

**Do Regulatory Loopholes Distort Technical Change?
Evidence from New Vehicle Launches under the Japanese Fuel Economy
Regulation ¹**

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

¹We thank Werner Antweiler, Brian Copeland, Corbett Grainger, Sumeet Gulati, Tomohiko Inui, Daiji Kawaguchi, Carol McAusland, Hikaru Ogawa, Hiroshi Ohashi, and other seminar participants at the University of British Columbia, the University of Tokyo, the National Institute of Environmental Studies (Japan), the National Graduate Institute for Policy Studies, 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 regulatory loopholes on technical change in the context of the Japanese fuel-economy regulation. We construct a model of firm's product choice under technology constraints (i) to substantiate the importance of the distinction between first-stage technology choice versus second-stage choice of product attributes for welfare and policy evaluation and (ii) to generate a testable economic hypothesis that attribute-based technology standards distort technical change when they create trade-offs between product attributes that differ from technically feasible trade-offs. We test this hypothesis in a unique data set on new product launches in the Japanese automobile industry. To control for the Ratchet effect and other time-varying confounders, we exploit the unique quasi-experimental variations in terms of width and stringency of standards across product segments over time, and construct two-fold treatment-control structures within each product segment. We find that contrary to earlier studies, the attribute-based standards did not induce sharp bunching at the notches, yet the differences in attribute trade-offs arising from these notches indeed induced different rates of technical progress, causing substantial distortion in technical change in the Japanese automobile industry.

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

Key Words: Attribute-based regulation, automobile, difference-in-difference-in-differences, energy policy, fuel economy regulation, Ratchet effect, regulatory loopholes, 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 firm’s technology choice, the effects of which can potentially persist over time via the market size effect of technical progress [Acemoglu (2002), Acemoglu *et al.* (2012), and Aghion *et al.* (2016)].

Earlier studies find clear and convincing evidence that firms indeed exploit regulatory loopholes in a variety of contexts — a flexible-fuel exemption under the U.S. CAFE regulation (Anderson and Sallee, 2011), a notched schedule of the U.S. Gas Guzzler tax (Sallee and Slemrod, 2012), and weight-basing under the Japanese fuel-economy regulation (Ito and Sallee, forthcoming). These studies, however, primarily focus on firm’s ‘second-stage’ product choice conditional on their ‘first-stage’ technology choice, and thus assume away potential distortion in technical change. Consequently, these studies implicate that regulatory loopholes are welfare-enhancing, offering regulated firms low-cost compliance strategies, under the given regulatory setup.¹ Such conclusion may change, however, once we account for their potential distortionary impacts on technical change.

In this manuscript, we empirically investigate if a regulatory loophole can indeed distort firm’s technology choice, in the context of automobile fuel-economy regulation in Japan. We first start with Knittel’s (2011) observations that automakers face technical trade-offs between fuel economy and other vehicle attributes and that these trade-offs change substantially over time. We incorporate these observations in a simple, static model of firms making a three-stage decision on technical upgrade, product choice, and pricing. The model gives us a clear, testable economic prediction: A fuel-economy regulation distorts technical change when it creates trade-offs between the targeted and secondary attributes that differ from technically feasible trade-offs. The model also allows us to clarify the important distinction between distortion in the first-stage technology upgrade versus that in the second-stage

¹Of course, the loopholes are never optimal in the first-best setting where the efficient regulation can be costlessly implemented. See **Section 6** for a more thorough review of the welfare effects in these studies.

product choice and their resulting implications for welfare and policy evaluation. We then further exploit Knittel’s insight, and identify the distortion in technology upgrade by its impact on observed technical trade-offs in the product-attribute space.

In empirical implementation, 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 vehicle weight: i.e., vehicles are classified into discrete weight segments with varying levels of fuel economy standards. Importantly, when revising the standards in 2007, the regulatory authority chose narrower weight segments, effectively creating two or more weight bins within each old weight segment. Consequently, we have substantial variations, in terms of stringency and width, across weight bins over time. We translate these variations into two measures of regulatory assignment: (i) the ‘stringency’ of fuel-economy standards, measured in relative terms to the old standards, and (ii) the ‘slope’ of fuel-economy regulation, measured as a decrease in the fuel economy standard per unit of weight increase. The latter is indeed a convenient measure of the attributes trade-offs induced 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 confounders that may be correlated with the regulatory assignments.

To that end, we employ a difference-in-difference-in-differences (DDD) research design, exploiting the three-fold control structures: (a) cross-sectional with models assigned to low regulatory slopes 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 standards as an additional control. This last control is intended to serve 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 last point is particularly important if firms engage in the Ratchet-type behavior (we discuss this in more depth in **Section 4**). We then combine this DDD strategy with several control strategies for other time-varying confounders. First, we construct treatment-control pairs within each old weight segment, and use maker- and segment-fixed effects interacted with time dummies. This allows us to control for confounders that arise from (segment-level) consumer demand, firm-level heterogeneity in technical progress, and tax/subsidy incentives offered during the post-2007 period. Second, we exploit variant-level vehicle attribute data to control for influence of variant-level attributes on fuel economy ratings. Third, we use two alternative measures, one exploiting bin-level variations and another exploiting model-level variations, for both the ‘slope’ and the ‘stringency’ of the regulation. This way, we are

able to attribute the difference in outcomes solely to the difference in attribute trade-offs created due to the fuel-economy regulation. We implement this DDD strategy using vehicle characteristics data for all domestic passenger vehicles introduced between 2004 and 2012, excluding electric, diesel, and hybrid cars as well as those launched in the interim regulatory period 2007-2009.

We have three empirical findings. First and foremost, we find strong evidence that regulatory loopholes in this context had a sizable, statistically significant distortionary effect on technical change. Our DDD estimate indicates that a one-unit increase in the steepness of the regulation causes a 17 percentage point (ppt) reduction in the rate of fuel-economy improvement. We emphasize that we obtain this qualitatively large impact after controlling for all vehicle-level observables (incl. vehicle weight), all time-varying brand and segment effects as well as other confounders that likely correlate with regulatory stringency. 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, the Japanese passenger cars had the same rate of technical progress just over the 8-year period between 2004 and 2012. The 17-ppt reduction represents roughly 1/3 of this technical change.

Second, we find that the notched schedule of the regulation does offer regulatory loopholes, but firm's incentives to exploit the loopholes vary substantially due to the differences in attribute trade-offs that arise from these notches. In other words, not every notch is equally important. We confirm this by examining the effect of regulatory slope on vehicle weights in a manner analogous to the main regression. Our results indicate that holding the stringency of the regulation, firms increase the weights of their vehicle models more when they are faced with steeper regulatory slopes than faced with flatter slopes. The magnitude of the impact is also qualitatively large — a one-unit increase in the steepness of the regulation causes a 12 ppt increase in vehicle weight. This finding also confirms the economic mechanism underlying our first finding.

Third, our analysis also reveals that sharp bunching at the notches of the regulation reported in Ito and Sallee (forthcoming) largely disappears when we construct a weight distribution using the data on actual product launches. Ito and Sallee use the fuel economy data published by the Ministry of Land, Infrastructure, and Transportation (MLIT). One critical drawback of the MLIT data is that they report only the range (i.e., minimum and maximum) of vehicle weights for roughly 3/4 of the observations. The weight range can be as large as 200 kg, averaging at around 35 kg. In contrast, we construct our dataset

using the web catalogue of one of the largest car dealers, Carsensor, in Japan. The dataset contains all vehicle offerings at the variant level *and* reports each vehicle’s weight with a unique value. Our analysis demonstrates that bunching behavior reported in Ito and Sallee is driven primarily by the minimums (rather than the maximums) of the weight ranges in the MLIT data, indicating the existence of firm’s incentive not to report their vehicles’ weights to cross over to a lighter weight segment. We therefore conclude that their finding is most likely the evidence of reporting bias by automakers, but not the evidence of distortion in actual vehicle offerings in the market.

This manuscript contributes to four strands of literature. First, this manuscript nicely complements the literature investigating regulatory loopholes in environmental regulation (e.g., Anderson and Sallee, 2011; Sallee and Slemrod, 2012; and Ito and Sallee, forthcoming). While they differ in context and focus, they all provide evidence that firms do exploit the loopholes when they reduce (marginal) costs of compliance. Sallee and Slemrod (2012) examine a notched schedule in the U.S. Gas Guzzler Tax and show that automakers adjusted their vehicles’ fuel economy ratings so as to qualify for lower tax rates. 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. Ito and Sallee (forthcoming) examine attribute-based regulations both theoretically and empirically, focusing on the same regulatory setup as ours. Our study closely parallels, but differ substantially from these studies on one important account. That is, we get at the effect of regulatory loopholes on firm’s first-stage technology/product choice while these studies get at firm’s second-stage choice of product attributes conditional on the first-stage choice.

Second, this study provides important implications for the literature examining the optimal design of environmental policy instruments under second-best settings. In first-best settings, a gasoline tax can fully restore economic efficiency, internalizing the negative externality associated with gasoline consumption (e.g., Fullerton and West, 2002). There are numerous obstacles that arise in real-world settings such as equity concerns in the presence of consumer heterogeneity, oligopolistic nature of automobile industries, and (possible) optimization failures in consumer choice. Under such second-best settings, the gasoline tax is not necessarily an ideal policy instrument. Economists have taken these real-world complexities seriously, and have initiated search for alternative policy instruments. In this regard, feebates and fuel-economy standards are now considered as two viable options in addressing some of these complexities (Anderson *et al.*, 2011).² In theory, feebates and fuel-economy

²A feebate is a fiscal instrument that charges a fee on purchase of high-emission products and gives a

standards with efficient credit trading are equivalent policy instruments. This may not be true, however, if firms exploit the regulatory loopholes that originate from attribute-basing in fuel-economy standards. In this context, Ito and Sallee (forthcoming) offers an important policy advice: A flat fuel-economy standard is optimal in the presence of efficient credit trading. This policy advice is well grounded in the theory of environmental regulation, which establishes that the efficiency of the permit trading market requires the allocation of permits to be independent of firm’s endogenous choice. Our results cast an important twist to this debate. Our results suggest that the flat standard is likely to bias firm’s technology/product choice substantially even with credit trading. To assure no bias, instead, regulators need to make regulatory trade-offs between attributes as close as technically feasible trade-offs, for otherwise, firms could endogenously choose the level of fuel-economy technologies to be eligible for more credits initially.

Third, a number of studies have empirically investigated the effects of U.S. Corporate Average Fuel Economy (CAFE) regulation on vehicle emissions and welfare (e.g., Austin and Dinan, 2005; Jacobsen, 2012; Goldberg, 1998). Jacobsen, for example, estimates the model of heterogeneous households and producers and explicitly accounts for the fact that the CAFE constraints may be binding only for a subset of the firms. Focusing on the distributional implications of the CAFE, he finds that the profit impacts of an increase in the CAFE standards fall almost entirely on domestic firms, and hence, concludes that harmonization of compliance costs across firms is likely to enhance economic efficiency of the CAFE regulation. In these studies, however, firms can respond to the CAFE regulation only by pricing or by improving fuel-economy technologies (with an arguably simplistic engineering function that relates the investment cost to the fuel-economy rating), largely ignoring endogenous product choice. Recent advances in the empirical industrial organization literature substantiate the importance of endogenous product choice (e.g., Seim, 2006; Hitsch, 2006, Fan, 2013, Crawford *et al.*, 2015, Wollmann, forthcoming), both qualitatively and quantitatively, for policy and welfare evaluation. Our results indeed indicate that the fuel-economy regulation alters firms’ technology/product choice, endogenously and in an intricate manner. Incorporating this aspect in the empirical study of the CAFE regulation could be an important direction for future research.

Lastly, since Porter (1991), a large number of empirical studies examined the effect of environmental regulation on innovation and technical change (see, for example, a comprehensive review by Jaffe *et al.*, 2002). This strand of literature essentially focuses on the ‘price effect’

rebate on purchase of low-emission products. Konishi and Zhao (2017) have indeed shown, in a real empirical context similar to ours, that feebates could be designed as a corrective tax instrument to address externality and imperfect competition simultaneously.

of technical change in the language of Acemoglu (2002). That is, the regulation-induced increase in the implicit cost of pollution promotes innovation in sectors that require pollution as by-product of their production activities. Fuel-economy regulation can indeed be cast as a regulation that raises the implicit price of pollution. Indeed, Newell *et al.* (1999) examine the effect of energy efficiency standards for home appliances in the product-characteristic approach similar to ours, and find that the standards induced modest technical changes in home appliance models offered in the market. We add another angle to this literature: design details matter in examining the price effect of regulation. While the environmental regulation is intended to raise the implicit cost of pollution, its design features can offer regulatory loopholes, which can distort the intended price signal.

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

Energy efficiency regulations around the world are often attribute-based: it relies on a secondary attribute that is not the direct target of the regulation. For example, fuel economy or carbon emissions standards are a function of vehicle footprint in the U.S. and of vehicle weight in Japan and the EU. 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.

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. In their model, however, there is no clear distinction between bias in technology choice versus bias

in product choice, and as a result, their empirics produce estimates that are difficult to interpret. The objective of this section, therefore, is to present a simple model of firm’s product choice under technology constraints in order to substantiate the importance of the distinction between the two types of bias for empirical analysis.

To do so, we first introduce the concept of a technology possibility frontier (TPF). We believe the concept is implicit in Ito and Sallee and other related studies, but has not been explicitly addressed. The concept somewhat originates from Knittel (2011), who finds 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. 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** displays 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.

We infer from this empirical regularity that there exists a technology possibility frontier (TPF), which we define as the set of product attributes that are technically feasible when technical inputs are used most efficiently given the technology capital. 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 in designing 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 should 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.³ Technical change can then be identified by changes in the level of the TPF. In what follows, we demonstrate the importance of the TPF in understanding the distortionary impact of attribute-based

³The TPF does not describe technical trade-offs that induce the same production cost or profit level like Knittel (2011) or Newell *et al.* (1999), who defines technically feasible trade-offs as the iso-cost curve in the attribute space. 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].

regulation.

Our model follows the convention in 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, which we treat as fixed. Each firm m produces a unique product. We ignore 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 firm’s two-period decision: each firm m chooses the next-period product attributes (f^1, w^1) given 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 attributes (f^1, w^1) .

Given this setup, firms compete in three stages. 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.⁴ We describe the third-stage product-specific profit for each firm 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 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.

To focus on the essentials, we make further simplifying assumptions on π, c, T and R . (A1) The TPF function T is linear and strictly decreasing in w , with $\partial^2 T / \partial w \partial s = 0$, in the

⁴We 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.

neighborhood of (f^0, w^0) from which (f^1, w^1) can be profitably chosen. That is, technical upgrade can only shift up and down the linear TPF schedule in the relevant decision space. (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 linear and decreasing in w , again in the neighborhood of (f^0, w^0) .

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, automakers can generally choose new car models around the neighborhoods of their pre-existing models given their platform designs, and are likely to face approximately linear technical trade-offs in the neighborhoods. In theory, it is known that (A2) may not hold in general. Without it, however, one cannot solve for and characterize the equilibrium. Assumption (A3) is 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) there are two types of distortion, one on technology choice and the other on product attributes given the technology choice. **Figure 2** helps us illustrate the economic significance of this proposition for empirical analysis. Imagine a product segment, say 'light sedan', which implicitly defines a (small) two-dimensional segment on the (f, w) space. On this segment, the current product offering, denoted O , is on the solid linear line, which represents the current-period TPF. The dashed curves 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. 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 line denoted R 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 us 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 . 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 southeast of A , like C . (C would lie to the southwest of A if R is flatter than the TPF). The distortionary impact of regulation (from A to C) can, therefore, be decomposed into one on technology choice (from A to A') and the other on product attribute holding the technology level (from A' to C). This figure also clarifies the importance of the TPF concept for empirical analysis. In **Figure 2**, C is further away from O than A , and hence, without accounting for TPF, one may be tempted to infer that C might be more costly than A . The concept of TPF clarifies, however, that A lies on a higher TPF than C does, and hence, A is indeed more costly to the firm than C .

2.B. From Theory to Empirics

Our empirical approach 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 thus 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 assumptions for identification. 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, although 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.⁵ 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, change the characteristics of car grades, 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* 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 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. 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

⁵Nowadays, 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.

Act for all manufacturing products that consume energy in utilization. Under the *Top-runner* system, the government first classifies each vehicle to a unique product category according to its vehicle weight, and then chooses the highest observed fuel economy rating 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.⁶ The Ministry of Land, Infrastructure, Transport and Tourism (MLIT) adopted a new fuel economy rating method, known as *JC08 Mode*, for the new standards. The figure reports the new standards in the old measure (known as *10.15 Mode*). The method of conversion between the two measures is described in **Subsection 3.B.** in more detail.

The new 2007 standards created an interesting regulatory setup, and thus, is a focus of our study. Under the new standards, the government chose a narrower weight segment to define each product category, resulting in 16 new weight segments in contrast to 9 under the old standards. As a result, each old weight segment was effectively divided into two or more bins, resulting effectively in 24 weight bins in total under the new standards. For some reason (not transparent in regulatory documents), the segment width differed substantially across weight segments. Furthermore, because the fuel economy performance of the top-runner relative to the peers in the same old weight segment 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.

There are a few more regulatory backgrounds that become important in our empirical implementation. First, the Japanese fuel-economy regulation does not permit credit trading, either across firms or segments. On one hand, this helps our identification since it eliminates the potential confounding effect from credit trading. On the other hand, it suggests the need to control for heterogeneous firm incentives because the marginal costs of compliance are unlikely to be equalized across weight segments and firms. **Figure A1 in the Appendix**

⁶In 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.

shows that at the beginning of the new standards, all domestic car makers were behind the required fuel economy standards, and hence, heterogeneity is probably more important at the car-model level. Second, the Japanese government 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. To isolate the confounding impact of these tax incentives, we make use of a treatment-control structure within each old segment, and also constrain our main empirical analysis up to year 2012 (see our identification strategies in **Section 4**).

Lastly, one more aspect of the Japanese regulation needs some discussion. Fines for non-compliance are only 1 million JPY (\approx \$10,000) per *firm*. Moreover, the Japanese standards are not enforced every year, and instead, firms are expected to meet the standards only by the (respective) target years. In contrast, under the U.S. CAFE, fines are \$55 *per vehicle sold, for every mile-per-gallon shortfall*. The National Highway Traffic Safety Administration reports that the U.S. automobile industry has been paying roughly \$20 million annually since 2010 (*AutomotiveNews*, July 16, 2016). Despite this weak incentive structure, however, Japanese firms take these standards very seriously, plausibly in fear of non-pecurinary sanctions such as damaging customer reputation and unfavorable treatment in public procurement. All firms met the 2001 standards by 2005 well ahead of its target year 2010 (and before the start of the tax/subsidy incentives in 2009). Hence, the new 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. Anecdotal evidence suggests that compliance with the standards is indeed very costly for at least some firms — a recent scandal revealed that Mitsubishi Motors had been inflating fuel economy ratings for nearly 20 models over the last 10 years (*Japan Times*, Jun 17, 2016). Furthermore, Ito and Sallee (forthcoming) show that firms do respond, very sharply indeed, to the weight cutoffs of the Japanese fuel economy standards.

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.⁷ In our placebo 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**.⁸ 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 regulation. 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.

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*.⁹ From here on, all fuel economy data are reported in the *10.15 mode*.

3.C. Reporting versus Product Offerings

We clarify important differences between our data and the data used in Ito and Sallee (forthcoming), and explain why Ito-Sallee data cannot be used in our analysis (even for

⁷The 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.

⁸There 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.

⁹The regression is surprisingly well fit with $R^2 \approx 0.99$.

cross-validation purposes). Their data come from the list of new cars published each year by the MLIT. The MLIT list reports data at the car configuration (or ‘*Katashiki*’) level, which is coarser than the grade level reported in the Carsensor catalog. There are two aspects of the MLIT data that make it unsuitable for our analysis. First, the MLIT list contains all cars sold as ‘new cars’ as of the end of each fiscal year. As a result, some cars 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. Second, and more importantly, the MLIT list reports only the *range* of vehicle weights for about 3/4 of the reported car models while reporting the *unique* fuel-economy rating for each model (see **Figure A2 in the Appendix** for the raw image of the MLIT table). This range can be as large as 200 kg, averaging at around 35 kg (see **Table 1**). Hence, if accurately reported, the MLIT data would imply each car model has a flat TPF in the short run.

Table 1 reports weight distributions in the two data sets. The MLIT data contain a smaller number of observations in each year than our Carsensor data (despite their possible double counting). Of 2,012 observations in the MLIT data between 2010 and 2012, only 25% report exact weights. The remaining 75% of observations report only minimum and maximum weights. 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. 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.

Figure 4 visualizes these differences in weight distributions between the two data sets. **Panel (a) of Figure 4** displays three weight distributions using the MLIT data, for all car configurations reported between 2010 and 2012: (i) observations reported with exact weights, (ii) minimum weights (for those with weight ranges), and (iii) maximum weights (for those with weight ranges). **Panel (b) of Figure 4** displays the same using our grade-level data. The figures confirm three points: (1) in the MLIT data, significant bunching occurs at the weight cutoffs, but the incidence of bunching disappears in our data; (2) in the MLIT data, bunching is primarily driven by the observations reported with range, and 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; and (3) bunching mostly corresponds

to the 2001 standards, not the 2007 standards. This last point 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. We believe the weight distribution in our data is more consistent with findings in the empirical IO 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 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.

The question arises naturally then: What explains the behavior in the MLIT data? That is, why do automakers they report the minimum weights so as *not* to cross over to the *lighter* weight category while they report the maximum weights so as *not* to cross over to the *heavier* weight category? Our explanation is as follows. First, the regulatory agency assigns a car model to the lightest weight category when it weights range over two or more weight categories. Therefore, automakers have a very strong incentive not to cross over to the lighter weight category. Second, automakers offer many different variants of the same car model/configuration, yet at the time of reporting the new model data to the MLIT, they do not know how well the new model would perform in the markets, and hence, how many variants of the model they wish to offer, over the course of the model year. 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 category. We take these weight distributions as suggesting that the MLIT data offer 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.

4. Identification and Estimation

4.A. Identification Strategy: An Overview

The goal of our empirical study is to estimate the causal (distortionary) impact of the technical attribute trade-offs imposed by the fuel-economy regulation. To do so, we need to control for other confounders that may be correlated with regulatory assignment — i.e.,

confounders that remain even after controlling for observable vehicle characteristics such as weight, horsepower, and manufacturer fixed effects. The first is the consumer demand. It is known that consumers who buy larger and heavier cars tend to care less about fuel economy ratings (Berry *et al.*, 1995; Konishi and Zhao, 2017). Car models under different weight categories may naturally have different rates of technical progress, irrespective of regulatory assignment. Second, as discussed in **Subsection 3.A.**, the government offered tax/subsidy incentives according to fuel economy ratings since 2009. We thus need to tease out the effect of these tax incentives. The third, and probably the most important, confounder concerns the Ratchet effect.¹⁰ The Ratchet effect refers to the phenomenon that the agent underperforms to avoid a demanding schedule in the future in a dynamic incentive scheme where the principal updates the scheme over time upon observing the agent’s performance (e.g., Freixas *et al.*, 1985; Laffont and Tirole, 1988). If exists, the Ratchet effect would imply that the rate of technical change in a weight category in the future may correlate with the current fuel-economy standard for that category since under the *Top-runner system*, the regulator chooses the best observed fuel-economy rating as the standard. It is difficult to isolate the Ratchet effect from the effect of regulatory trade-offs because both arise from the same design features of the fuel-economy regulation.

We attempt to control all these confounders in the following ways. First, we create treatment-control pairs within each old weight segment, exploiting the fact that the 2007 standards created new and narrower weight categories. In addition, we include weight-segment dummies, interacted with time-series dummies, in our regression analysis. By this, we are able to compare the outcomes of regulatory assignments for car models that faced roughly the same demand shocks and the same tax incentives over time. Second, we also include maker fixed effects, again interacted with the time dummies, to control for firm-specific technical progress. Third, we exploit two types of variations in regulatory assignment, i.e., changes in the level of the fuel-economy standards and changes in the width of the weight segments. As discussed in **Subsection 3.A.**, some weight segments under the 2007 standard are more stringent than others in terms of required improvements relative to the old standards. At the same time, some segments are narrower than others, resulting in variations in the width of weight segments. We transform these variations into two measures of regulatory assignment: (1) the ‘stringency’ of fuel-economy standards, measured in relative terms to the old standards, and (2) the ‘slope’ of fuel-economy standards, measured as a decrease in fuel-economy standards per unit of decrease in vehicle weight. For robustness, we construct two alternative variables for each of these measures. These measures, we hope, would get at two types of economic incentives separately, the Ratchet effect and the effect

¹⁰We thank Hiroshi Ohashi for suggesting this point.

of attribute trade-offs (we call it the ‘slope effect’ henceforth). We discuss these measures in more depth in the next section. Lastly, we combine these with a difference-in-difference-in-differences (DDD) approach, exploiting the three-fold control structures as follows:

- (a) Cross-sectional between-group variation in regulatory slope
- (b) Temporal variation over years (with years 2004-2006 as control)
- (c) Cross-sectional within-group variation in regulatory stringency

Here, we treat car models faced with the same (or similar) regulatory slope(s) as a group. By using temporal variation with years 2004-2006 as an additional control, we are able to control for any stationary differences across groups as well as time-varying factors that are common to the groups. However, this difference-in-differences (DD) structure is not sufficient to control for the Ratchet effect or other time-varying confounders that affect these groups differently over time. To take care of this concern, we use another within-group variation. As we shall discuss more below, the stringency of the standards most likely captures the Ratchet behavior, and hence, serves as an additional within-group control. That is, we compare the outcomes of vehicles assigned to different regulatory slopes, but with the same (or similar) regulatory stringency level(s). The resulting DDD estimate is consistent under a weaker identifying assumption: *i.e., unobservables that affect the rate of technical progress differently across car models assigned to different regulatory slopes do not systematically differ across car models assigned to different stringency levels.* 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 same stringency level is also consistent.

4.B. Slope vs. Ratchet

To illustrate our empirical strategy, let us take an old weight segment 970-1,265 kg as an example for our exposition. In **Figure 5**, both old and new fuel economy standards are drawn (red and blue lines, respectively). The intersections of the two standards create three weight bins on this old weight segment. 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 standards approximately trace out the technology frontier of ‘the most fuel efficient’ vehicles that were available as of 2007. Consider a line connecting the two endpoints A and B of this segment. For the moment, let this line represent the technical frontier of a ‘typical’ or average firm. Again for the moment, let us define the regulatory slope of each weight bin as the slope of the line connecting endpoints of each bin. Then the regulatory slope is clearly steeper than the technical frontier for some bin(s), and flatter for other bin(s). In **Figure 5**, 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. As discussed in **Subsection 2.B.**, we assume that 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 low-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 points A and B of the old segment represents the technical frontier of a ‘typical’ firm. What matters for our empirical analysis is that vehicle models that lie on the same old segment are likely to face, on average, roughly the same market demand, the same regulatory incentives other than the slope, and the same technology frontier prior to the 2007 standards.

The question then is, what would be the most appropriate measure of the regulatory trade-offs? We offer two alternatives. The first measure directly applies the above logic, and calculates the slope of each bin b as the slope of a line connecting the two endpoints of the weight bin:

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

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. For the high-slope bin in **Figure 5**, this measure

is simply the slope of the line connecting A and C . The advantage of this measure is that it uses only the variations in regulatory assignment, and hence, it is unlikely to be correlated with other confounders at the firm or the model level, especially after controlling for the Ratchet effect (which we discuss below).

The disadvantage, however, is that it fails to account for firm-level or model-level heterogeneity. Even within a weight bin, different car models have different fuel-economy ratings and vehicle weights at the onset of the new standards, and these differences in initial positioning are likely to present different regulatory trade-offs. For example, a car model located at position O would be able to lower the standard by Δh by increasing its weight by Δw . This reduction Δh represents a large gain relative to its required fuel-economy improvement. In contrast, another car model located at position O' would be able to attain the same benefit, but by increasing its weight more by $\Delta w'$. Our second measure, therefore, accounts for this heterogeneity arising from initial positioning. That is, we calculate the regulatory slope for each car model j as a reduction in the fuel-economy standard, expressed as a percentage of the required fuel-economy improvement for that model, per unit of weight increase required for that model:

$$\Gamma_j^* = \left| \frac{(h_{b+1}^{\text{new}} - h_b^{\text{new}})/(h_b^{\text{new}} - f_j)}{w_{b+1} - w_j} \right|. \quad (3)$$

Albeit its merit, the potential disadvantage of this measure is that it may be correlated with unobservables that affect the rate of technical progress, even after controlling for the Ratchet effect, because it explicitly uses initial fuel-economy information in its calculation.

We now turn to a more intricate confounder, the Ratchet effect. By construction, the regulatory slope depends on the stringency of the fuel-economy standards, and we have two reasons to believe that it may capture the Ratchet effect, rather than or in addition to, the slope effect. First, firms cannot observe competitors' technical progress prior to their product launches, and hence, are likely to base their Ratchet strategy on the regulatory standards. Second, firms are also likely to base their future Ratchet behavior on their past Ratchet behavior. Because the fuel-economy standards are the outcome of the past Ratchet behavior, they can be directly associated with the future Ratchet behavior. This logic suggests that we might be able to control for the Ratchet effect by controlling for the stringency of the fuel-economy standards. The key identifying assumption here is that firms can influence the level of fuel-economy standards by manipulating the rate of technical change, but cannot affect how the weight category is chosen, so the width (and the resulting slope) of each weight segment is an exogenous shock to the firms.

The remaining question is, what would be the appropriate measure of regulatory stringency? Like the slope effect, we consider two alternatives. The first measure simply computes

the difference between the old standard and the new standard for each weight bin b :

$$\Lambda_b = h_b^{\text{new}} - h_b^{\text{old}}, \quad (4)$$

whereas the second measure computes the difference between the new standard and the pre-policy fuel-economy rating for each car model j :

$$\Lambda_j^* = h_b^{\text{new}} - f_j. \quad (5)$$

In other words, the first measure simply evaluates the absolute stringency of the new standard for each weight bin while the second measure evaluates the relative stringency for each car model. The pros and cons of these stringency measures are analogous to those of the two slope measures. Because the *Top-runner* system chooses the highest observed fuel-economy rating as the standard for that segment, the first measure is likely to be directly related only to the top performer’s rate of technical progress. In contrast, the second measure is likely related to the own rate of progress (relative to the top performer). The latter is more likely to closely capture the Ratchet-type incentives for each car model, but is more likely to be endogenous than, the former.

4.C. Estimation Strategy

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 with different design features would respond differently to regulatory assignment. To exploit the variant-level variations, we employ parametric DDD regression in a manner analogous to Gruber (1994), augmenting the Cobb-Douglass

specification of the technical possibility frontier à la Knittel (2011):¹¹

$$\begin{aligned}
\ln f_{ijmst} = & \alpha + \beta_1 R_t + \beta_2 T_j + \beta_3 H_j \\
& \dots + \beta_4 (R_t \times T_j) + \beta_5 (R_t \times H_j) + \beta_6 (T_j \times H_j) \\
& \dots + \beta_7 (R_t \times T_j \times H_j) + \mathbf{X}'_{ijt} \boldsymbol{\gamma} + \boldsymbol{\eta}_{mst} + \epsilon_{ijst},
\end{aligned} \tag{6}$$

where $\ln f_{ijt}$ is a logged fuel economy of vehicle variant i of model j , introduced by firm m in year t , assigned to old weight segment s during the pre-2007 period, R_t indexes a regulatory period and equals 1 during the post-2007 period, T_j and H_j are our key treatment variables ('slope' and 'stringency' of the 2007 standards, respectively), X_{ijt} is a vector of observable vehicle characteristics [incl. weight (w), horsepower (hp), size ($size$), torque (tq) (all in logged values), and transmission type], and $\boldsymbol{\eta}_{mst}$ denotes maker- and segment-fixed effects and their interactions with R_t .

We emphasize here that several control strategies are nested in (6). First and foremost, we employ the triple difference incorporating H_j so as to isolate the slope effect from the Ratchet effect. Second, we exploit treatment-control pairs within each old weight segment, and use maker- and segment-fixed effects and their interactions with R_t . This allows us to control for confounders that originate from (segment-level) consumer demand, firm-level heterogeneity in technical progress, and tax/subsidy incentives during the post-2007 period.¹² Third, we exploit all variant-level attribute data to control for influence of variant-level attributes other than vehicle weight. Fourth, as discussed above, we consider two alternative measures for both T_j and H_j , each offering different advantages and disadvantages. Lastly, the regression also nests the DD structure: i.e., the DDD estimates are also consistent under the standard common-trend assumption that makes the DD estimates consistent. Combining

¹¹There is an alternative strategy. That is, 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. Our earlier attempt to employ PSM estimator confirms this prediction.

¹²As 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 discussed 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 segment.

these control strategies with the standard DDD argument, we see that the OLS estimate of β_7 identifies the causal impact of the regulatory slope under a much weaker assumption than the common trend assumption. One (potential) disadvantage is that our specification assumes the regulatory assignment can affect only the *level* of the technical frontier, not the *slope*, 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, our empirical strategies primarily exploit *within-segment* variations, and we doubt that we have large enough within variations to credibly identify the impact on the slope parameter(s).

We clarify two practical issues in estimating (6). First, our regulatory variables T_j and H_j vary at the car-model level, not the variant level. We can only do this because one variant introduced in a year cannot be credibly identified with another introduced in a different year.¹³ Therefore, we trace out model histories, so that all vehicle variants introduced during the post-2007 period can be associated, via model identifiers, with those introduced during the pre-2007 period. For models that continue to exist, this is easy because they can be easily matched by model identifier. For discontinued models, we search through publicly available articles and company reports to see if there is any successor model for each retired model. What complicates the issue is that not all variants of a model necessarily fall in a single weight bin because there are many variants of each vehicle model. To address it, we calculate the unweighted mean of vehicle weights of all variants for each vehicle model during the pre-2007 period, and then classify the vehicle model according to that mean.¹⁴ Second, as noted above, we would like to ensure a treatment-control pair in each of old weight segments. For the second measures, this can be easily achieved since we let $T_j = \Gamma_j^*$ and $H_j = \Lambda_j^*$ for each j and as a result, there are several models, with sufficient variations, in every segment. The problem is with the first measures, which vary only by weight bin. To ensure treatment-

¹³It is highly questionable to identify vehicle variants according to their attribute data, at least in our context. For example, suppose we observe two variants of Honda Civic, one introduced in 2004 and another in 2012, that have the same displacement, horsepower, etc. Suppose further that Honda Civic went through a significant platform change between the two years — many models would indeed go through such model change during such a long period. In that case, it seems natural to treat these variants as different variants.

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.

Figure A3 in the Appendix reports the summary of model histories and box diagrams describing the distribution of variant-level curb weights 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.

control pairs in all segments, we classify weight bins into high- versus low-slope bins according to whether their slopes are steeper than the slope of the joint segment connecting all bins within each old segment as illustrated in **Figure 5**. By this, we are assuming that vehicles within each old weight segment faced roughly the same technical frontier and that a new segment steeper than this average slope give more incentives to manipulate on vehicle weight. Similarly, we also classify weight bins into quartiles of stringency levels, with 1 denoting bins that fall in the lowest 25th percentile and 4 that fall in the highest 25th percentiles.

Table 2 clarifies these points, highlighting the main sources of variation we exploit in our analysis. Each row represents a 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 of these weight bins, we report old and new fuel-economy standards, regulatory slope and stringency in two alternative definitions, the number of vehicle variants, and the mean and standard deviation of fuel-economy ratings during the pre- and the post-2007 periods. The first measure of slope Γ [eq. (2)] calculates the slope using the ‘height’ and ‘width’ of weight bins, and therefore, has a unique value for each weight bin (recall **Figure 5**). The next column reports 1 if this slope is higher than the overall slope of the weight segment joining all weight bins that belong to the weight segment. This defines our main treatment variable T_j for our first measure. The second measure Γ^* [eq. (3)] instead computes the slopes for all car models that belong to each bin, and therefore, we report the mean and standard deviation in each row. We report two measures of regulatory stringency Λ [eq. (4)] and Λ^* [eq. (5)] in an analogous manner. There is no high-slope weight bin that falls in either the 1st quartile or the 3rd quartile of stringency levels. Hence, for cleaner results, we drop vehicle models that fall in the 1st and the 3rd stringency quartiles. This also eliminates bins that are too narrow to have any product offerings (i.e., rows 17 and 23). Once we remove these bins, we have substantial variations in both slope and stringency measures 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. 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 another reason why we think our reduced-form approach to identify only the TPFs is more viable than a structural approach.

4.D. Descriptive Evidence

Before moving to our main analysis, we take a glance at graphical evidence. We first make use of weight-bin-level variations in regulatory assignment. **Panel A-(a) of Figure 6** displays an unconditional scatter plot of logged fuel-economy ratings against logged vehicle weights for vehicle grades introduced before the 2007 standards. 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 stringency levels during the pre-2007 period. Variants of vehicle models assigned to the high-slope weight bins are marked with circle; those assigned to the low-slope bins are marked with \times [Here, the definition of high- versus low-slope follows **column (5) in Table 2**]. The figure indicates no sign of a significant difference in the technical trade-offs between fuel economy and weight prior to the new standards.

Panel A-(b) of Figure 6 repeats the same 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 some difference in technical trade-offs between the two groups. For better visibility, we condition out the influence of vehicle attributes other than those that directly relate to vehicle weight. The figures in **Panel B** essentially are the same as those in **Panel A**, except that **Panel B** plots the residuals from a regression of logged fuel economy on key vehicle attributes (in logged values) after removing the linear projection from terms involving horsepower, torque, transmission, and brand dummies. We now see the TPF of those assigned to the high-slope bins lie *far below* the TPF of those assigned to the low-slope bins after the new standards despite the fact that the former lie *slightly above* the latter before the new standards.

Next, we turn to the model-level variations in regulatory assignment. We first compute the (unweighted) means of fuel-economy ratings (over vehicle variants) for each model, for each of the pre- and the post-2007 periods. We then calculate the changes in these means between the two periods. **Figure 7** then plots these changes against the regulatory slopes that account for initial positioning at the model level [per equation (3)] for different levels of regulatory stringency. The figure indeed shows the patterns consistent with our economic predictions. Many of the models improved (average) fuel-economy ratings after the new standards. And these improvements are indeed greater for those faced with more demanding fuel-economy targets (relative to their initial positions). Yet, the improvements seem to decline, and turn even negative in some cases, with the increase in regulatory slope.

To offer support for the common-trend assumptions, 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 differences in the mean fuel economy ratings between the high-stringency and the low-stringency groups by year by treatment in **Figure 8**. **Figure 8-(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 8-(b)** demonstrates that the differences in average fuel economy between the high- and the low-stringency 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 old 2001 standards expired in 2009.

5. DDD Regression Estimates

5.A. Main Results

Table 3 reports the results of four regression models for each of the two alternative measures of regulatory assignment. **Panel A** displays the results using the bin-level regulatory variations in Γ and Λ for constructing our treatment variables T and H . The first model (in columns 1 and 2) in this panel estimates DD regressions on the pooled sample, with high-slope bins against low-slope bins as the primary treatment. The estimates from these regressions would be biased downward if vehicles assigned to the low-stringency 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 third models estimate the same regressions, but on subsamples consisting only of those of high-cost bins and of low-cost bins, respectively. The last model estimates full DDD regressions on the pooled sample. Each of these models is estimated with or without segment dummies interacted with time dummies. All specifications include weight (w), horsepower (hp), size ($size$), torque (tq) (all in logged values) and AT/CVT dummy as well as brand dummies interacted with time dummies. **Panel B** essentially repeats the same, but using the model-level regulatory variations in Γ and Λ for T and H .

We first discuss **Panel A**. The DD estimates of the impact of the high slope on the pooled sample are negative and statistically significant. The magnitude of the estimates gets much

larger when the same regressions are run on a subsample consisting only of high-stringency weight bins. In contrast, the DD estimates turn either significantly positive or insignificant on a subsample consisting only of those assigned to the low-cost bins. These results are consistent with our expectation, and are indeed suggestive of the success in our empirical strategies. Firms have a greater incentive to exploit regulatory loopholes when faced with more stringent standards. Per our theory presented in **Section 2**, this incentive result in a lower rate of progress in fuel-economy technology. However, the fact that the DD estimate on the low-stringency subsample is positive if time-varying segment controls are not included, but turns negative (and insignificant) once these controls are included implies that the rates of technical progress do vary across segments, irrespective of regulatory assignment, and are indeed higher for vehicles assigned to the low-stringency bins. This in turn suggests that vehicles assigned to the high-stringency bins might have been those with a lower rate of technical progress, and hence, if uncontrolled, this might confound the DD estimates since the regulatory slope correlates with the regulatory stringency. Hence, this gives support for our DDD strategy. The DDD estimate is indeed negative, highly statistically significant, and qualitatively very large: With the Cobb-Douglas specification, the estimate implies that the assignment to high-slope weight bins slows down fuel-economy improvements by roughly 13-19%. 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.¹⁵ 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) or (d)**. 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.

Next, we turn to **Panel B**. The results here are quantitatively very similar to those in **Panel A**, but differ qualitatively on one important account. Recall that our regulatory variables T and H in this panel incorporate variations in vehicle models' initial positioning (prior to the 2007 standards) relative to the 2007 standards. Hence, there remains substantial variation in regulatory stringency H across vehicles within a subsample consisting only of either the high-stringency or the low-stringency bins. Therefore, if the Ratchet-type effect indeed exists and arises due to this model-level regulatory stringency, the DD regression

¹⁵We also estimated the same regressions using translog specifications. The results are virtually identical. The translog specifications improve the fit by a small margin, but they also seems over-parameterized as in Knittel (2011). Since the Cobb-Douglas specifications are already well fit, we do not report the results here.

estimates would be biased on all samples (i.e., in all of columns 1-6). On each subsample, the bin-level regulatory stringency is roughly controlled, and hence, much of the remaining variation arises from the variation in vehicle models’ initial positioning. Since our slope measure is a direct function of this model-level regulatory stringency, vehicles faced with steeper (model-level) slopes may be simply those that had initially low fuel-economy ratings. The DD estimates on these subsamples, therefore, may be simply picking up the effect of having initially low fuel-economy ratings relative to the new standards. Though the direction of the bias is hard to predict a priori, our DD results seem to suggest a plausible direction. Our DD estimates are negative (and statistically significant) on the subsample consisting of the high-stringency bins whereas they are positive (and statistically significant) on the low-stringency bins. We may interpret this result as follows. An additional increase in regulatory stringency (at the model level) induces the Ratchet-type effect and slows down the rate of technical progress only when the regulatory stringency is already very high. The Ratchet-type incentive disappears, however, when the stringency level is low. Therefore, an additional increase in stringency simply leads to a higher rate of technical progress. Our DDD strategy helps us control for this effect, giving us an unbiased estimate of the slope effect. The DDD estimate on the pooled sample is indeed negative and highly statistically significant. The magnitude is also large — a one-unit increase in the regulatory slope slows down fuel-economy improvements by roughly 17-28 ppt.

5.B. 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 mechanism caused that effect. The economic mechanism outlined in **Section 2** is that weight bins that have steeper slopes relative to the pre-existing TPFs would induce firm to increase vehicle weights. 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 intricate 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 stringency \times slope subsamples. Hence, we run DD and DDD regressions of logged curb weight on the same set of covariates as in **Table 3** (excluding logged weight, of course). **Table 4** reports the results of these regressions, for each of the two alternative measures of regulatory assignment as in **Table 3**.

In **Panel A**, the DD estimate is positive and statistically highly significant on the high-stringency weight bins, but is not significant on the low-stringency bins. These are consistent with our results on fuel-economy ratings. The DD estimate on the pooled sample averages out these two, and hence, is positive but statistically insignificant. On the other hand, the DDD estimate gets at the differences between the two, and hence, is positive and statistically highly significant. These results seem to confirm that high regulatory slopes do create incentives to increase curb weight, particularly in high-stringency bins. **Panel B** essentially confirms the same point. Recall that all DD estimates in **Panel B** are likely to give us biased estimates for essentially the same reason discussed in the previous subsection — the DD estimates may simply pick up the effect of having initially low fuel-economy ratings relative to the new standards rather than the effect of assignment to high slopes. Hence, we focus on the DDD estimate. There, we again see a large and statistically significant, positive effect of high-slope assignment on curb weight.

5.C. Placebo Checks

We run two placebo checks to verify our regression results. For ease of interpretation, these placebo experiments perturb only the bin-level regulatory assignments. Hence, our placebo results are directly comparable to those reported in **Panel A of Table 3**. First, we arbitrarily perturb our weight-bin assignment and see if our results continue to hold.

Specifically, we shift weight cutoffs w_b in (2) by an arbitrary number k (in kg) and run the same regressions as before. The parameter estimates from this fictitious assignment should qualitatively differ from those of the factual weight-bin assignment. This placebo assignment, however, may not simply result in the disappearance of the statistical significance since we expect some influence of regulatory assignment in almost every weight bin. **Panel A of Table 5** below reports the DD and DDD estimates when $k = 25$ and all the same covariates as in **Table 3** (incl. brand and segment dummies interacted with time dummies) are used.¹⁶ The DD estimates on the pooled sample and on the high-stringency subsample (columns 1 and 2) are statistically insignificant. Furthermore, the DD estimate on the low-stringency subsample turns negative and statistically highly significant. We take these as a support for our main results. The placebo experiments make the distortionary effect of the high-slope assignment go away on the samples where we expect it to be large, while making it stronger on the sample where we expect it to be small. Consequently, the DDD estimate is negative and statistically significant. The estimate, however, results from essentially taking the difference between these two delusive estimates, and therefore, should not be taken as an objection to our main results.

Next, we perturb on temporal dimension, holding weight bin assignment. In **Panel B of Table 5**, we report the DD and DDD estimates, using 2003-2004 (instead of 2004-2006) as the control period and 2005-2006 (instead of 2010-2012) as the fictitious treatment period (again, all the other covariates stay the same as in **Panel A of Table 3**). Because the new standards are adopted in July 2007, the estimates on this placebo treatment should not capture the effects of differential regulatory treatments due to the new standards. This placebo exercise should, instead, capture the pre-trends across different weight-bin assignments, and therefore, also serve as the check for the common-trend assumptions. The DD estimates on both the high-stringency and the low-stringency subsamples are positive and statistically significant at conventional levels. Accordingly, the DD estimate on the pooled sample is also positive, though not statistically significant. On one hand, these results support our main results — the fictitious treatment gives us the results that are qualitatively much different from the main results. On the other hand, these results also imply that those assigned to high-slope bins are associated with higher rates of technical progress during the pre-2007 period. If the same trends continue to the actual treatment period (2010-2012), our DD estimates would be biased toward zero and the true impacts would be larger. In contrast, the DDD

¹⁶Note 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.

estimate on this fictitious treatment is negative and statistically significant, implying that the pre-trends between car models assigned to different regulatory slopes do indeed differ across car models assigned to different stringency levels. This implies that our DDD estimate might be biased upward if the pre-trends extend to the actual treatment period, although the magnitude of the bias is quite small (-0.045). These reasonings suggest that the true impact might fall somewhere between the DD estimate (-0.047) and the DDD estimate (-0.139). The difficulty, however, is that if firms indeed engage in the Ratchet-type behavior, they might intentionally change the rates of technical progress between these periods, and therefore, the true trends during the post-2007 period could be quite different from these pre-trends. The DDD regression exploiting the model-level variations might be more robust as it control for the Ratchet behavior more directly.

6. Welfare Implications

In this section, we discuss two important implications of our empirical findings for policy and welfare evaluation. These implications also help us substantiate the important differences between ours and 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)].

First, 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. The size of the losses may not be necessarily large, however, if we do not account for their impacts on technical change.¹⁷ 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 in the absence of an efficient credit trading mechanism. All these studies, however, assume the distortion shows up in firm's or consumer's *second-stage choice conditional on firm's first-stage 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

¹⁷This 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.

same principle is also applied in Ito and Sallee (forthcoming), who estimate the discrete choice model of firms' new product offerings under the assumption of perfect competition. In their model, there is no distinction between the first-stage choice of technology and the second-stage choice of fine-tuning product attributes, and 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, any distance from the existing offerings is counted as the cost to the firms. Therefore, their welfare estimates ignores the gain in consumer surplus resulting from technical change. In contrast, the recent IO literature substantiates the importance of accounting for the gain in consumer surplus from new or improved products [e.g., Goolsbee and Petrin (2004), Nevo (2003), and Petrin (2002)]. Failure to account for consumer loss from the distortion in technical change, therefore, is likely to understate the welfare effects of regulatory loopholes, which may lead to misguided policy advice. To demonstrate this point, let us calculate a naive estimate of the welfare loss from the distortion in technical change. We simply take the (marginal) value of fuel-economy ratings estimated in Konishi and Zhao (2017), and estimates the consumer gain from a 20-ppt increase in fuel-economy ratings for vehicle models that are assigned to high-slope bins. This naive estimate indicates that the welfare loss is \$23,183 per vehicle model in 2012 alone. In contrast, Ito and Sallee estimate that the welfare gain due to the Japanese fuel-economy regulation (relative to a flat standard) is \$694-\$1,167 per vehicle model. Therefore, accounting for the welfare loss from distortion in technical change could potentially overturn the policy evaluation.

Second, by design, fuel-economy standards are intended to address the externality arising from the consumption of fossil fuel in second-best settings. The textbook theory suggests that in first-best settings, a gasoline tax can fully restore economic efficiency, correcting the externality from both car utilization and car ownership. Yet, real-world complexities make the gasoline tax a less ideal policy instrument. Examples of such complexities include consumer heterogeneity, imperfect competition, and possible optimization failures in consumer choice. There is a growing interest among economists investigating the significance of these real-world complexities for environmental regulation, recognizing the fact that the governments around the world often rely on complementary instruments such as feebates and fuel-economy standards, either in addition to or in place of the gasoline tax (Anderson *et al.*, 2011). For example, Konishi and Zhao (2017) demonstrate, in the real empirical context related to ours, that a simple gasoline tax can never achieve the social optimum in the presence of imperfect market competition, and that the second-best policy must undertake a daunting task of coordinating tax incentives across products, accounting for product-level

consumer demand and imperfect competition in the market.¹⁸ A question arises naturally in this regard. In first-best settings, feebates and fuel-economy standards with efficient credit trading are equivalent policy instruments. The question then is in second-best settings, what *regulatory* conditions would constitute the equivalence between the two policy instruments? In this context, Ito and Sallee correctly point out the importance of striking a balance between marginal cost harmonization versus bias minimization in firm’s attribute choice. They also show that no attribute-biasing is optimal in the presence of an efficient credit trading. This result is likely to change, however, once we account for our findings: the flat standard is too flat relative to the slope of firm’s technical possibility frontiers, and hence, is likely to bias firm’s technology/product choice. This bias still exists even when an efficient market for fuel-economy credits is in place because firms can change the amount of initial credits by changing the level of technical upgrade or by changing the product attributes — violation of the condition for the independence of credit-trading equilibrium from initial allocation of credits. It is largely an empirical question to what extent imperfect competition in the market either reduce or increase the inefficiency loss from such a bias. Future research might take in recent advances in the empirical IO literature, fully endogeneizing firm’s technology/product choice, to empirically investigate this effect.

7. Concluding Remarks

We examine the distortionary impact of regulatory loopholes on technical progress in the well-defined context of the Japanese weight-based fuel economy regulation. 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 13-17 ppt. Hence, we conclude that the attribute-based regulation has significantly distorted technological change in the Japanese automobile industry.

¹⁸Konishi and Zhao provide a detailed discussion of this point in their online appendix.

The findings of the paper deliver four important messages. First, it is not just the notched schedule but rather the regulatory attribute trade-offs 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 bias in reporting to the government rather than bias 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] 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.
- [4] 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.
- [5] Anderson, Soren T., Ian W. H. Parry, James M. Sallee, and Carolyn Fischer 2011. Automobile Fuel Economy Standards: Impacts, Efficiency, and Alternatives. *Review of Environmental Economics and Policy* 5 (1): 89-108.

- [6] Austin, D. and T. Dinan. (2005). Clearing the Air: The Costs and Consequences of Higher CAFE Standards and Increased Gasoline Taxes. *Journal of Environmental Economics and Management* 50: 562–582
- [7] Berry, Steven, James Levinsohn, and Ariel Pakes. 1995 Automobile Prices in Market Equilibrium. *Econometrica* 63 (4): 841-890.
- [8] 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.
- [9] Copeland, Brian R. and M. Scott Taylor. (1994). North-South Trade and the Environment. *The Quarterly Journal of Economics* 109 (3): 755-787.
- [10] Fan, Y. (2013). Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market. *American Economic Review* 103 (5): 1598-1628.
- [11] Freixas, X., Guesnerie, R., & Tirole, J. (1985). Planning under Incomplete Information and the Ratchet Effect. *The Review of Economic Studies* 52 (2), 173-191.
- [12] Fullerton, Don and Sarah E. West. 2002. Can Taxes on Cars and Gasoline Mimic an Unavailable Tax on Emissions? *Journal of Environmental Economics and Management* 43: 135-157.
- [13] Goldberg, Pinelopi Koujianou. 1998. The Effects of the Corporate Average Fuel Efficiency Standards in the US. *Journal of Industrial Economics* 46 (1): 1-33.
- [14] Goolsbee, A. and A. Petrin. (2004). The Consumer Gains from Direct Broadcast Satellites and the Competition with Cable TV. 72 (2): 351-381.
- [15] Gruber, J. (1994) The Incidence of Mandated Maternity Benefits. *American Economic Review* 84(3): 622-641.
- [16] Hitsch, G. J. (2006). An Empirical Model of Optimal Dynamic Product Launch and Exit Under Demand Uncertainty. *Marketing Science* 25 (1): 25–50.
- [17] 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*.

- [18] Jacobsen, Mark R. (2013). Evaluating US Fuel Economy Standards in a Model with Producer and Household Heterogeneity. *American Economic Journal: Economic Policy* 5 (2): 148-87.
- [19] Jaffe, A.B., Newell, R.G. and Stavins, R.N. (2002). Environmental Policy and Technological Change. *Environmental and Resource Economics* 22: 41-70.
- [20] 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.
- [21] 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.
- [22] Laffont, Jean-Jacques and Tirole, Jean. (1988). The Dynamics of Incentive Contracts. *Econometrica* 56 (5): 1153-75.
- [23] Nevo, A. (2003). New Products, Quality Changes and Welfare Measures Computed from Estimated Demand Systems. *The Review of Economics and Statistics*. 85 (2): 266-275.
- [24] 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.
- [25] Petrin, A. (2002). Quantifying the Benefits of New Products: The Case of the Minivan. 110 (4): 705-729
- [26] Popp, D. (2002). Induced Innovation and Energy Prices. *The American Economic Review* 92 (1): 160–180.
- [27] Porter, M. (1991). America’s Green Strategy. *Scientific American* 264(4), 168.
- [28] 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,
- [29] Seim, K. (2006). An Empirical Model of Firm Entry with Endogenous Product-type Choices. *RAND Journal of Economics* 37 (3): 619-640.
- [30] Wollmann, T. *Forthcoming*. Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicle. *American Economic Review*.

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


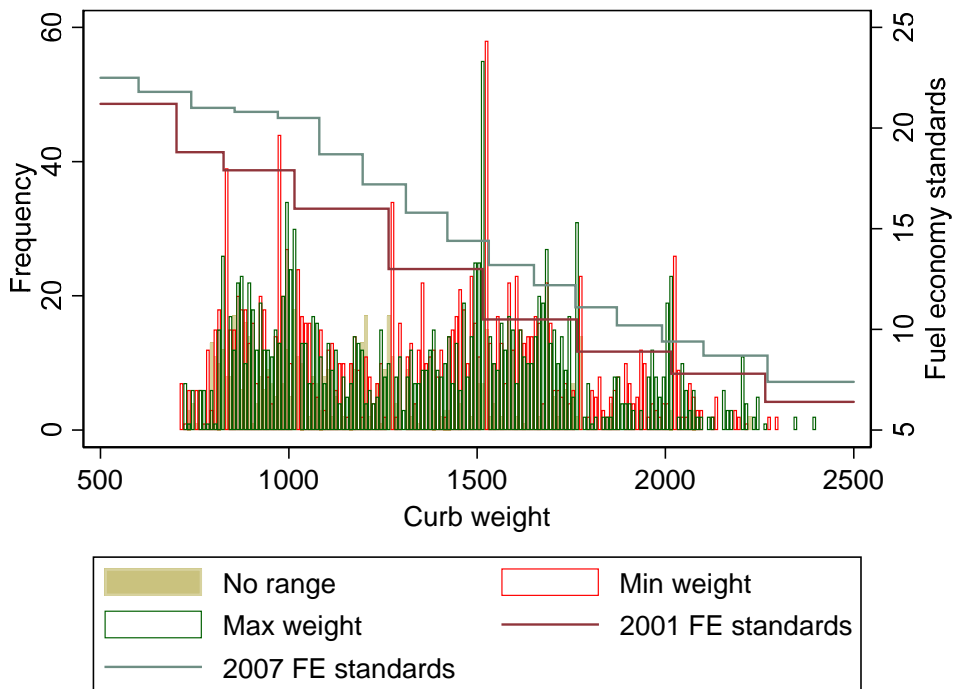


Figure 4. Vehicle Weight Distributions, Years 2010-2012

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



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

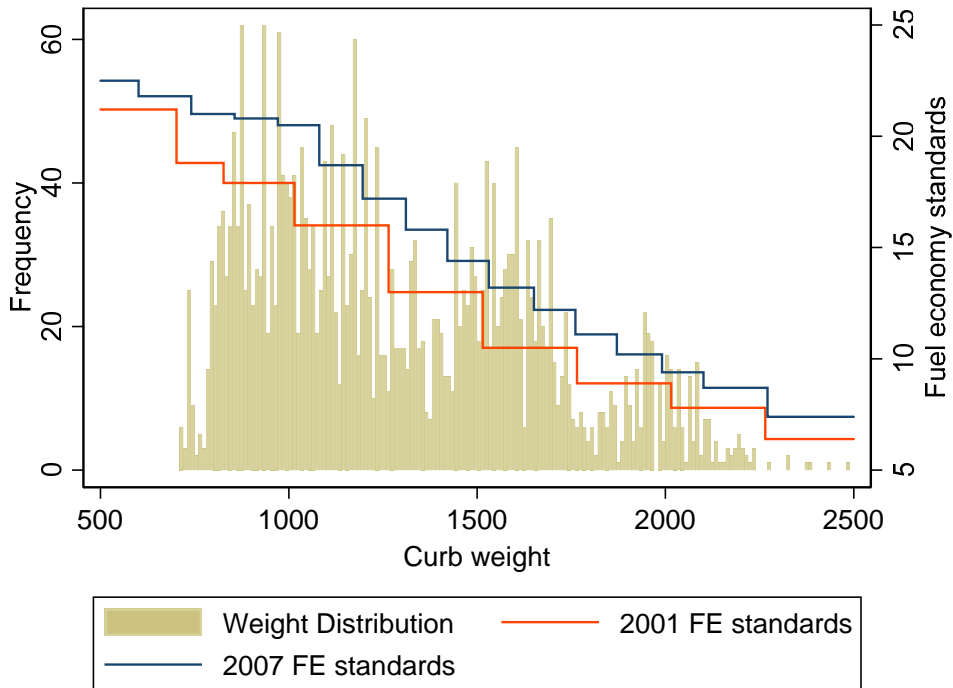


Figure 5. Variation in Regulatory Assignments: An Illustration



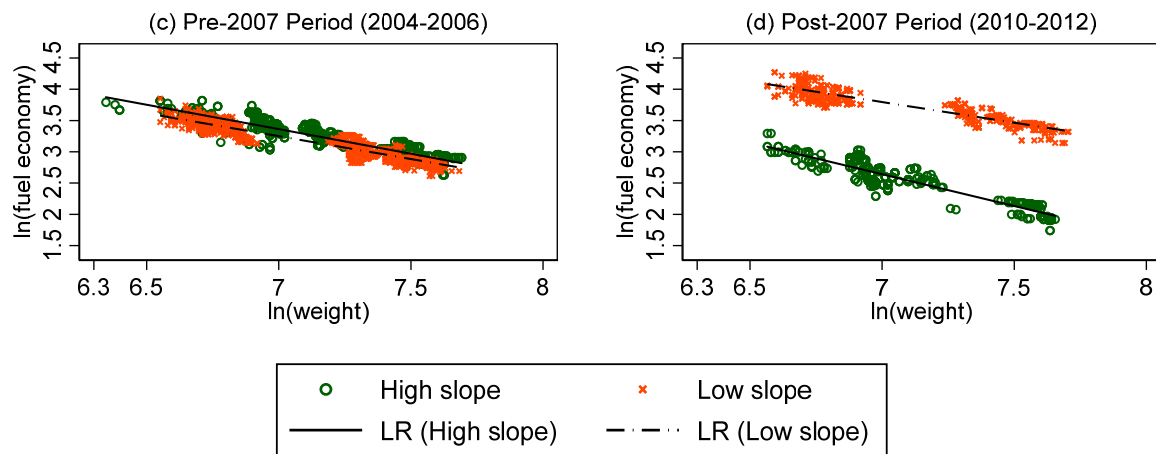


Figure 6. Technology Possibility Frontiers Before and After the New Standards

A. Unconditional TPFs




B. Conditional TPFs



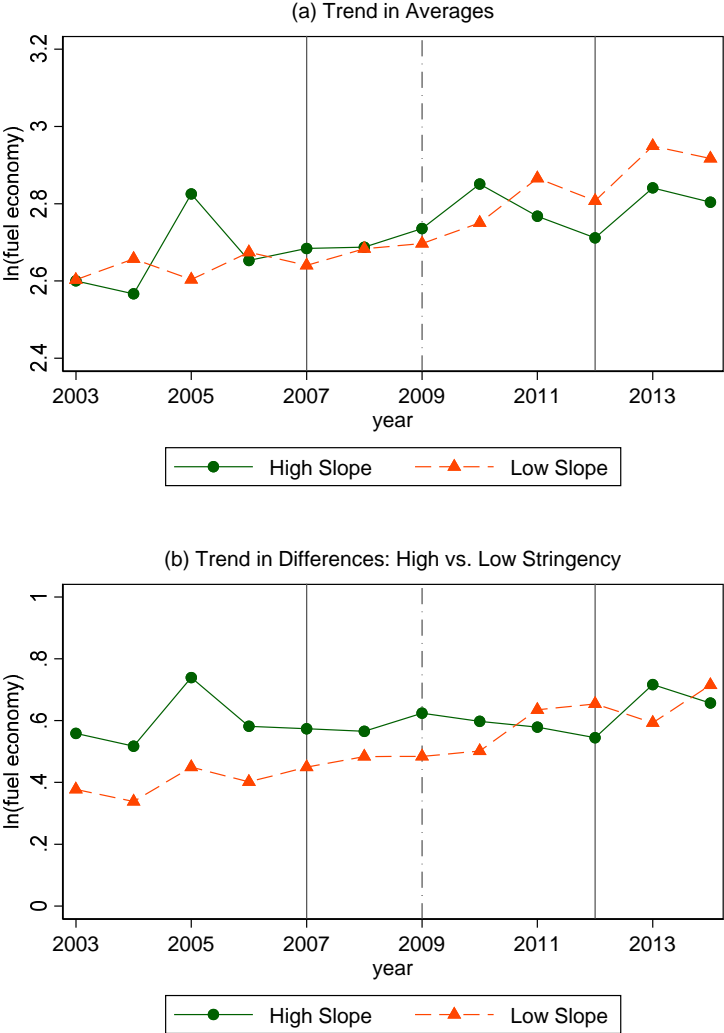
Note: The figure excludes commercial vans and trucks, imported brands, diesel, hybrid, and electric cars as well as observations that fall in weight segments with the first and the third quartiles of compliance costs during the pre-2007 period. Panel A displays scatter plots using raw data. Panel B displays scatter plots using the residuals from a regression of logged fuel economy on key vehicle attributes (in logged values) after removing the linear projection from terms involving horsepower, torque, transmission, and brand dummies.

Figure 7. Changes in Fuel-economy Ratings by Regulatory Slope and Stringency



Note: The changes in fuel-economy ratings are calculated as the difference in the model-level means between the pre-2007 and the post-2007 periods. Slope and stringency measures are calculated per equations (3) and (5), respectively.

Figure 8. Trends in Average Fuel Economy between and within Groups



Note: Panel (a) plots average fuel economy ratings in logged values for the high-slope and the low-slope groups. Panel (b) plots the differences in average fuel economy ratings between the high-cost and the low-cost groups for the high-slope and the low-slope groups.

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	Slope of Regulation				Stringency of Regulation				Pre-2007 (2004-06)		Post-2007 (2010-12)				
	H22 (10.15M)	H27 (JC08M)	Γ	Larger Than Seg. TPF?	Γ^*	Λ	Quartile of Δ FE	Λ^*	N	Mean	S.D.	N	Mean	S.D.	
	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Fuel Econ. (km/L)	Fuel Econ. (km/L)	Fuel Econ. (km/L)	Fuel Econ. (km/L)	Mean	S.D.	
1 0 - 600	21.2	22.5	0.12	1	7.95	0.00	1.3	2	1.10	0.00	5	23.40	1.95	0	--
2 600 - 702	21.2	21.8	0.00	0	3.99	2.77	0.6	1	1.47	0.47	8	25.88	2.34	0	--
3 702 - 740	18.8	21.8	2.11	1	1.74	2.43	3.0	4	1.37	0.50	42	22.12	1.67	43	27.82
4 740 - 827	18.8	21	0.00	0	3.97	14.18	2.2	2	2.67	1.48	255	20.75	1.94	146	26.12
5 827 - 856	17.9	21	0.69	1	6.01	18.19	3.1	4	3.35	1.47	150	19.29	1.82	108	23.66
6 856 - 970	17.9	20.8	0.26	0	4.12	14.97	2.9	4	3.94	1.47	438	18.43	1.70	435	22.80
7 970 - 1015	17.9	20.5	0.00	0	1.34	1.09	2.6	3	5.10	2.14	130	17.88	3.23	160	20.74
8 1015 - 1080	16	20.5	2.77	1	1.60	2.36	4.5	4	4.25	1.85	202	17.20	1.70	202	19.37
9 1080 - 1195	16	18.7	1.30	0	1.71	2.19	2.7	3	3.94	1.56	488	16.35	1.65	318	18.72
10 1195 - 1265	16	17.2	0.00	0	1.54	1.50	1.2	1	3.42	1.07	310	15.83	1.58	161	17.49
11 1265 - 1310	13	17.2	3.11	1	1.97	2.68	4.2	4	3.66	1.43	131	13.50	1.82	91	16.20
12 1310 - 1420	13	15.8	1.27	1	1.89	1.93	2.8	3	3.55	1.44	301	12.97	1.78	162	15.26
13 1420 - 1515	13	14.4	0.00	0	1.34	1.37	1.4	2	2.98	1.45	339	12.49	1.48	197	13.76
14 1515 - 1530	10.5	14.4	8.00	1	1.05	0.97	3.9	4	2.99	1.20	75	11.47	1.15	57	12.91
15 1530 - 1650	10.5	13.2	0.83	0	1.01	0.76	2.7	3	2.58	1.01	419	11.29	1.08	295	12.66
16 1650 - 1760	10.5	12.2	1.00	0	2.10	5.17	1.7	2	2.58	0.97	281	10.32	1.20	174	12.24
17 1760 - 1765	10.5	11.1	0.00	0			0.6	1			0	--	--	0	--
18 1765 - 1870	8.9	11.1	0.86	1	3.04	7.19	2.2	2	2.07	0.79	193	9.50	0.87	48	10.39
19 1870 - 1990	8.9	10.2	0.67	0	0.91	0.62	1.3	2	1.60	0.57	141	9.16	0.59	113	10.50
20 1990 - 2015	8.9	9.4	0.00	0	1.00	0.53	0.5	1	1.72	0.31	16	8.57	0.26	28	10.22
21 2015 - 2100	7.8	9.4	0.82	1	1.26	0.69	1.6	2	1.65	0.37	32	8.24	0.40	59	9.04
22 2100 - 2265	7.8	8.7	0.00	0	1.45	0.39	0.9	1	1.59	0.13	14	8.10	0.15	14	8.45
23 2265 - 2270	6.4	8.7	26.00	1			2.3	3			0	--	--	0	--
24 2270 - 3500	6.4	7.4	0.00	0	0.00	0.00	1.0	1	1.19	0.30	11	6.41	0.30	4	7.00

Table 3. DDD Regression Results using 2004-2012 Passenger Cars

	DD (Pooled)		DD (Stringency: High)		DD (Stringency: Low)		DDD (Pooled)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel A: Bin-level Assignment</i>								
<i>T</i> [Slope: High = 1]	0.035 *** (0.005)	-0.021 *** (0.007)	0.036 *** (0.008)	0.040 *** (0.013)	-0.001 (0.007)	0.020 ** (0.009)	-0.002 (0.008)	0.014 (0.011)
<i>R</i> [Post-2007 = 1]	0.005 (0.032)	-0.047 (0.035)	0.206 *** (0.025)	0.558 *** (0.046)	-0.002 (0.026)	0.178 *** (0.035)	0.004 (0.031)	-0.046 (0.041)
<i>R</i> × <i>T</i>	-0.035 *** (0.009)	-0.047 *** (0.011)	-0.115 *** (0.013)	-0.239 *** (0.023)	0.082 *** (0.014)	-0.002 (0.020)	0.089 *** (0.016)	0.019 (0.024)
<i>H</i> [Stringency: High = 1]							0.001 (0.007)	-0.011 (0.013)
<i>H</i> × <i>T</i>							0.052 *** (0.010)	-0.043 *** (0.014)
<i>H</i> × <i>R</i>							0.088 *** (0.011)	0.123 *** (0.028)
<i>H</i> × <i>R</i> × <i>T</i>							-0.210 *** (0.019)	-0.139 *** (0.031)
ln(weight)	-1.354 *** (0.033)	-1.387 *** (0.041)	-1.429 *** (0.053)	-1.761 *** (0.058)	-0.979 *** (0.047)	-0.748 *** (0.053)	-1.350 *** (0.034)	-1.386 *** (0.041)
Variant-level controls for observables	✓	✓	✓	✓	✓	✓	✓	✓
Time-varying Brand-effect controls	✓	✓	✓	✓	✓	✓	✓	✓
Time-varying Segment-effect controls		✓		✓		✓		✓
R ²	0.917	0.930	0.699	0.745	0.933	0.949	0.921	0.932
Obs.	3,253	3,253	1,516	1,516	1,737	1,737	3,253	3,253
<i>Panel B: Model-level Assignment</i>								
<i>T</i> [Slope]	-0.098 *** (0.025)	-0.083 *** (0.024)	0.005 (0.040)	0.010 (0.038)	0.153 ** (0.066)	0.266 *** (0.060)	0.780 *** (0.146)	0.433 *** (0.127)
<i>R</i> [Post-2007 = 1]	0.012 (0.032)	-0.041 (0.035)	0.162 *** (0.025)	0.343 *** (0.041)	-0.099 *** (0.030)	0.067 *** (0.018)	-0.043 (0.031)	-0.098 *** (0.031)
<i>R</i> × <i>T</i>	-0.200 *** (0.054)	-0.187 *** (0.053)	-0.443 *** (0.074)	-0.402 *** (0.070)	4.699 *** (0.727)	2.905 *** (0.763)	1.414 *** (0.377)	0.710 ** (0.330)
<i>H</i> [Stringency]							-0.021 *** (0.002)	-0.044 *** (0.002)
<i>H</i> × <i>T</i>							-0.184 *** (0.031)	-0.098 *** (0.028)
<i>H</i> × <i>R</i>							0.014 *** (0.003)	0.012 *** (0.003)
<i>H</i> × <i>R</i> × <i>T</i>							-0.330 *** (0.079)	-0.187 *** (0.069)
ln(weight)	-1.363 *** (0.033)	-1.324 *** (0.039)	-1.430 *** (0.054)	-1.734 *** (0.059)	-0.946 *** (0.045)	-0.815 *** (0.050)	-1.340 *** (0.031)	-1.032 *** (0.034)
Variant-level controls for observables	✓	✓	✓	✓	✓	✓	✓	✓
Time-varying Brand-effect controls	✓	✓	✓	✓	✓	✓	✓	✓
Time-varying Segment-effect controls		✓		✓		✓		✓
R ²	0.917	0.930	0.694	0.733	0.933	0.950	0.927	0.951
Obs.	3,247	3,247	1,516	1,516	1,731	1,731	3,247	3,247

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 4. Regression Results on Vehicle Weight

	DD (Pooled)	DD (Stringency: High)	DD (Stringency: Low)	DDD (Pooled)
	(1)	(2)	(3)	(4)
<i>Panel A: Bin-level Assignment</i>				
DD or DDD Estimate	0.002 (0.005)	0.052 *** (0.016)	-0.004 (0.008)	0.038 ** (0.017)
R ²	0.989	0.975	0.975	0.990
Obs.	2,014	658	1,356	2,014
<i>Panel B: Model-level Assignment</i>				
DD or DDD Estimate	0.044 * (0.023)	0.075 *** (0.024)	0.818 ** (0.350)	0.118 ** (0.058)
R ²	0.988	0.975	0.973	0.988
Obs.	2,008	658	1,350	2,008

Note: All regressions use a subsample consisting only of bins with width less than 40 (in kg). In all regressions, logged curb weight is regressed on the same set of covariates as in Table 3, excluding logged weight. 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 (Stringency: High)	DD (Stringency: Low)	DDD (Pooled)
	(1)	(2)	(3)	(4)
<i>Panel A. Perturbing Bin Assignments</i>				
DD or DDD Estimate	-0.015 (0.013)	0.014 (0.019)	-0.149 *** (0.029)	-0.119 *** (0.033)
R ²	0.824	0.728	0.969	0.826
Obs.	2,006	1,130	437	2,006
<i>Panel B. Perturbing Treatment Periods</i>				
DD or DDD Estimate	0.014 (0.009)	0.036 * (0.020)	0.049 *** (0.017)	-0.045 * (0.027)
R ²	0.938	0.737	0.953	0.940
Obs.	2,827	1,183	1,644	2,827

Note: Panel A reports the results of regressions on a placebo treatment where weight cutoffs are shifted by $k = 25$ kg. Panel B reports on another placebo treatment where the control period is 2003-2004 and the treatment period is 2005-2006. All regressions use the same covariates as in Table 3 including time-varying brand and segment effects. In parentheses are standard errors. The asterisks *, **, and *** indicate significance at 0.1, 0.05, and 0.01 levels.

Online Appendix

Appendix A. 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

$$\sigma = \rho + \alpha, \quad \text{for some } \alpha \in \mathbb{R}, \tag{1}$$

where $\sigma = -dR/dw$ is the slope of the regulatory constraint and $\rho = -\partial T/\partial w$ is the slope of the TPF. Using (1), 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. ■

Appendix B. Firm-level Compliance as of 2007

We control for firm-level heterogeneity in technical progress by including maker-fixed effects interacted with the time-period dummy in our DDD regression. Yet, if firms differ substantially in compliance level prior to the new standards, they may respond to regulatory treatment quite differently. Since we construct two-fold control-treatment pairs within a narrow weight segment, we do not have sufficient variation to allow for interaction effects to capture this response heterogeneity. We, therefore, check whether firms' compliance levels differed substantially at the beginning of the new standards.

In **Figure A1**, the red bar displays the sales-weighted average fuel economy of vehicles sold in 2007 for each domestic car maker. The green shaded bar reports 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 A1. Sales-weighted Fuel Economy and Standards by Maker in 2007



Appendix C. MLIT Data

Figure A2 is a raw image of the original table reported in the Ministry of Land, Infrastructure, and Transportation fuel-economy data. The table captions are in Japanese, so we highlight the relevant section in red. As shown, vehicle weights are reported in range for a majority of entries.

Figure A2. MLIT Data Image

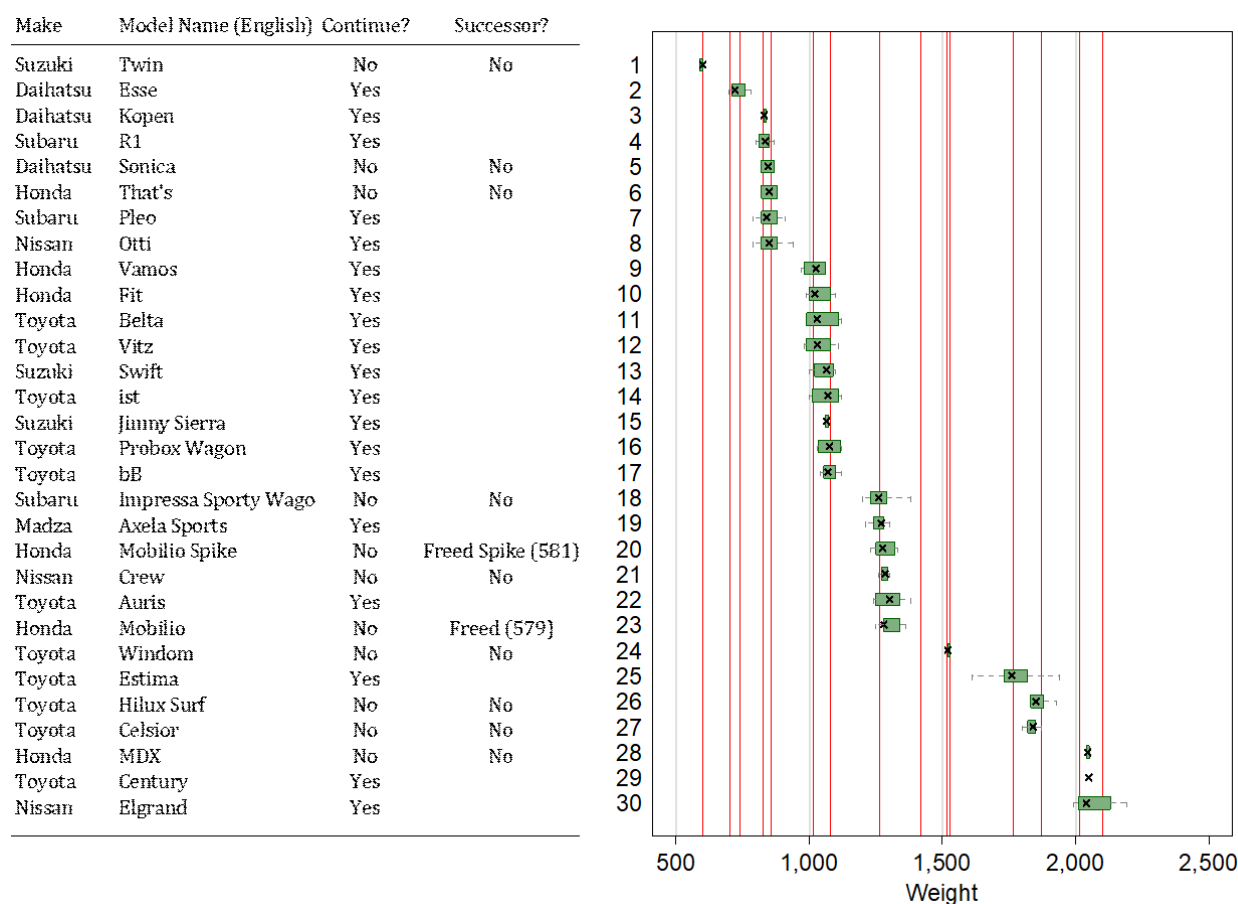
自動車燃費の算出方法及び輸入の事業を行う者の届出又は名称 日本建設工業株式会社 月間燃費(平成22年度)

車名	通称名	原動機		型式	総排気量(L)	変速機の種類及び構造	車重(kg)	燃費定率(%)	10・15E-F			主要燃費効率	その他の燃費効率			燃費効率	燃費効率
		型式	新式						燃費定率(%)	燃費定率(%)	燃費定率(%)		燃費定率(%)	燃費定率(%)	燃費定率(%)		
ホンダ	インサイト	DBA-ZE2	LEA(内排気筒)MF(電動機)	CVT(E)	1339	CVT(E)	1190	5	31.0	75	16.0	C-VT-EP-C-H	3W+EGR	F	タイプ 185/65R15	燃費効率	125
		DBA-ZE2	LEA(内排気筒)MF(電動機)	CVT(E)	1339	CVT(E)	1190	5	30.0	77	16.0	C-VT-EP-C-H	3W+EGR	F		燃費効率	125
		DBA-ZE2	LEA(内排気筒)MF(電動機)	CVT(E)	1339	CVT(E)	1200	5	28.0	83	16.0	C-VT-EP-C-H	3W+EGR	F	タイプ 185/65R16	燃費効率	125
	インサイト エクスターループ	DBA-ZE2	LEA(内排気筒)MF(電動機)	CVT(E)	1496	CVT(E)	1200~1210	5	26.5	88	16.0	V-TEC-EP-C-H	3W+EGR	F		燃費効率	125
		DBA-ZE3	LEA(内排気筒)MF(電動機)	CVT(E)	1496	CVT(E)	1210	5				3W+EGR	F	タイプ 185/65R16	燃費効率	125	
		DBA-GF1	LEA(内排気筒)MF(電動機)	CVT(E)	1339	CVT(E)	1150~1165	5				3W+EGR	F		燃費効率	125	
	フィット	DBA-GE6	L13A	CVT(E+LTC)	1339	CVT(E+LTC)	1010	5	24.5	95	17.9	C-V-EP	3W+EGR	F	CVTパワーマ-	燃費効率	125
			L13A	CVT(E+LTC)	1339	CVT(E+LTC)	1010	5	24.0	97	17.9	C-V-EP	3W+EGR	F		燃費効率	125
			L13A	CVT(E+LTC)	1339	CVT(E+LTC)	1030~1030	5	22.0	106	16.0	C-V-EP	3W+EGR	F	CVTパワーマ-	燃費効率	125
		DBA-GE6	L13A	SMT	1339	SMT	990	5	21.5	108	16.0	C-V-EP	3W+EGR	F		燃費効率	125
			L13A	SMT	1339	SMT	990	5	21.0	111	17.9	V-EP	3W+EGR	F		燃費効率	115
			L13A	SMT	1339	SMT	990~1010	5	21.0	111	17.9	V-EP	3W+EGR	F		燃費効率	115
DBA-GE6		L15A	CVT(E+LTC)	1496	CVT(E+LTC)	1070~1080	5	20.0	116	16.0	C-V-EP	3W+EGR	F	走行駆動改善	燃費効率	125	
		L15A	CVT(E+LTC)	1496	CVT(E+LTC)	1080~1100	5	20.0	116	16.0	C-V-EP	3W+EGR	F		燃費効率	125	
		L15A	CVT(E+LTC)	1496	CVT(E+LTC)	1080~1100	5	19.6	118	16.0	C-V-EP	3W+EGR	F		燃費効率	120	
DBA-GE6		L15A	CVT(E+LTC)	1496	CVT(E+LTC)	1080~1100	5	19.2	121	16.0	C-V-EP	3W+EGR	F	CVTパワーマ-	燃費効率	120	
		L15A	CVT(E+LTC)	1496	CVT(E+LTC)	1080~1100	5	18.8	123	16.0	C-V-EP	3W+EGR	F	タイプ 185/65R16	燃費効率	115	
		L15A	SMT	1496	SMT	1050~1050	5	17.4	133	16.0	V-EP	3W+EGR	F		燃費効率	105	
DBA-GE7		L13A	SAT(E+LTC)	1339	SAT(E+LTC)	1140~1170	5	17.2	135	16.0	V-EP	3W+EGR	A	走行駆動改善	燃費効率	105	
		L13A	SAT(E+LTC)	1339	SAT(E+LTC)	1140~1170	5	17.0	137	16.0	V-EP	3W+EGR	A		燃費効率	105	
		L15A	SAT(E+LTC)	1496	SAT(E+LTC)	1170~1180	5	16.4	142	16.0	V-EP	3W+EGR	A	走行駆動改善	燃費効率	100	
DBA-GE8		L15A	SAT(E+LTC)	1496	SAT(E+LTC)	1180~1170	5	16.2	143	16.0	V-EP	3W+EGR	A		燃費効率	100	
		DBA-GQ2	LEA(内排気筒)MF(電動機)	CVT(E)	1339	CVT(E)	1150~1230	5	30.0	77	16.0	C-VT-EP-C-H	3W+EGR	F		燃費効率	125
		DBA-GQ7	L15A	CVT(E+LTC)	1496	CVT(E+LTC)	1140~1170	5	20.0	116	16.0	C-V-EP	3W+EGR	F		燃費効率	120
DBA-GQ6	L15A	SAT(E+LTC)	1496	SAT(E+LTC)	1210~1240	5	16.4	142	16.0	V-EP	3W+EGR	A		燃費効率	120		
	DBA-CB3	L15A	CVT(E+LTC)	1496	CVT(E+LTC)	1270~1300	4.5/7/8	17.0	137	13.0	C-V-EP	3W+EGR	F	CVTパワーマ-	燃費効率	120	
	DBA-CB3	L15A	SAT(E+LTC)	1496	SAT(E+LTC)	1270~1300	4.5/7/8	16.4	142	13.0	C-V-EP	3W+EGR	F		燃費効率	125	
DBA-CB4	L15A	SAT(E+LTC)	1496	SAT(E+LTC)	1360~1480	4.5/8/7	16.0	166	13.0	V-EP	3W+EGR	A		燃費効率	105		
	L15A	SAT(E+LTC)	1496	SAT(E+LTC)	1360~1440	4.5/7	14.0	166	13.0	V-EP	3W+EGR	A		燃費効率	105		
	DBA-GP3	LEA(内排気筒)MF(電動機)	CVT(E)	1496	CVT(E)	1380~1420	5/8/7	24.0	97	13.0	V-TEC-EP-C-H-CV	3W+EGR	F		燃費効率	120	
DBA-CB3	L15A	CVT(E+LTC)	1496	CVT(E+LTC)	1270~1280	4.5/8/7/8	17.0	137	13.0	C-V-EP	3W+EGR	F	CVTパワーマ-	燃費効率	125		
	L15A	CVT(E+LTC)	1496	CVT(E+LTC)	1270~1350	4.5/7/8	16.4	142	13.0	C-V-EP	3W+EGR	F		燃費効率	125		
	L15A	SAT(E+LTC)	1496	SAT(E+LTC)	1360~1460	4.5/8/7	14.0	166	13.0	V-EP	3W+EGR	A		燃費効率	105		
メテオプロ	DBA-RK1	R20A	CVT(E+LTC)	1997	CVT(E+LTC)	1600~1680	7/8	16.0	145	10.5	C-V-H-EP	3W+EGR	F		燃費効率	125	
		R20A	CVT(E+LTC)	1997	CVT(E+LTC)	1630~1680	7/8	15.8	147	10.5	C-V-H-EP	3W+EGR	F	タイプ 205/65R16	燃費効率	125	
		R20A	CVT(E+LTC)	1997	CVT(E+LTC)	1580~1670	8	14.2	163	10.5	C-V-EP	3W+EGR	F		燃費効率	120	
	DBA-RK1	R20A	CVT(E+LTC)	1997	CVT(E+LTC)	1580~1670	8	14.2	163	10.5	C-V-EP	3W+EGR	F		燃費効率	125	
		R20A	CVT(E+LTC)	1997	CVT(E+LTC)	1610~1680	8	14.0	166	10.5	C-V-EP	3W+EGR	F	タイプ 205/65R16	燃費効率	125	
		R20A	CVT(E+LTC)	1997	CVT(E+LTC)	1610~1680	8	14.0	166	10.5	C-V-EP	3W+EGR	F	タイプ 205/65R16	燃費効率	125	
	DBA-RK3	R20A	CVT(E+LTC)	1997	CVT(E+LTC)	1680~1700	7	16.0	145	10.5	C-V-H-EP	3W+EGR	F		燃費効率	125	
		R20A	CVT(E+LTC)	1997	CVT(E+LTC)	1710	7	15.8	147	10.5	C-V-H-EP	3W+EGR	F	タイプ 205/65R16	燃費効率	125	
		R20A	CVT(E+LTC)	1997	CVT(E+LTC)	1680~1700	7	14.2	163	10.5	C-V-EP	3W+EGR	F		燃費効率	120	
	DBA-RK3	R20A	CVT(E+LTC)	1997	CVT(E+LTC)	1680~1700	7	14.2	163	10.5	C-V-EP	3W+EGR	F		燃費効率	125	
		R20A	CVT(E+LTC)	1997	CVT(E+LTC)	1660~1700	7	14.2	163	10.5	C-V-EP	3W+EGR	F		燃費効率	120	
		R20A	CVT(E+LTC)	1997	CVT(E+LTC)	1680~1700	7	14.0	166	10.5	C-V-EP	3W+EGR	F	タイプ 205/65R16	燃費効率	125	

Appendix D. Notes on Model-level Assignment

Figure A3 reports the summary of model histories and box diagrams describing the distribution of variant-level curb weights 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. The graph demonstrates that for virtually all models, the mean and the median values lie within a single weight bin.

Figure A3. Model History and Distribution of Curb Weights for Vehicle Models
Assigned to High Slope Weight Bins



Appendix E. Raw DDD Table

Table A1 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.

Table A1. Unconditional Difference-in-difference-in-differences (DDD) Estimates

	1. High Compliance Cost Segments (Treated)			2. Low Compliance Cost Segments (Control)		
	i. High Slope (Treated)	ii. Low Slope (Control)	Cross-sectional Difference (i - ii)	i. High Slope (Treated)	ii. Low Slope (Control)	Cross-sectional Difference (i - ii)
<i>A. Pre-2007 Period (2002-2006)</i>						
Number of models	23	13		7	45	
Number of variants	513	442		217	1,089	
Variants per model	22	34		31	24	
Fuel Economy (km/L)	17.5 (3.2)	18.8 (1.8)	-1.3 (0.2)	9.9 (2.4)	13.3 (4.6)	-3.4 (0.3)
<i>B. Post-2007 Period (2010-2012)</i>						
Number of models	14	10		3	25	
Number of variants	275	290		90	341	
Variants per model	20	29		30	14	
Fuel Economy (km/L)	18.9 (3.3)	23.9 (3.4)	-5.0 (0.2)	10.6 (1.1)	14.1 (5.3)	-3.4 (0.5)
<i>Time difference (B - A)</i>						
Fuel Economy (km/L)	1.4 (0.2)	5.1 (0.2)		0.7 (0.6)	0.8 (0.3)	
<i>Difference-in-differences (DD)</i>						
Fuel Economy (km/L)			-3.7 (0.3)			-0.1 (0.6)
<i>Difference-in-difference-in-differences (DDD)</i>						
Fuel Economy (km/L)						-3.6 (0.7)

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.

Appendix F. Naive Welfare Estimate

This text documents how we computed a ‘naive’ estimate of the welfare loss from the distortion in technical change. We consider a counterfactual in which the fuel economy ratings of vehicles assigned to the high-slope bins would improve the same way as those assigned to the low-slope bins, 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. This is a ‘naive’ 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 one unit of fuel economy improvement, $-\theta/\lambda p = 18.16$ (in 2012 Japanese yen). Assuming that prices and other attributes (and hence, consumer demand for each car model) stay the same, we can 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.

Figure A4 visualizes what is driving our welfare estimate. **Top panel (a) of Figure A4** 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. We obtain a larger welfare loss estimate per vehicle than Ito and Sallee do precisely because ours account for the size of sales and for the technical change (only) while they do not account for the size of sales but account only for the adjustment to weight and fuel economy ratings conditional on technical change.

Figure A4. Model-level Sales and Welfare Impact of DTC in the Attribute Space

