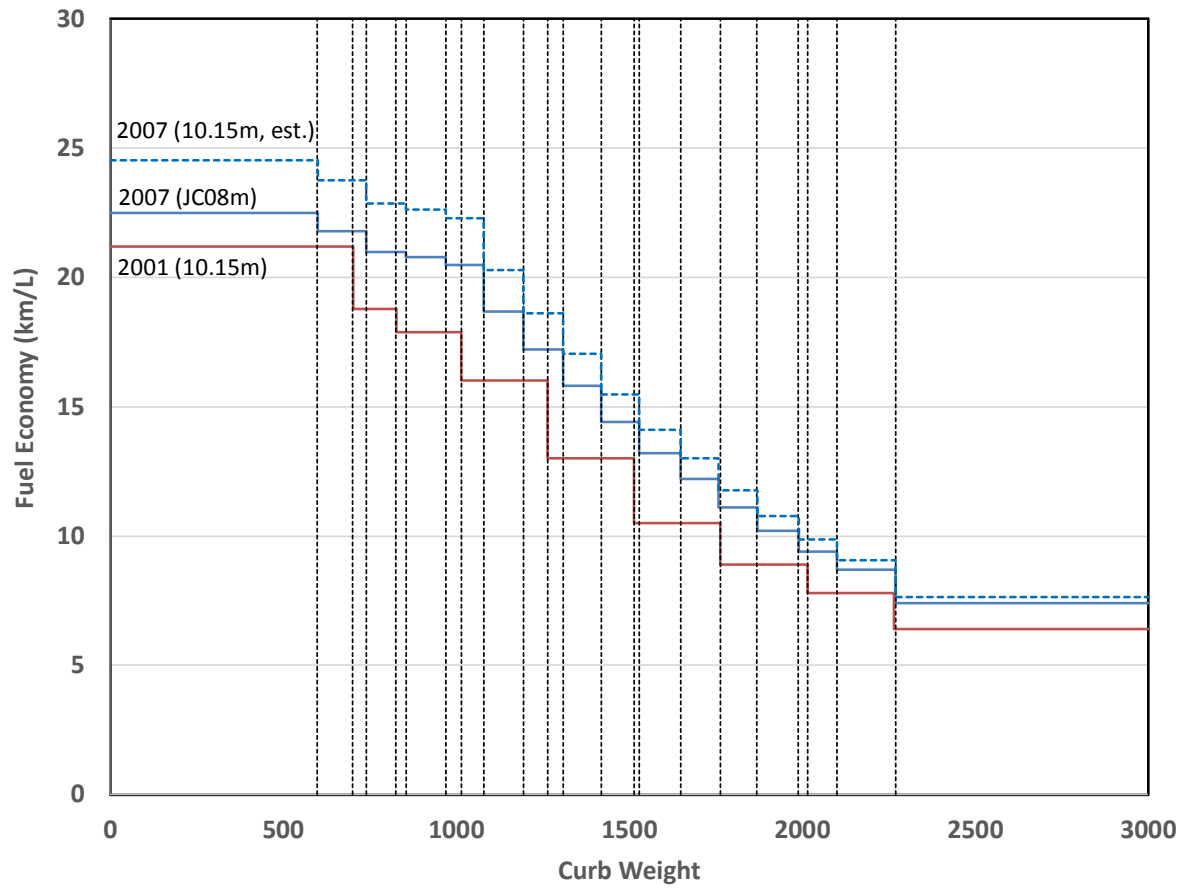


Figure 4. Variation in Regulatory Assignments: An Illustration

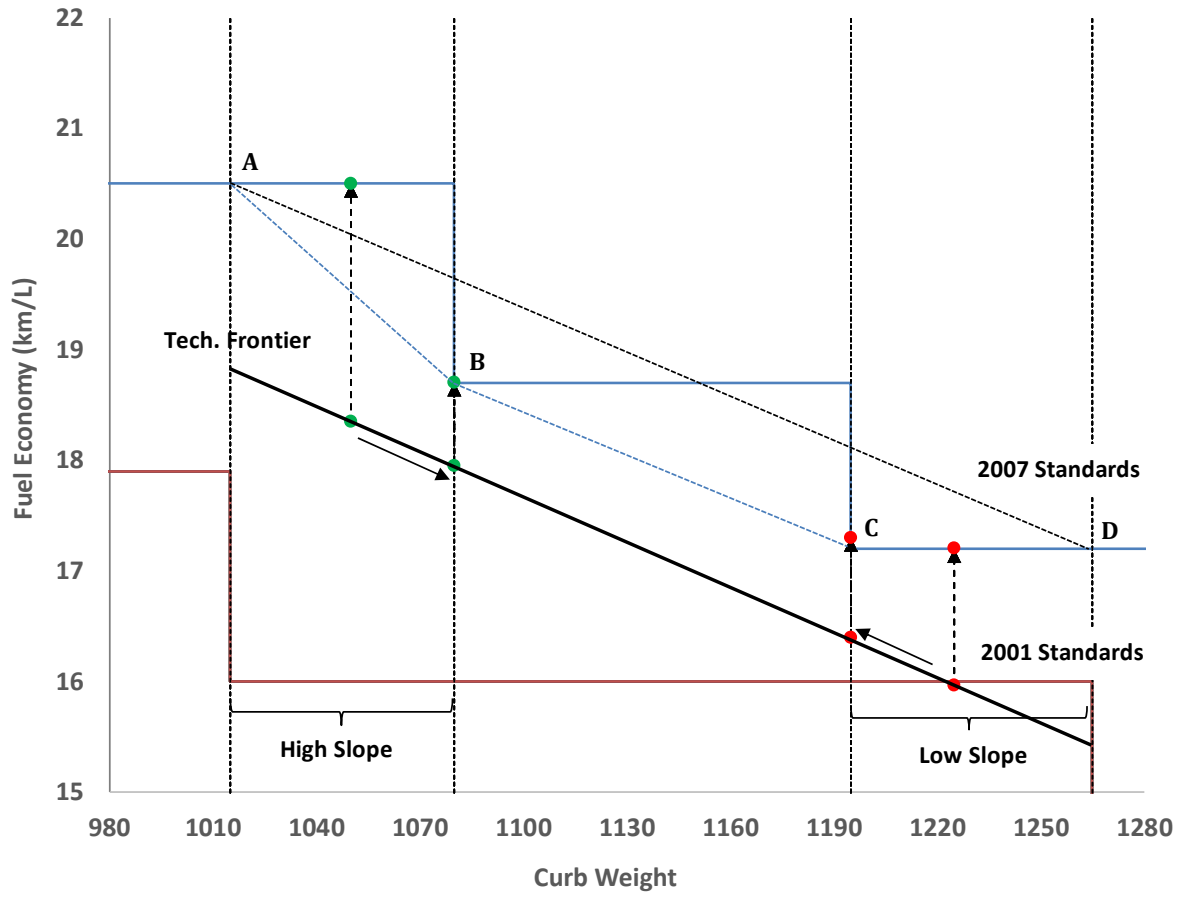
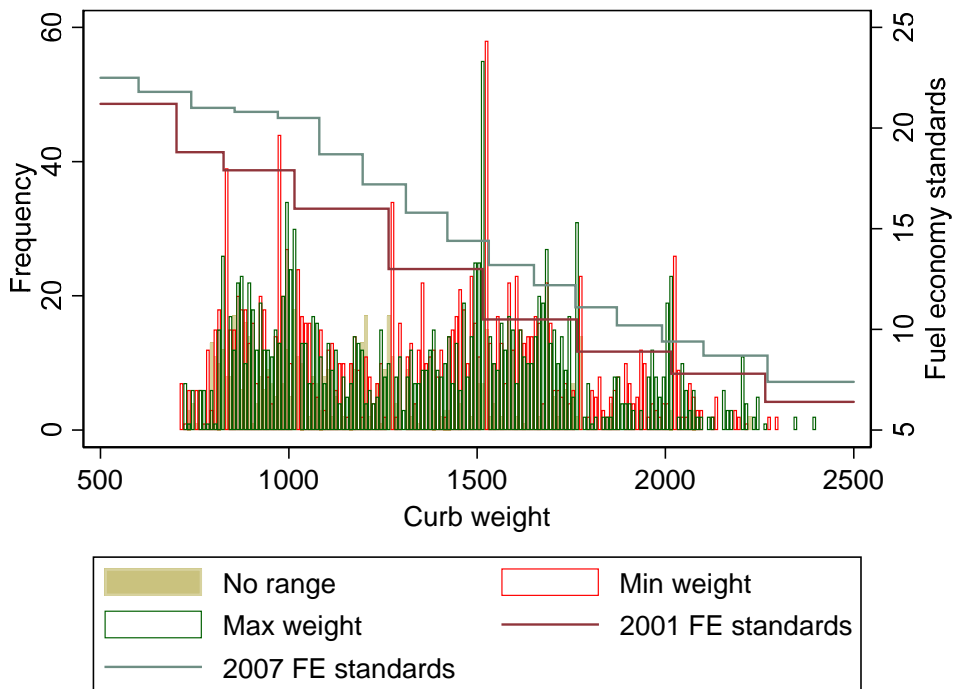


Figure 5. Vehicle Weight Distributions, Years 2010-2012

(a) Model/Configuration-level Data reported to MLIT



(b) Variant/Grade-level Data using Carsensor Catalog


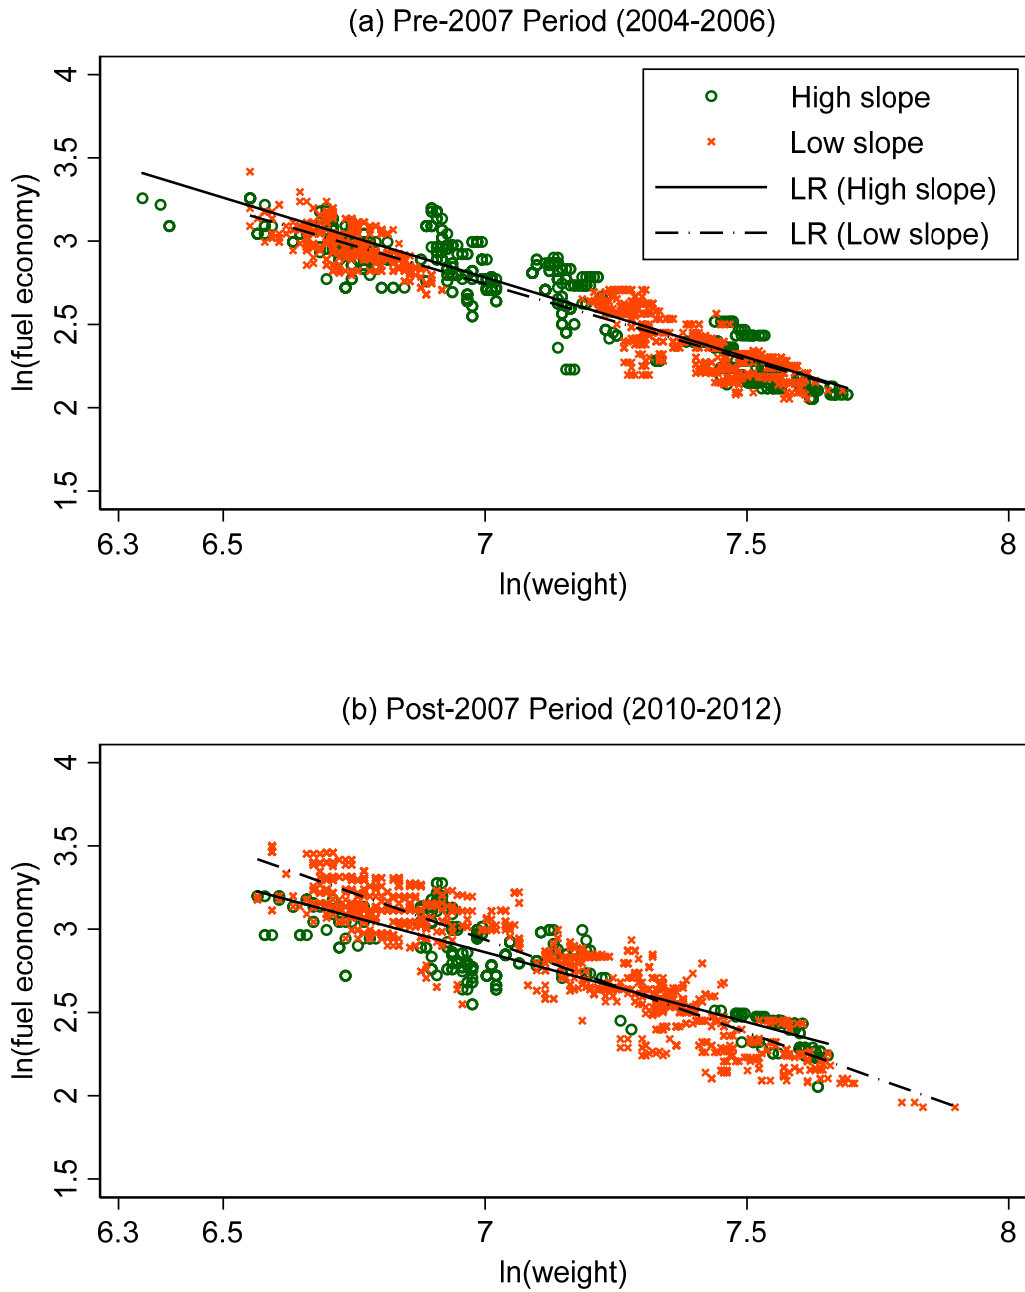

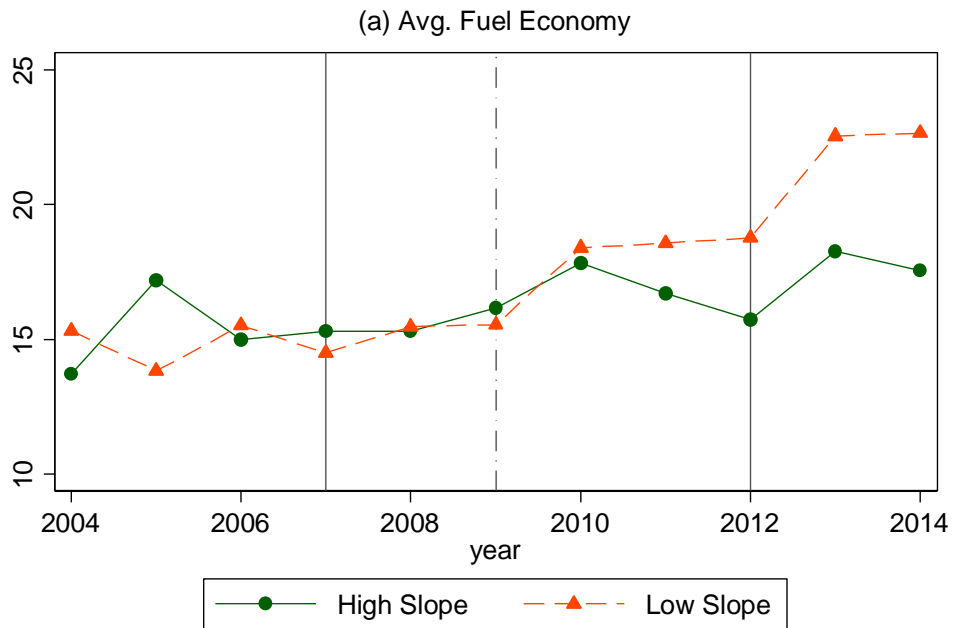


Figure 6. Technology Possibility Frontiers Before and After the New Standards



Note: The figure excludes commercial vans and trucks, imported brands, diesel, hybrid, and electric cars as well as observations that fall in weight segments with the first and the third quartiles of compliance costs during the pre-2007 period.

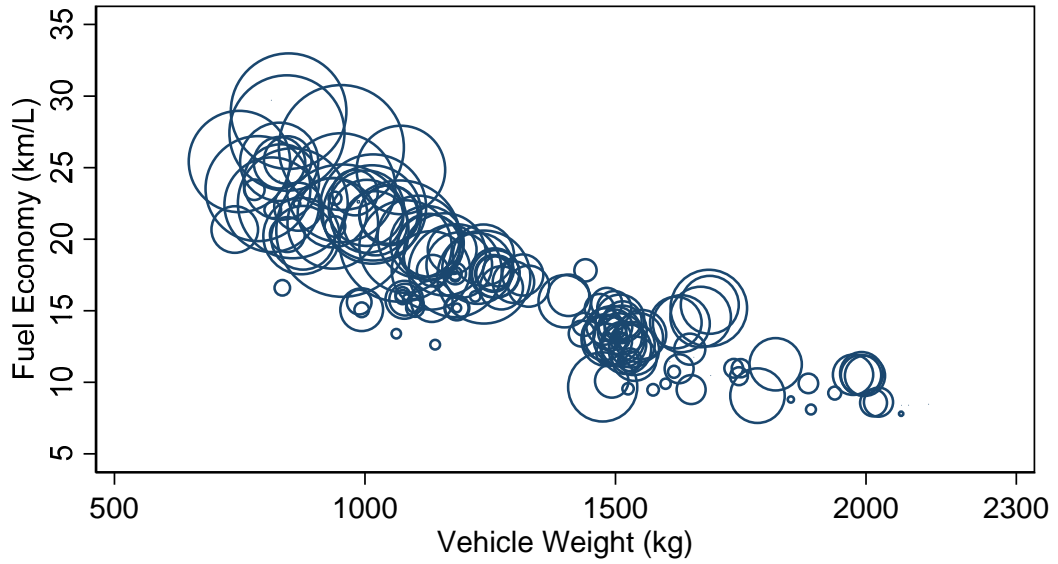
Figure 7. Trends in Average Fuel Economy between and within Groups



Note: Panel (a) plots average fuel economy ratings for the high-slope (treatment) and the low-slope (control) groups. Panel (b) plots the ratio of the average fuel economy of the high-cost group to that of the low-cost for the treatment and the control groups.

Figure 8. Model-level Sales and Welfare Impact of DTC in the Attribute Space

(a) Model-level Sales in the Attributes Space



(b) Visualizing Impact of DTC in the Attribute Space

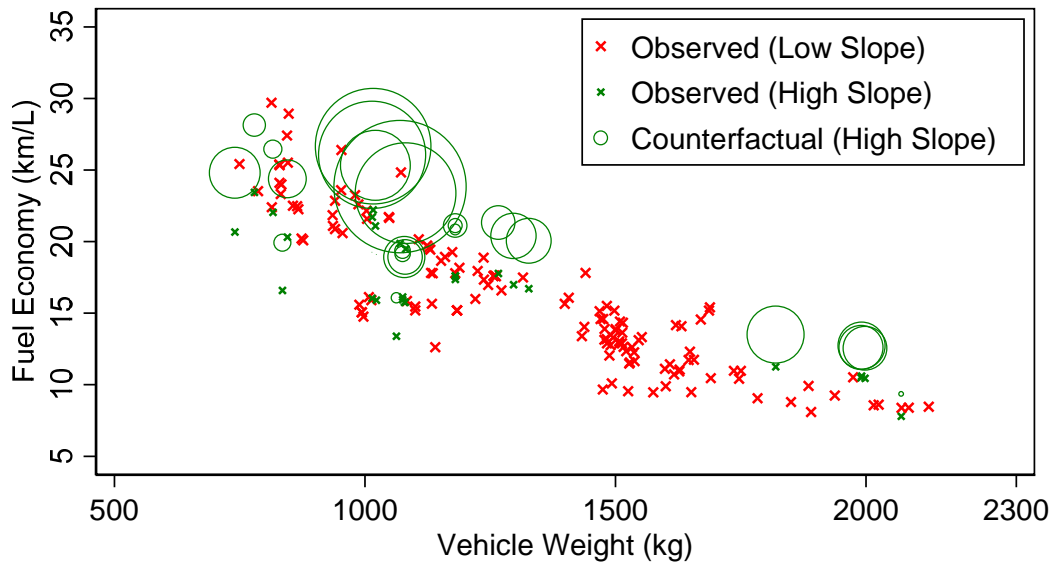


Table 1. Distribution of Vehicle Weight in MLIT Data vs. Catalog Data
2010 - 2012

	MLIT Data		Our Data	
	Obs.	Percent	Obs.	Percent
All	2,011		4,303	
Range = 0	508	(0.25)	4,303	(1.00)
Range > 0	1,503	(0.75)	0	(0.00)
Avg. Range (in kg)		35.5		
Min. Range (in kg)		10		
Max. Range (in kg)		200		
Of those reported with range = 0				
<i>Weight at the cutoffs of 2001 standards</i>				
To the right	59	(0.04)	214	(0.05)
To the left	59	(0.04)	192	(0.04)
<i>Weight at the cutoffs of 2007 standards</i>				
To the right	55	(0.04)	439	(0.10)
To the left	159	(0.11)	1,033	(0.24)
Of those reported with range > 0				
<i>Minimum weight at the cutoffs of 2001 standards</i>				
To the right	272	(0.18)		
To the left	59	(0.04)		
<i>Maximum weight at the cutoffs of 2001 standards</i>				
To the right	59	(0.04)		
To the left	235	(0.16)		
<i>Minimum weight at the cutoffs of 2007 standards</i>				
To the right	196	(0.13)		
To the left	555	(0.37)		
<i>Maximum weight at the cutoffs of 2007 standards</i>				
To the right	182	(0.12)		
To the left	560	(0.37)		

Table 2. Fuel Economy Ratings
by Weight Band under the New 2007 Standards

Weight Segments	Pre-2007 Period (2004-06)										Post-2007 Period (2010-12)		
	H22 (10.15M)	H27 (10.08M)	Chg. in FE	Quartile of ΔFE	Slope	High Slope	N	Fuel Economy (km/L)		N	Fuel Economy (km/L)		
								Mean	S.D.		Mean	S.D.	
1 0 – 600	21.2	22.5	1.3	2	0.001	1	5	23.40	1.95	0	--	--	
2 600 – 702	21.2	21.8	0.6	1	0.000	0	8	25.88	2.34	0	--	--	
3 702 – 740	18.8	21.8	3.0	4	0.021	1	42	22.12	1.67	43	27.82	4.74	
4 740 – 827	18.8	21	2.2	2	0.000	0	255	20.75	1.94	146	26.12	3.36	
5 827 – 856	17.9	21	3.1	4	0.007	1	150	19.29	1.82	108	23.66	3.41	
6 856 – 970	17.9	20.8	2.9	4	0.003	0	438	18.43	1.70	435	22.80	2.87	
7 970 – 1015	17.9	20.5	2.6	3	0.000	0	130	17.88	3.23	160	20.74	3.49	
8 1015 – 1080	16	20.5	4.5	4	0.028	1	202	17.20	1.70	202	19.37	3.15	
9 1080 – 1195	16	18.7	2.7	3	0.013	0	488	16.35	1.65	318	18.72	2.56	
10 1195 – 1265	16	17.2	1.2	1	0.000	0	310	15.83	1.58	161	17.49	1.99	
11 1265 – 1310	13	17.2	4.2	4	0.031	1	131	13.50	1.82	91	16.20	1.30	
12 1310 – 1420	13	15.8	2.8	3	0.013	1	301	12.97	1.78	162	15.26	1.83	
13 1420 – 1515	13	14.4	1.4	2	0.000	0	339	12.49	1.48	197	13.76	2.29	
14 1515 – 1530	10.5	14.4	3.9	4	0.080	1	75	11.47	1.15	57	12.91	1.06	
15 1530 – 1650	10.5	13.2	2.7	3	0.008	0	419	11.29	1.08	295	12.66	1.58	
16 1650 – 1760	10.5	12.2	1.7	2	0.010	0	281	10.32	1.20	174	12.24	2.34	
17 1760 – 1765	10.5	11.1	0.6	1	0.000	0	0	--	--	0	--	--	
18 1765 – 1870	8.9	11.1	2.2	2	0.009	1	193	9.50	0.87	48	10.39	1.34	
19 1870 – 1990	8.9	10.2	1.3	2	0.007	0	141	9.16	0.59	113	10.50	1.19	
20 1990 – 2015	8.9	9.4	0.5	1	0.000	0	16	8.57	0.26	28	10.22	0.85	
21 2015 – 2100	7.8	9.4	1.6	2	0.008	1	32	8.24	0.40	59	9.04	0.55	
22 2100 – 2265	7.8	8.7	0.9	1	0.000	0	14	8.10	0.15	14	8.45	0.61	
23 2265 – 2270	6.4	8.7	2.3	3	0.260	1	0	--	--	0	--	--	
24 2270 – 3500	6.4	7.4	1.0	1	0.000	0	11	6.41	0.30	4	7.00	0.12	

Table 4. Full DDD Regression Results using 2004-2012 Passenger Cars

	DD (Pooled)		DD (High-cost Only)		DD (Lowt-cost Only)		DDD (Pooled)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Constant	-2.7131 *** (0.5961)	-3.3684 *** (0.6272)	-6.0696 *** (1.0495)	-7.0484 *** (1.0797)	1.3132 * (0.7605)	2.4093 *** (0.7715)	-3.1886 *** (0.5849)	-3.6552 *** (0.6170)
T [High Slope = 1]	0.0371 *** (0.0049)	0.0362 *** (0.0049)	0.0363 *** (0.0083)	0.0451 *** (0.0086)	-0.0243 *** (0.0070)	-0.0064 (0.0071)	0.0001 (0.0069)	-0.0054 (0.0069)
R [Pol. Period = 1]	0.1396 *** (0.0051)	0.1414 *** (0.0051)	0.2077 *** (0.0086)	0.2073 *** (0.0084)	0.0905 *** (0.0057)	0.0988 *** (0.0057)	-0.0244 *** (0.0079)	-0.0108 (0.0081)
$R \times T$	-0.0551 *** (0.0085)	-0.0514 *** (0.0085)	-0.1521 *** (0.0120)	-0.1554 *** (0.0120)	0.0589 *** (0.0127)	0.0480 *** (0.0128)	0.0875 *** (0.0064)	0.0887 *** (0.0064)
H [High Cost = 1]							0.0882 *** (0.0104)	0.0690 *** (0.0106)
$H \times T$							0.1182 *** (0.0101)	0.1206 *** (0.0100)
$H \times R$							0.0715 *** (0.0144)	0.0860 *** (0.0147)
$H \times R \times T$							-0.2238 *** (0.0181)	-0.2376 *** (0.0185)
$\ln(\text{weight})$	-1.2473 *** (0.0331)	-1.3032 *** (0.0339)	-1.2813 *** (0.0524)	-1.3885 *** (0.0536)	-0.9083 *** (0.0470)	-0.9411 *** (0.0488)	-1.2376 *** (0.0332)	-1.3076 *** (0.0345)
$\ln(\text{hp})$	0.0177 (0.0214)	-0.0054 (0.0215)	0.017268 (0.0368)	-0.0092 (0.0366)	0.0883 *** (0.0270)	0.0569 *** (0.0272)	0.0476 * (0.0209)	0.0190 (0.0210)
$\ln(\text{torque})$	-0.1721 *** (0.0206)	-0.1359 *** (0.0206)	-0.1277 *** (0.0291)	-0.1051 *** (0.0285)	-0.2842 *** (0.0279)	-0.2655 ** (0.0288)	-0.1878 *** (0.0200)	-0.1455 *** (0.0202)
$\ln(\text{size})$	1.6817 *** (0.0896)	1.7891 *** (0.0939)	2.0653 *** (0.1519)	2.2501 *** (0.1552)	0.9902 *** (0.1177)	0.9011 *** (0.1185)	1.7217 *** (0.0880)	1.8202 *** (0.0925)
AT/CVT	0.0276 *** (0.0066)	0.0308 *** (0.0066)	0.0370 *** (0.0086)	0.0380 *** (0.0087)	-0.0025 (0.0096)	0.0049 (0.0092)	0.0239 *** (0.0064)	0.0265 *** (0.0064)
Maker fixed effects		✓		✓		✓		✓
R^2	0.9043	0.9083	0.6551	0.6784	0.9180	0.9257	0.9111	0.9145
Obs.	3,257	3,257	1,520	1,520	1,737	1,737	3,257	3,257

Note: Regressions exclude commercial vans and trucks, imported brands, diesel, hybrid, and electric cars as well as observations that fall in weight segments with the first and the third quartiles of compliance costs during the pre-2007 period. In parentheses are standard errors. The asterisks *, **, and *** indicate significance at 0.1, 0.05, and 0.01 levels.

Table 5. Regression Results on Placebo Treatments

	DD (Pooled)	DD (High-cost)	DD (Lowt-cost)	DDD (Pooled)
	(1)	(2)	(3)	(4)
<i>A. Perturbing Bin Assignments</i>				
DD or DDD Estimate	-0.0257 ** (0.0100)	0.0076 (0.0126)	-0.0132 (0.0182)	-0.0886 *** (0.0268)
R ²	0.9248	0.6732	0.941	0.805
Obs.	2,006	1,130	437	2,006
<i>B. Using 2004-2006 as Treatment Period</i>				
DD or DDD Estimate	-0.0165 *** (0.0057)	-0.0315 *** (0.0085)	0.0043 (0.0079)	-0.0375 *** (0.0121)
R ²	0.9248	0.6732	0.941	0.9266
Obs.	4,340	1,893	2,447	4,340

Note: Top panel A reports the results of regressions on a placebo treatment where weight cutoffs are shifted by $k = 25$ kg. Bottom panel B reports on another placebo treatment where the control period is 2001-2003 and the treatment period is 2004-2006. All regressions use the same covariates as in Table 3 including brand dummies. In parentheses are standard errors. The asterisks *, **, and *** indicate significance at 0.1, 0.05, and 0.01 levels.

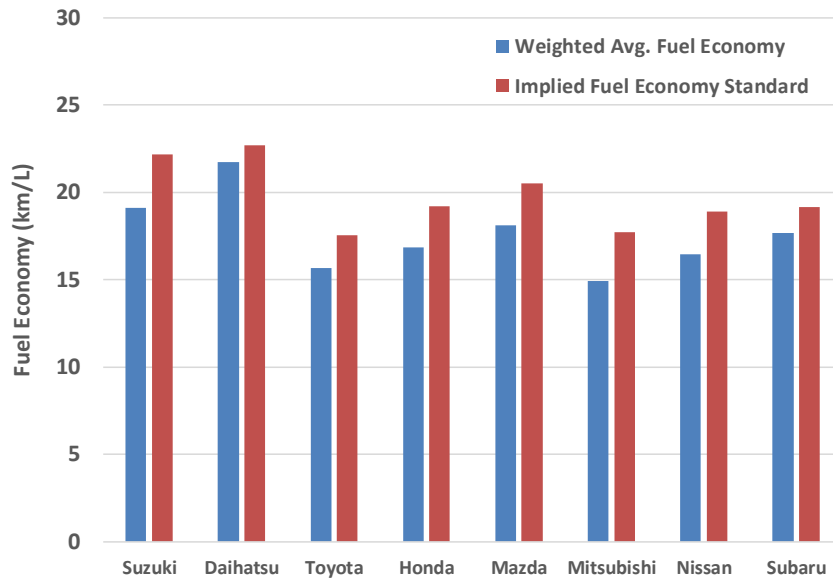
Table 6. Regression Results on Vehicle Weight

	DD (Pooled)	DD (High-cost Only)	DD (Lowt-cost Only)	DDD (Pooled)
	(1)	(2)	(3)	(4)
DD or DDD Estimate	0.0169 *** (0.0065)	0.0256 *** (0.0078)	0.0331 *** (0.0108)	-0.0142 (0.0138)
R ²	0.9778	0.9594	0.9443	0.979
Obs.	2,018	662	1,356	2,018

Note: All regressions use a subsample consisting only of bins with width less than 40 (in kg). In all regressions, logged curb weights are regressed on the same set of covariates as in Table 3, excluding brand dummies and logged weights. In parentheses are standard errors. The asterisks *, **, and *** indicate significance at 0.1, 0.05, and 0.01 levels.

Online Appendix.

Figure A1. Sales-weighted Fuel Economy and Standards by Maker, Year 2007



Note: The blue bar indicates the sales-weighted average fuel economy of vehicles sold in 2007 for each domestic car maker. The red bar indicates the estimated fuel economy standard for each maker, using the 2007 sales weights and the 2007 fuel economy standards. These statistics are estimates because we average out fuel economy and weight data over variants of each vehicle model. The exact sales data at the car variant level are not available. The figure demonstrates that at the beginning of the new standards, all domestic car makers were far behind the required fuel economy standards, and hence, are likely to have made some efforts to meet the standards during the post-2007 period.

Figure A2. Distribution of Curb Weights for Vehicle Models Assigned to High Slope Weight Bins

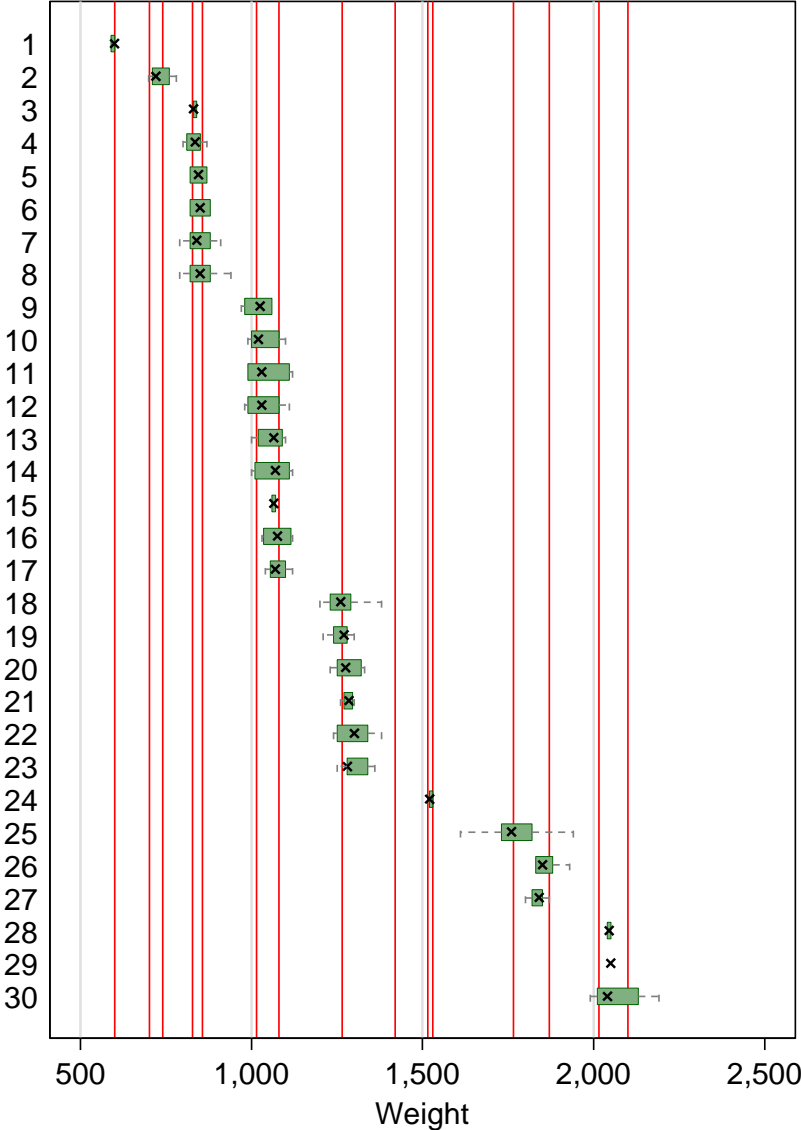


Table A1. Raw Image of Sample MLIT Table Used in Ito and Sallee (2018)

当該自動車の製造又は輸入の事業者の氏名又は名称

本田技研工業株式会社

ガソリン乗用車(普通・小型)

目標年度(平成22年度)

車名	通称名	原動機		変速装置の型式及び変速段数	車両重量(kg)	乗車定員(名)	10・15モード			主要燃費改善対策	その他燃費値の異なる要因		(参考)低排出ガス認定レベル	燃費基準達成・向上達成レベル		
		型式	型式				総排気量(L)	燃費値(km/L)	1km走行におけるCO2排出量(g-CO2/km)		燃費基準値(km/L)	主要排出ガス対策			駆動形式	その他
ホンダ	インサイト	DAA-ZE2	LDA(内燃機関) -MF6(電動機)	1.339	CVT(E)	1190	5	31.0	75	16.0	CY-V-I, EP-C-H	3W+EGR	F	タイヤ 185/60R15	☆☆☆☆	125
		DAA-ZE2	LDA(内燃機関) -MF6(電動機)	1.339	CVT(E)	1190	5	30.0	77	16.0	CY-V-I, EP-C-H	3W+EGR	F		☆☆☆☆	125
		DAA-ZE2	LDA(内燃機関) -MF6(電動機)	1.339	CVT(E)	1200	5	28.0	83	16.0	CY-V-I, EP-C-H	3W+EGR	F	タイヤ 185/55R16	☆☆☆☆	125
インサイト エクスクルーシブ	DAA-ZE3	LEA(内燃機関) -MF6(電動機)	1.496	CVT(E)	1200~1210	5	26.5	88	16.0	V-I, EP-C-H	3W+EGR	F		☆☆☆☆	125	
	DAA-ZE3	LEA(内燃機関) -MF6(電動機)	1.496	CVT(E)	1210	5					3W+EGR	F	タイヤ 185/55R16	☆☆☆☆	125	
フィット	DAA-GP1	LDA(内燃機関) -MF6(電動機)	1.339	CVT(E)	1130~1150	5					EP-C-H	3W+EGR	F		☆☆☆☆	125
	DBA-GE6	L13A	1.339	CVT(E-LTC)	1010	5	24.5	95	17.9	C-V-EP	3W+EGR	F	CVTウォーマー	☆☆☆☆	125	
	DBA-GE6	L13A	1.339	CVT(E-LTC)	1010	5	24.0	97	17.9	C-V-EP	3W+EGR	F		☆☆☆☆	125	
	DBA-GE6	L13A	1.339	CVT(E-LTC)	1030~1080	5	22.0	106	16.0	C-V-EP	3W+EGR	F	CVTウォーマー	☆☆☆☆	125	
	DBA-GE6	L13A	1.339	CVT(E-LTC)	1030~1080	5	21.5	108	16.0	C-V-EP	3W+EGR	F		☆☆☆☆	125	
	DBA-GE6	L13A	1.339	5MT	990	5	21.0	111	17.9	V-EP	3W+EGR	F		☆☆☆☆	115	
	DBA-GE6	L13A	1.339	5MT	990~1010	5	21.0	111	17.9	V-EP	3W+EGR	F		☆☆☆☆	115	
	DBA-GE8	L15A	1.496	CVT(E-LTC)	1070~1090	5	20.0	116	16.0	C-V-EP	3W+EGR	F	走行抵抗改善	☆☆☆☆	125	
	DBA-GE8	L15A	1.496	CVT(E-LTC)	1080~1100	5	20.0	116	16.0	C-V-EP	3W+EGR	F		☆☆☆☆	125	
	DBA-GE8	L15A	1.496	CVT(E-LTC)	1080~1100	5	19.6	118	16.0	C-V-EP	3W+EGR	F		☆☆☆☆	120	
	DBA-GE8	L15A	1.496	CVT(E-LTC)	1080~1100	5	19.2	121	16.0	C-V-EP	3W+EGR	F	CVTウォーマー	☆☆☆☆	120	
	DBA-GE8	L15A	1.496	CVT(E-LTC)	1080~1100	5	18.8	123	16.0	C-V-EP	3W+EGR	F	タイヤ 185/55R16	☆☆☆☆	115	
	DBA-GE8	L15A	1.496	6MT	1050~1080	5	17.4	133	16.0	V-EP	3W+EGR	F		☆☆☆☆	105	
	DBA-GE8	L15A	1.496	5MT	1050~1080	5	17.2	135	16.0	V-EP	3W+EGR	F		☆☆☆☆	105	
	DBA-GE7	L13A	1.339	5AT(E-LTC)	1140~1170	5	17.2	135	16.0	V-EP	3W+EGR	A	走行抵抗改善	☆☆☆☆	105	
	DBA-GE7	L13A	1.339	5AT(E-LTC)	1140~1170	5	17.0	137	16.0	V-EP	3W+EGR	A		☆☆☆☆	105	
	DBA-GE9	L15A	1.496	5AT(E-LTC)	1170~1180	5	16.4	142	16.0	V-EP	3W+EGR	A	走行抵抗改善	☆☆☆☆	100	
	DBA-GE9	L15A	1.496	5AT(E-LTC)	1160~1170	5	16.2	143	16.0	V-EP	3W+EGR	A		☆☆☆☆	100	
	フィット シヤトル	DAA-GP2	LDA(内燃機関) -MF6(電動機)	1.339	CVT(E)	1190~1230	5	30.0	77	16.0	CY-V-I, EP-C-H	3W+EGR	F		☆☆☆☆	125
		DBA-GG7	L15A	1.496	CVT(E-LTC)	1140~1170	5	20.0	116	16.0	C-V-EP	3W+EGR	F		☆☆☆☆	125
DBA-GG7		L15A	1.496	CVT(E-LTC)	1160~1180	5	19.2	121	16.0	C-V-EP	3W+EGR	F	タイヤ 185/55R16	☆☆☆☆	120	
DBA-GG8		L15A	1.496	5AT(E-LTC)	1210~1240	5	16.4	142	16.0	V-EP	3W+EGR	A		☆☆☆☆	100	
フリード	DAA-GP3	LEA(内燃機関) -MF6(電動機)	1.496	CVT(E)	1380~1420	5/6/7	24.0	97	13.0	V-I, EP-C-H, CY	3W+EGR	F		☆☆☆☆	125	
	DBA-GB3	L15A	1.496	CVT(E-LTC)	1270~1380	4/5/6/7/8	17.0	137	13.0	C-V-EP	3W+EGR	F	CVTウォーマー	☆☆☆☆	125	
	DBA-GB3	L15A	1.496	CVT(E-LTC)	1270~1350	4/5/7/8	16.4	142	13.0	C-V-EP	3W+EGR	F		☆☆☆☆	125	
	DBA-GB4	L15A	1.496	5AT(E-LTC)	1360~1460	4/5/6/7	14.0	166	13.0	V-EP	3W+EGR	A		☆☆☆☆	105	
	DBA-GB4	L15A	1.496	5AT(E-LTC)	1350~1440	4/5/7	14.0	166	13.0	V-EP	3W+EGR	A		☆☆☆☆	105	
フリード スパイク	DAA-GP3	LEA(内燃機関) -MF6(電動機)	1.496	CVT(E)	1380~1420	5/6/7	24.0	97	13.0	V-I, EP-C-H, CY	3W+EGR	F		☆☆☆☆	125	
	DBA-GB3	L15A	1.496	CVT(E-LTC)	1270~1380	4/5/6/7/8	17.0	137	13.0	C-V-EP	3W+EGR	F	CVTウォーマー	☆☆☆☆	125	
	DBA-GB3	L15A	1.496	CVT(E-LTC)	1270~1350	4/5/7/8	16.4	142	13.0	C-V-EP	3W+EGR	F		☆☆☆☆	125	
	DBA-GB4	L15A	1.496	5AT(E-LTC)	1360~1460	4/5/6/7	14.0	166	13.0	V-EP	3W+EGR	A		☆☆☆☆	105	
	DBA-GB4	L15A	1.496	5AT(E-LTC)	1350~1440	4/5/7	14.0	166	13.0	V-EP	3W+EGR	A		☆☆☆☆	105	
ステップワゴン	DBA-RK1	R20A	1.997	CVT(E-LTC)	1600~1680	7/8	16.0	145	10.5	C-V-I-EP	3W+EGR	F		☆☆☆☆	125	
	DBA-RK1	R20A	1.997	CVT(E-LTC)	1630~1660	7/8	15.8	147	10.5	C-V-I-EP	3W+EGR	F	タイヤ 205/60R16	☆☆☆☆	125	
	DBA-RK1	R20A	1.997	CVT(E-LTC)	1580~1670	8	14.2	163	10.5	C-V-EP	3W+EGR	F		☆☆☆☆	125	
	DBA-RK1	R20A	1.997	CVT(E-LTC)	1580~1670	8	14.2	163	10.5	C-V-EP	3W+EGR	F		☆☆☆☆	125	
	DBA-RK1	R20A	1.997	CVT(E-LTC)	1610~1680	8	14.0	166	10.5	C-V-EP	3W+EGR	F	タイヤ 205/60R16	☆☆☆☆	125	
	DBA-RK1	R20A	1.997	CVT(E-LTC)	1610~1680	8	14.0	166	10.5	C-V-EP	3W+EGR	F	タイヤ 205/60R16	☆☆☆☆	125	
	DBA-RK3	R20A	1.997	CVT(E-LTC)	1690~1700	7	16.0	145	10.5	C-V-I-EP	3W+EGR	F		☆☆☆☆	125	
	DBA-RK3	R20A	1.997	CVT(E-LTC)	1710	7	15.8	147	10.5	C-V-I-EP	3W+EGR	F	タイヤ 205/60R16	☆☆☆☆	125	
	DBA-RK3	R20A	1.997	CVT(E-LTC)	1680~1700	7	14.2	163	10.5	C-V-EP	3W+EGR	F		☆☆☆☆	125	
	DBA-RK3	R20A	1.997	CVT(E-LTC)	1680~1700	7	14.2	163	10.5	C-V-EP	3W+EGR	F		☆☆☆☆	125	
DBA-RK3	R20A	1.997	CVT(E-LTC)	1680~1700	7	14.0	166	10.5	C-V-EP	3W+EGR	F	タイヤ 205/60R16	☆☆☆☆	125		

Curb weights in kg

