

# Who Benefits When Firms Game Corrective Policies?

Mathias Reynaert

Toulouse School of Economics, University of Toulouse Capitole

James M. Sallee

University of California, Berkeley  
and the NBER\*

August 12, 2019

## Abstract

Firms sometimes comply with externality-correcting policies by gaming the measure that determines policy. This harms buyers by eroding information, but it benefits them when cost savings are passed through into prices. We develop a model that highlights this tension and use it to analyze gaming of automobile carbon emission ratings in the EU. We document startling increases in gaming using novel data. We then analyze the effects of gaming in calibrated simulations. Over a wide range of parameters, we find that pass through substantially outweighs information distortions; on net, buyers benefit from gaming, even when they are fooled by it.

Keywords: gaming, corrective taxation, Goodhart’s Law, environmental regulation, carbon emissions, automobiles, fuel economy

JEL: Q5, H2, L5

---

\*A previous version of this paper was circulated with the title “Corrective Policy and Goodhart’s Law: The Case of Carbon Emissions from Automobiles.” We thank seminar participants at Arizona State, ENSAE-CREST, Fundação Getúlio Vargas, the NTA Annual Conference on Taxation, the NBER Summer Institute, the Stanford Institute for Theoretical Economics (SITE), the Norwegian School of Economics, Resources for the Future, Texas A&M, the Toulouse School of Economics, UC Davis, Yale, CEPR Applied IO conference Jerusalem, ETH Zurich, University of Antwerp, Erasmus University, Heidelberg University and IIOC Boston (discussant Alex MacKay) for helpful comments. Special thanks to Meredith Fowlie for introducing us to the term Goodhart’s Law in this policy context. Sallee thanks support from the USDA National Institute of Food and Agriculture. Reynaert and Sallee thank support from the France-Berkeley Fund. Reynaert acknowledges funding from ANR under grant ANR-17-EURE-0010 (Investissements d’Avenir program).

© 2016 Mathias Reynaert and James Sallee. All rights reserved.

# 1 Introduction

Sometimes firms comply with a regulation by gaming the measure targeted by policy, rather than by changing their true behavior as intended. This relates to Goodhart’s Law, which posits that “when a measure becomes a target, it ceases to be a good measure.”<sup>1</sup> We study how gaming of this sort can impact buyers who rely on manipulable measures for making choices in a market. Gaming erodes buyer information and induces mistakes. But, when the gaming is done in reaction to a regulatory constraint, sellers lower their costs by gaming, which benefits buyers via pass through. The impact of gaming on buyer surplus is thus ambiguous, even when the gaming completely fools buyers, and the net effect depends critically on whether or not gaming is done in response to a policy.

This paper explores the impact of gaming on buyer surplus both theoretically and empirically for the case of automobile fuel consumption ratings. We do three things in this paper. First, we use a novel data set to measure on-road fuel consumption and document gaming of fuel consumption ratings, which escalates dramatically following the introduction of regulations that target this rating. Second, we develop a theoretical model that derives the impact of gaming on buyer surplus in a setting where sellers game energy-efficiency ratings, which buyers use to evaluate products. Third, we conduct analysis using structural estimates of the automobile demand system to quantify the effects on buyer surplus identified by our theory in calibrated simulations.

Our empirical analysis considers the introduction of stringent corrective policy in the EU automobile market. Prior to 2007, there were no policies in Europe that hinged directly on fuel consumption ratings, but since then both EU-standards and nation-specific tax schemes have created policy incentives that reward lower laboratory fuel consumption test ratings. To measure gaming, we compare the laboratory ratings, which form the basis of policy, with direct measures of on-road fuel consumption that we construct from a data set that tracks fuel consumption and kilometers traveled for a panel of more than 250,000 drivers over twelve years in the Netherlands. Using these data, we estimate the percentage difference between the laboratory test and on-road performance, which we call the *performance gap*, for each vehicle vintage and model.

We document a sharp rise in the performance gap coincident with policy change. Vehicles produced before 2007 show a small, and relatively stable, performance gap. Vehicles produced after that exhibit a large and rising performance gap, so that 2014 model year vehicles have performance gaps in excess of 50% on average. These results are robust to a number of specification checks and various controls and are similar for all automakers. The rise in the performance gap implies that around 65% of the gains in fuel economy since the introduction of policy, as measured by laboratory tests, are false. Using conventional estimates of lifetime distance traveled and a social cost of carbon of \$40 per ton of carbon, the difference between apparent and actual emissions reductions amounts to \$1.2 billion annually from 2010 to 2014 when extrapolated to all of Europe. We interpret the rise in the performance gap as evidence of gaming in response to policy incentives, in the spirit

---

<sup>1</sup>Goodhart’s original concern was monetary policy (Goodhart 1981).

of Goodhart’s Law, and then turn our attention to our central question—who benefits when firms game corrective policy?—both theoretically and in welfare simulations.

Our theoretical model considers a monopolist who sells a good to a representative buyer. The good has some attribute that is desirable to buyers, but it also creates a negative externality that motivates corrective policy. The attribute closely matches the role of fuel economy ratings in the automobile market.<sup>2</sup> The attribute is not directly observable, however, so buyer demand and government regulation are based on a measure provided by the seller. The seller can change the measure either by changing the true attribute or by gaming, both of which are costly.

In our model, we allow that some fraction of gaming is undetected by buyers. In the absence of policy, this means that gaming lowers buyer surplus for two reasons. First, gaming causes buyers to mis-optimize (choose the wrong quantity of the good), which leads to a loss in buyer surplus that we call choice distortion.<sup>3</sup> Second, gaming causes the seller to raise price because buyers perceive an improvement in the product. This price effect further reduces buyer surplus.

Corrective policy disrupts this logic by flipping the sign of price effects. Regulation raises the cost of production. Gaming allows the seller to lower its costs, and this benefits buyers through lower prices in the same way that a reduction in a tax would. When this price effect dominates choice distortions from faulty information, buyers *benefit* from the seller’s gaming, even when they are fooled by it. We believe that our identification of these two competing effects, which highlights how the impact of gaming hinges on the existence of policy, is unique to the literature on gaming.

Throughout the paper we focus on buyer surplus as a notion of the private surplus of consumers that consider buying the good. This is narrower than consumer surplus which would encompass the externality. We focus on incidence and thus on buyer surplus, but gaming will also impact the level of the externality. The ultimate effect of gaming on the environment will turn on the degree of sophistication of the policymaker, who may increasingly tighten policy to achieve real gains, and of the buyers, who may expand the overall size of the market when they mistakenly perceive lower costs of ownership. By definition, in terms of overall social welfare, gaming of an optimal policy is harmful (or at least not helpful). Rather, our point is that buyer surplus may rise from gaming in a variety of circumstances, not that gaming is good for society, except in cases where the regulation is too stringent.

Next, we set out to quantify the price effect and choice distortions in our empirical setting. We demonstrate that the price effects and choice distortions identified by our theory have direct empirical analogs in a discrete choice setting. We then estimate a demand model of the European car market that provides estimates of consumer preferences and the marginal costs of products.

---

<sup>2</sup>In the automobile market, fuel economy is a characteristic valued by buyers, and it is directly linked to an externality (carbon emissions). In terms of the model, fuel economy ratings, which are isomorphic to carbon emissions ratings, are the true attribute. But, true on-road fuel economy is observed by neither car buyers nor the regulator. Instead, regulations and consumer-facing fuel economy labels are based on laboratory tests. Firms can improve a test result either by increasing a vehicle’s true fuel economy or by gaming the test.

<sup>3</sup>The choice distortion that we identify is conceptually identical to the consumer surplus loss due to cognitive frictions or information limitations in several prior papers, including Leggett (2002), Allcott (2013), Sallee (2014) and Dubois, Griffith, and O’Connell (2018).

Given these preferences and costs we calibrate the incidence of gaming for a range of alternative assumptions regarding consumer awareness, policy stringency, and the degree of gaming.

We find robust results that align with our theoretical predictions. When there is no corrective policy and when consumers are fooled by gaming, we find that lowering perceived fuel costs through gaming leads to modest losses in buyer surplus. A significant majority of these losses come through the price effect. As a result, firm profits rise with gaming, and this comes at the expense of lower buyer surplus.

As suggested by the theory, the welfare effects of gaming change when we introduce a corrective policy. We model a mandated decrease in average fuel consumption ratings that firms comply with via shifting their sales mix towards more efficient models. When firms comply honestly, private consumer surplus falls substantially, as consumers are forced into less desirable products. When we allow firms to relax this regulatory constraint by gaming to meet the standard, we find beneficial price effects for buyers that consistently dominate choice distortions so that the net impact of gaming is to raise buyer surplus. Gaming with or without a policy induces a similar sized choice distortion, but we find that this is an order of magnitude smaller than the price effect for a significant range of parameter choices. These results provide empirical validation for our theoretical prediction: gaming benefits consumers in the presence of stringent policy, even when buyers are fooled. We also show that a corrective policy roughly triples the private benefit to a single firm that games when all others are honest, which implies that policy amplifies competitive pressures that incentivize gaming.

Our paper makes several contributions. First, a large literature evaluates the efficiency and market implications of fuel economy standards.<sup>4</sup> We show that gaming in response to standards is important, though our focus is on establishing the incidence of gaming, not on providing a policy evaluation of the EU fuel economy standard. For such an evaluation, see [Reynaert \(2017\)](#), which shows that the standard has induced firms to abate emissions with a mixture of technology adoption and gaming.

Second, we bridge separate literatures on the effects of gaming on market outcomes and the incidence of externality-correcting policies. The gaming literature has considered what we call the choice distortion, and the policy evaluation literature has considered what we call price effects. We combine these in our theory model and then determine which one dominates in our empirical analysis.

A theoretical literature on gaming has considered how gaming can impact market outcomes in a variety of models (e.g., [Gabaix and Laibson 2006](#), [Heidhues, Köszegi, and Murooka 2017](#), [Frankel and Kartik Forthcoming](#)), but to the best of our knowledge, only [Rhodes and Wilson \(Forthcoming\)](#)

---

<sup>4</sup>Key studies include [Goldberg \(1998\)](#), [Austin and Dinan \(2005\)](#), [Bento, Goulder, Jacobsen, and von Haefen \(2009\)](#), [Anderson and Sallee \(2011\)](#), [Klier and Linn \(2012\)](#), [Jacobsen \(2013\)](#), [Gillingham \(2013\)](#), [Parry, Evans, and Oates \(2014\)](#), [Whitefoot, Fowlie, and Skerlos \(2017\)](#), and [Durmeyer and Samano \(Forthcoming\)](#). Recent articles have critiqued efficiency of standards due to interactions with the used car market ([Jacobsen and van Benthem 2015](#)), safety ([Jacobsen 2013](#), [Bento, Gillingham, and Roth 2017](#)), unintended consequences of attribute basing ([Ito and Sallee 2018](#), [Whitefoot and Skerlos 2012](#)), and heterogeneity in vehicle lifetimes ([Jacobsen, Knittel, Sallee, and van Benthem 2016](#)).

and [Drugov and Troya-Martinez \(2015\)](#) consider gaming in the context of regulatory enforcement. In contrast to [Rhodes and Wilson \(Forthcoming\)](#), we consider gaming in response to externality-correcting policy when strong enforcement is lacking. In contrast to [Drugov and Troya-Martinez \(2015\)](#), we include endogenous prices for the good. Empirically, there are a number of papers that show how market measures can be manipulated so as to create market distortions, including test scores ([Dee, Dobbie, Jacob, and Rockoff 2016](#)), mortgage markets ([Ben-David 2011](#), [Jiang, Nelso, and Vytlačil 2014](#)), or corporate earnings ([Roychowdhury 2006](#)). None of these weigh the competing effects of information-based losses and positive price effects on buyer surplus, which is our focus.

On the other hand, many papers have analyzed the incidence of externality-correcting policies, showing, for example, how the costs of complying with environmental regulations are passed through into product prices ([Jacobsen 2013](#)), energy prices ([Fabra and Reguant 2014](#)), or labor earnings ([Walker 2013](#)). We study the flip side of this coin, which is how cost savings from regulatory avoidance can pass through to buyers of a regulated good. We add to that literature by integrating the welfare implications of information erosion into incidence analysis.

Third, our analysis is also relevant to the specific issue of emissions testing reliability, which has become a major policy topic in the wake of the Volkswagen emissions testing scandal. Our reduced-form estimates of the degree of gaming are consistent with a growing set of media accounts, with analysis performed by the International Council for Clean Transportation, which has documented a similar gap in a number of European countries ([The International Council on Clean Transportation 2014, 2015](#)), and with quasi-experimental evidence on fuel consumption gaming in Japan ([Tanaka 2017](#)).<sup>5</sup> We add further evidence of gaming to that body of work, and also introduce the first analysis of how emissions test gaming impacts the private welfare of car buyers.

Finally, recent research has documented a failure of energy efficiency programs to realize ex ante expected savings in appliances ([Davis, Fuchs, and Gertler 2014](#)) and buildings ([Fowlie, Greenstone, and Wolfram Forthcoming, Levinson 2016](#)). We provide related evidence of a lack of efficacy in energy-efficiency policies for automobiles, which have delivered only a fraction of expected savings because of gaming.

The paper is structured as follows. We begin with our reduced form analysis of the performance gap in section 2. In section 3 we describe our theoretical model, which delineates choice distortions and price effects. Section 4 estimates a structural model of the automobile market and uses the estimates to calculate changes in consumer surplus due to choice distortions and price effects from gaming in our calibrated simulations. Section 5 concludes.

---

<sup>5</sup>The ICCT reports include analysis of the Netherlands based on the same data that we use here. Our results are broadly similar, but our analysis contributes by providing an independent academic analysis and through a number of methodological improvements, including the calculation of standard errors, the introduction of time of driving fixed effects and other controls, a variance decomposition, an Empirical Bayes correction, and extrapolation to the entire European market through matching.

## 2 Automakers game EU fuel consumption ratings

In this section we discuss Goodhart’s Law in the EU automobile market. We document a sharp decrease in the correlation between official and on-road fuel consumption ratings that coincides with those ratings becoming subject to regulation. We start by describing policy changes in the EU automobile market and detailing how carmakers report official fuel consumption ratings. We then describe data from a large sample of drivers in the Netherlands that we use to construct estimates of on-road consumption. Finally, we quantify the gap between official ratings and on-road consumption and relate that pattern to the rollout of policy.

### 2.1 European fuel economy regulations are relatively new

Prior to 2007, automobiles in Europe were not subject to explicit regulation based on their carbon emissions. Since then, automobiles have become subject to both an EU-wide regulation and nation-specific tax and subsidy schemes that are explicitly tied to carbon emissions. The EU-wide regulation mandates that automakers sell vehicles that have a sales-weighted average carbon emission rate below a certain level. A vehicle’s carbon emission rate is measured as a linear transformation (to convert units) of fuel consumption, which itself is just the inverse of fuel economy. Thus, while the EU policy is described as regulating carbon, it is directly analogous to the US Corporate Average Fuel Economy standards, which are described as regulating fuel economy.

Prior to 2007, there was no legally binding fuel economy regulation in Europe. The standard was announced in 2007, passed into law in its final form in 2009, and phased in over several subsequent years. The first year of enforcement was 2012, with a ramp-up in the standard taking place from 2012 to 2015. Fully phased-in, the regulation is quite aggressive by historical and international standards. Fines for non-compliance are stiff. When the standard was announced no automaker was complying with the standard. If automakers had left their fleets unchanged, they would have faced an average fine of 1,250 euros per vehicle.<sup>6</sup>

In addition to the EU standard, all member states have separate tax schemes for new vehicles. During the rollout of the EU-wide regulation, many member states adjusted their national policies so that tax schemes were based explicitly on carbon emission rates. France introduced a feebate, which taxes heavily-polluting cars and subsidizes cleaner models, in 2008. In 2009, Germany switched its annual road taxes so that they depend on carbon emissions rates, rather than engine cylinders. In 2008, Spain introduced registration taxes that depend on emissions ratings. The Netherlands did the same in 2010. [Gerlagh, van den Bijgaart, Nijland, and Michielsen \(Forthcoming\)](#) present a detailed overview of the move towards carbon taxation in EU member states after 2007.

### 2.2 How are fuel economy ratings measured?

Carbon emissions ratings for all of these policies are based on a laboratory test, called the New European Driving Cycle (NEDC). This test procedure, which is conducted by third-party facili-

---

<sup>6</sup>We describe the EU standard in more detail in the appendix [A.2](#).

ties that are funded by the automakers, measures fuel consumption in liters per 100 kilometers (L/100km). A test vehicle is put onto a chassis dynamometer (a treadmill for cars), and a professional driver “drives” the car through a specified series of speeds and accelerations. Emissions are captured directly from the tailpipe and used to determine gaseous outputs, which are used to determine fuel consumption. Two different test cycles are run to simulate city and highway driving. A coast down test is also performed to measure aerodynamics.

The NEDC is not only the basis of CO<sub>2</sub> regulation, but also consumer-facing labels and emissions limits on local air pollutants, such as NO<sub>x</sub>, PM and CO. The test procedure captures local pollutants and measures their quantities to determine vehicle compliance with emissions limits. In terms of consumer information, the NEDC rating is the rating that automakers are required to use in consumer advertising, and it is the rating that appears on mandatory energy efficiency labels for new vehicles. As such, the NEDC was used as a regulatory and market instrument for roughly a decade before the rollout of carbon regulation.

We are interested in how automakers might game this test. Tests are performed in third-party facilities, but these are funded by the automakers and may not have incentive to maintain test integrity. According to media and industry accounts, the European test procedure offers the tester considerable “flexibility” in test procedures.<sup>7</sup> For example, automakers are not expressly prohibited from submitting test vehicles that have been modified. Firms remove optional equipment, thereby changing the weight of the vehicle to improve performance. They also tape down seams in the vehicle, remove side mirrors and roof racks, and over-inflate tires to improve results of the coast down test. Alternatively, automakers may install technologies that perform particularly well on the test cycle, or they could even calibrate an engine to perform in a particularly efficient way during the highly specific test cycle’s series of speeds and accelerations.<sup>8</sup> To detect gaming, we compare laboratory test ratings with a direct measure of on-road fuel consumption, which we describe next.

## 2.3 On-road fuel consumption data

We obtained data from TravelCard NV, a company providing fuel services in the Netherlands. These panel data contain information on 66 million fuel station visits from drivers using a TravelCard NV card between January 2004 and May 2015. Most of the individuals in this sample drive a vehicle provided to them by their employer, who also pays directly for fuel. This implies that we have a selected sample, though the provision of a company car is quite common in the Netherlands due to tax advantages and the high cost of personal vehicle ownership. Nearly half of new vehicles sold in the Netherlands are registered as company cars (Booz & Company 2012).

When visiting a fuel station, Travelcard NV users swipe a smart card to pay for fuel. When a driver swipes her card we observe the drivers’ license plate and the date, time and location of the fuel station visit. We also observe the exact amount and the type of fuel purchased and a self-

---

<sup>7</sup>For example, see “Europe’s Auto Makers Keep Test Firms Close” in the March 21, 2016 *Wall Street Journal*: <http://on.wsj.com/1o5h47B>.

<sup>8</sup>These “flexibilities” differ significantly from analogous tests in the US, where the law is much more explicit about the details of the test vehicle and test procedures.



reported odometer reading at the time of fueling. Regarding the latter, drivers are asked to enter their current odometer reading into the fuel pump’s keypad during the transaction. Transactions are linked to an account, which is a unique combination of a driver with a particular vehicle, which corresponds to a unique license plate. Unfortunately, if the same individual drives one vehicle for some time and then switches to another vehicle, we have no way of linking those data and must treat them as separate accounts. We refer to this panel variable as a driver, which should be understood as a unique combination of a driver with a particular vehicle.

TravelCard NV provides us with a second dataset that matches each license plate with the vehicle brand, model name, weight, fuel type and the official fuel consumption rating of the vehicle. These characteristics allow us to match the Dutch data with a panel on European car sales and prices from 1998-2011 used in [Grigolon, Reynaert, and Verboven \(Forthcoming\)](#) and [Reynaert \(2017\)](#). We match each vehicle to sales volume, list price and a broader set of characteristics, including length, width and several engine characteristics. We define a vehicle as the unique combination of brand (Volkswagen), model name (Golf), fuel type (Diesel) and official fuel consumption. We define the release year for a vehicle as the first time we observe a unique combination of these variables in the data between 1998 and 2014. We observe vehicles at the engine level and in the empirical analysis we will be able to control for vehicle models as the combination of brand and model name. Odometer readings are self-reported and many individual entries are unreliable. In [appendix A.1](#) we detail steps we take to purge unreliable entries. Our final sample includes 24 million transactions. These data include over 2,500 unique vehicles driven by 266,000 different drivers.

[Table A.1](#), in the appendix, gives summary statistics for the raw data. The average vehicle in our sample has a fuel consumption of 6.65 L/100km and a weight of 1,354kg. Somewhat less than half of the vehicles (46%) have diesel engines. We observe an average of 107 drivers per car with a maximum of 3,228 drivers. For each of these drivers we observe an average of 134 visits to the pump with a maximum of 1,135. Driver mean total consumption is 6,015 liters of fuel purchased corresponding to 111,726 km traveled. Finally, the average fuel station visit involves 45 liters of fuel purchased, corresponding to an odometer increase of 671 km with a standard deviation of 192 km.

Our sample is drawn from company cars. We compare these to the overall market in [appendix Table A.2](#). Vehicles in the Travelcard data are cheaper, lighter, more fuel efficient and more likely to be a diesel than the average new vehicle in the Netherlands. One explanation of the difference is that our data contain almost no luxury vehicles or sports cars. Also, company cars are driven more than the average household car, which likely explains the higher share of diesels in our sample.

Nevertheless, the models in our final sample are the models that account for 76% of all sales in the Netherlands. Moreover, the Dutch car market is very similar to the entire EU market; it features a mix of French, German and foreign brands that is very similar to the EU average. Thus, our data provide estimates of on-road performance that cover the broader market fairly well. However, we wish to stress that we have no way of directly assessing whether drivers of company cars drive their vehicles differently, and thereby exhibit a different performance gap, than the average Dutch



driver. Fortunately, we are focused on *changes* in the performance gap over time, so as long as differences between company car drivers and other drivers are stable over time, this will not be a major concern for our conclusions.

## 2.4 On-road fuel consumption and the performance gap

We construct a measure of on-road fuel consumption  $r_{nij}$  for each pump visit  $n$  of driver  $i$  in car  $j$  as the ratio of the liters purchased and the change in reported odometer between the visit and the previous visit:

$$r_{nij} = \frac{\text{liters}_n}{\text{odometer}_n - \text{odometer}_{n-1}} * 100. \quad (1)$$

This measure of on-road fuel consumption, in units of L/100km, will vary between pump visits of a driver for three reasons. First, variable driving conditions such as outside temperature, route choice, driving style and congestion will differ across observations. Second, the driver may over or understate the odometer reading. We are not aware of any incentive for the drivers to deliberately misreport distance traveled, but from the data it is obvious that there are many mistakes. Third, there might be variability due to stockpiling effects. If the consumer does not always fill the tank of the vehicle completely there will be variation in  $r_{nij}$ . If a driver visits the fuel station with an empty tank and fills half of it we will observe a very low fuel consumption for visit  $n$  and a higher fuel consumption for the next visit if she refills the tank completely.

Next, we construct the percentage gap between on-road and tested fuel consumption as:

$$d_{nij} = \frac{r_{nij} - li_j}{li_j}, \quad (2)$$


in which the official rating  $li_j$  is constant for each car  $j$  and the on-road rating varies across observations. We call this the performance gap. We are interested in estimating the mean and variance of  $d_j$ , defined as the average  $d_{nij}$  across  $n$  and  $i$  for a given vehicle type  $j$ . In particular, we are interested in the mean and variance of  $d_j$  across vehicle types from the same vintage (release year). We now discuss this performance gap by presenting its evolution over time, by estimating a fixed effects model and by analyzing heterogeneity in the gaps over time.

## 2.5 The mean performance gap rises over time

Figure 1 plots the mean official rating  $li_j$ , the mean on-road rating  $r_{nij}$  and the percentage-gap  $d_{nij}$  for each release year. Between 1998 and 2006 we see that both official fuel consumption and on-road consumption vary between 6.1 and 7.5 L/100km. The percentage gap fluctuates between 12% and 20% and shows evidence of a slight upward trend after 2004. From 2006 onward we see a spectacular drop in official consumption from 6.2 to less than 4 L/100km. This translates to a rise from the already high value of 38 mpg in 2004 to a truly remarkable 67 mpg in 2014. Official fuel consumption decreases by almost 50% over the sample period.

The on-road ratings follow a trend similar to the official ratings up until 2008. After 2008

**Figure 1:** On-road and official fuel consumption per release year



the on-road fuel consumption decreases much more slowly than the official rating, going from 6.5 to 5.9 L/100km. As a result, the percentage gap between the official and on-road ratings increases dramatically: from around 12% at the beginning of the sample to almost 55% by 2014. This divergence is remarkable in magnitude, and it coincides exactly with the new EU-wide fuel consumption regulation, which was announced in 2007, finalized in 2009, and phased-in over the remainder of the sample, as well as the phase-in of national policies described above.

The divergence in test ratings and on-road performance is not isolated to a particular automaker. Figure 2 plots the estimated mean gap  $d_{nij}$  for three sets of model years (early, middle and late) separately for each automaker. All automakers show a substantial increase in the gap over time, and all show an economically important gap in the later years. The performance gap is a global phenomenon: the three largest gaps are for a European firm (Renault), a Japanese firm (Toyota), and an American firm (Ford).<sup>9</sup>

## 2.6 The trend in the performance gap is robust to controls

To account for various potential confounders, we next estimate the performance gap through a series of regressions on our microdata that take the following form:

$$d_{nij} = RELEASE_j + FUEL_j + MODEL_j + f(DIST_i) + f(YEAR_n, MONTH_n, AGE_j) + \epsilon_{nij},$$

<sup>9</sup>We also show in Appendix Figure A.4 and A.5 that the increase in gaps is similar across weight quartiles.

**Figure 2:** Gap between on-road and official fuel consumption per firm




Figure shows estimated coefficients and standard errors from regressing the performance gap on three sets of model release years (early, middle and late) per automaker, the data is restricted to the ten brands with most observations.

where  $RELEASE_j$  are release year (i.e., vintage) dummies;  $FUEL_j$  is a dummy for diesel engines;  $MODEL_j$  are model dummy variables;  $f(DIST_i)$  is a third order polynomial in the distance per year driven by each driver  $i$  and  $f(YEAR_n, MONTH_n, AGE_j)$  are time and age controls explained in more detail below.

We interpret the release year dummy coefficients as average measures of the performance gap for each vintage. These are our coefficients of interest. The omitted category is vehicles present in the first year of our data, which implies that they were released in 1998 or before. Our identification challenge is to separate the release year coefficients from other changes in the fleet over time. Table 1 and Figure 3 show the coefficients of the release dummies for five different specifications that introduce different controls.<sup>10</sup>

Column (1) in Table 1 shows results from a regression of the percentage-gap  $d_{nij}$  on release year dummies without any controls. The  $R^2$  for this model is .28, which suggests that release years explain a considerable amount of variation given that the unit of observation is an individual fuel transaction.<sup>11</sup> The regression shows a clear upward trend in the release year dummies. Our interpretation of this time trend in the performance gap is that the accuracy of official fuel consumption ratings declines over time, due to the introduction of European fuel economy regulations.

We next estimate a series of alternative specifications that aim to control for other factors that might bias the release year coefficients. The specification in column (2) adds our preferred controls, including model fixed effects, time of driving (i.e., the year and month in which a fuel transaction is observed) fixed effects, and a polynomial in annual mileage for each driver. If some vehicle types have a bigger gap than others, then our time trend could reflect compositional changes. We add a fuel type dummy for diesel (versus gas) and dummies for each model to control for compositional changes. In our parlance, a model is more aggregate than a vehicle  $j$ . For example, a Toyota Camry is a model, whereas a 2010 Toyota Camry with 2.0L engine is a vehicle type. Inclusion of model fixed effects isolates variation between engines that differ over vintages of a model.<sup>12</sup> Year and month of driving dummies control for driving conditions that change seasonally or over time. These time controls account for changes in driving behavior, sample selection and other factors such as fuel prices that vary over the sample period. We can separate out release year and time effects

---

<sup>10</sup>Ideally one would like to leverage variation in the regulatory incentives across vehicles or firms to aid causal identification, but the features of the policy preclude several approaches. First, the regulation was announced in 2007 and was made fully binding in 2015, which means that there is not a sudden surprise policy shock to examine. Second, the regulation binds for the fleet wide average for an automaker, so automakers have common incentives to improve all vehicles. Moreover, the program lacks the nonlinear notches used for identification in related contexts (Ito and Sallee 2018, Sallee and Slemrod 2012, Tanaka 2017). Third, the regulation was explicitly designed with a weight-based standard to reduce variation in compliance burdens across automakers, which sharply reduces the cross-firm variation in policy impacts. Moreover, the policy allows for “joint compliance” between automakers, which further limits cross-firm variation. As a result, no firms were already in compliance, and all firms had to make similar improvements.

<sup>11</sup>The estimate for release year 2006 is considerably less precise than other years. The Travelcard data contain a larger number of observations from new model releases that year, which might explain the imprecision. Regardless of the cause, the anomaly is eliminated once other controls are added.

<sup>12</sup>Note that every engine type is given a unique fuel consumption rating. These engine types are sold until there is a redesign which comes with a new fuel consumption rating. The redesign does not happen yearly, so that we do not have the same strict ‘model-year’ definition as in the US.

because we see the same vehicle driven over multiple years. For example, we observe both 2008 and 2009 vehicles being driving in 2010, and so we can estimate vintage effects (2008 versus 2009) while controlling for time of driving (2010). Finally, we control for the distance that each driver covers per year with a 3rd order polynomial as a way of capturing differences in driving patterns across individuals. As shown in column (2) of Table 1, adding these controls has very little impact on the release year coefficients.

Column (3) shows results from a specification that includes month by year (i.e., month of sample) dummies, rather than year and month (i.e., season) dummies. This controls more flexibly for time effects. Again, the trend in the release year fixed effects is unaffected by the introduction of these richer time of controls. The point estimates are almost identical to those in the second column.

Our specifications above identify release year effects controlling for time period shocks, but an additional concern arises if fuel consumption changes with the age of a vehicle. In that case, we face a version of the well known age-cohort-time identification problem (see, e.g., [Card and Lemieux 2001](#)). Emissions of local air pollution are known to degrade substantially with vehicle age, but fuel economy is typically assumed to be stable. Even so, there could be a form of selection through which older vehicles are driven in different patterns or styles on average, which will manifest as an age effect in our on-road performance data. Regardless, we would like to be able to control for age effects, but age fixed effects are perfectly collinear with release year and time fixed effects. We cannot separately identify all effects without some additional restriction. One approach is to assume that release year effects are completely flat prior to 2006, so that a single release year dummy for that period is included. A full set of age effects can be estimated based on that period, and age and time effects can both be included while still identifying the remaining release year effects. Column (4) reports results from a specification that does that. Results are qualitatively similar, though the magnitude of our estimated performance gap does change somewhat, falling from 41% to 34% as compared to column (3), which has the same time period controls. As discussed below, we provide alternative approaches of addressing age effects in the appendix, which yield qualitatively similar results to the results reported here.

Finally, column (5) reports results from a regression where we use long differences to construct the fuel performance gap, instead of differences between fuel station visits. Odometer readings are self-reported and errors in the readings result in positive measures of the gaps. Resulting attenuation could be responsible for some of the estimated performance gap prior to the increase in gaming, and correlation between reporting errors and release years could bias our findings about the trend over release years. The analysis with long differences is more robust to individual reporting errors. The drawback is that it does not allow us to control for time of driving effects in the regressions. Regression (5) shows that our findings are robust: we get nearly identical estimates when we use long differences, taken between the last and first fuel station visit for each driver.<sup>13</sup>

---

<sup>13</sup>Note that the measurement error in the data is an additional reason, on top of sample selection, that makes it difficult to interpret the level of the gap in the beginning of the sample.

**Table 1:** Release year fixed effects from regression of performance gap on controls

	All (1)	Controls (2)	MbY (3)	Age (4)	Long (5)
Constant	11.69 (0.66)	13.52 (0.65)			7.80 (0.61)
1999	2.00 (1.45)	0.53 (0.88)	0.52 (0.88)		2.45 (1.29)
2000	0.32 (0.88)	0.75 (0.72)	0.75 (0.72)		0.38 (0.88)
2001	2.329 (0.98)	-0.18 (0.95)	-0.18 (0.95)		2.70 (1.02)
2002	1.77 (1.56)	1.27 (1.07)	1.28 (1.07)		1.48 (1.48)
2003	2.46 (1.32)	1.31 (0.80)	1.31 (0.80)		1.49 (1.46)
2004	3.30 (1.41)	2.57 (0.78)	2.57 (0.78)		1.24 (1.33)
2005	4.71 (1.43)	4.20 (0.967)	4.20 (0.97)		3.60 (1.25)
2006	10.59 (5.09)	3.84 (1.09)	3.83 (1.09)	0.28 (1.03)	7.57 (4.52)
2007	7.38 (1.39)	8.46 (1.04)	8.44 (1.04)	4.40 (1.07)	6.02 (1.21)
2008	9.53 (1.37)	8.86 (1.00)	8.83 (1.00)	4.27 (1.08)	6.21 (1.20)
2009	18.50 (2.24)	13.74 (1.41)	13.72 (1.41)	8.72 (1.72)	13.65 (1.89)
2010	22.57 (2.10)	17.93 (1.44)	17.91 (1.44)	12.47 (1.54)	18.17 (2.11)
2011	23.43 (1.57)	18.34 (1.13)	18.33 (1.13)	12.34 (1.39)	18.29 (1.24)
2012	35.16 (2.95)	28.62 (2.09)	28.65 (2.09)	22.09 (2.31)	29.55 (2.94)
2013	41.24 (1.89)	35.86 (1.66)	35.93 (1.65)	28.87 (1.96)	38.58 (1.80)
2014	42.68 (2.47)	40.72 (3.50)	40.71 (3.50)	33.68 (3.66)	46.25 (3.84)
Year + Month F.E.		Yes			
Year by Month F.E.			Yes	Yes	
Fueltype		Yes	Yes	Yes	
Model F.E.		Yes	Yes	Yes	
Distance (3 orders)		Yes	Yes	Yes	
#Obs. (*10 <sup>6</sup> )	23.98	23.98	23.98	23.98	0.24
R <sup>2</sup>	0.281	0.400	0.401	0.401	0.35

Table reports coefficients from a regression of the performance gap ( $d_{nij}$ ) on release year fixed effects. The unit of observation is an individual refueling transaction. Standard errors clustered by car type in parentheses. Columns vary as follows: (1) contains all data, (2) all the data with distance, year and month, fuel type and vehicle model fixed effects, (3) add month by year fixed effects to (2), (4) adds vehicle age and constrains release fixed effect to be constant prior 2006, (5) dependent variable is long difference between first and last observation for each driver instead of differences between visits.

**Figure 3:** Release year coefficients from fixed effect regressions




Figure plots coefficients from a regression of the performance gap ( $d_{nij}$ ) on release year fixed effects. Coefficients correspond to the five regressions in Table 1.

In Appendix Table A.5 we report the coefficients for all control variables in the specifications reported here. In addition, we report several additional specifications. We estimate the release year fixed effects separately for gasoline and diesel engines.<sup>14</sup> We also introduce three additional methods of controlling for age. First, we define age as the year a license plate (combination of vehicle and driver) is first observed in our data set, which differs from the release year. This controls for the length of time a driver has had the vehicle. In this way we can include a full set of release, age and time fixed effects because there is variation in age within release year and time, based on when a driver obtained a vehicle. Second, we define age in terms of odometer instead of calendar years and introduce the level of the odometer as a control.<sup>15</sup> Third, we keep only observations of  $d_{nij}$  that take place in the release year of the vehicle, so that we are capturing fuel consumption gaps only among the newest cars. This eliminates concern about age as a confounder, but it does not allow

<sup>14</sup>If more recent vintages are more fuel efficient, consumers may respond by driving less carefully or using temperature controls or other equipment more often. In this case, reduced average fuel consumption rates will lead to an increasing gap. One might expect the same mechanism to create a significant difference in the gap between gasoline and diesel powered vehicles, as diesel vehicles are about 30% more energy efficient. If the rebound effect were a major driver of the performance gap, we would expect to see a bigger gap among diesels that grows more quickly. We find no such pattern.

<sup>15</sup>Note that the age dummies have surprising negative magnitudes in this regression. These can only be due to compositional changes within model groups because we find age trends to be flat within engine versions.



us to control for time period effects. As a final robustness check, we add almost twenty thousand fixed effects that control for the location of the fuel station for each transaction. Figure plots the release year coefficients for all specifications not included in Table 1. In every specification we find a qualitatively similar trend in the performance gap by release year, with a sharp increase in the middle of the sample that ends with a large gap on the order of 30 to 50% by 2014.

Finally, in Appendix Table A.6 and Figure A.2 we show the parallel results that measure the performance gap in levels (liters per 100 kilometers) rather than percentages. These regressions clearly confirm the upward trend in the gap.

In sum, we conclude that the increase over time in the performance gap is robust to changes in fuel prices or congestion over time, to composition of the fleet, to rebound effects, to aging or selection over time, and to bias from stockpiling or reporting errors.

## 2.7 Heterogeneity in the performance gap does not rise over time

The welfare effects of gaming that we explore in the remainder of the paper potentially depend on whether there is significant heterogeneity in gaming within a release year across vehicle types.

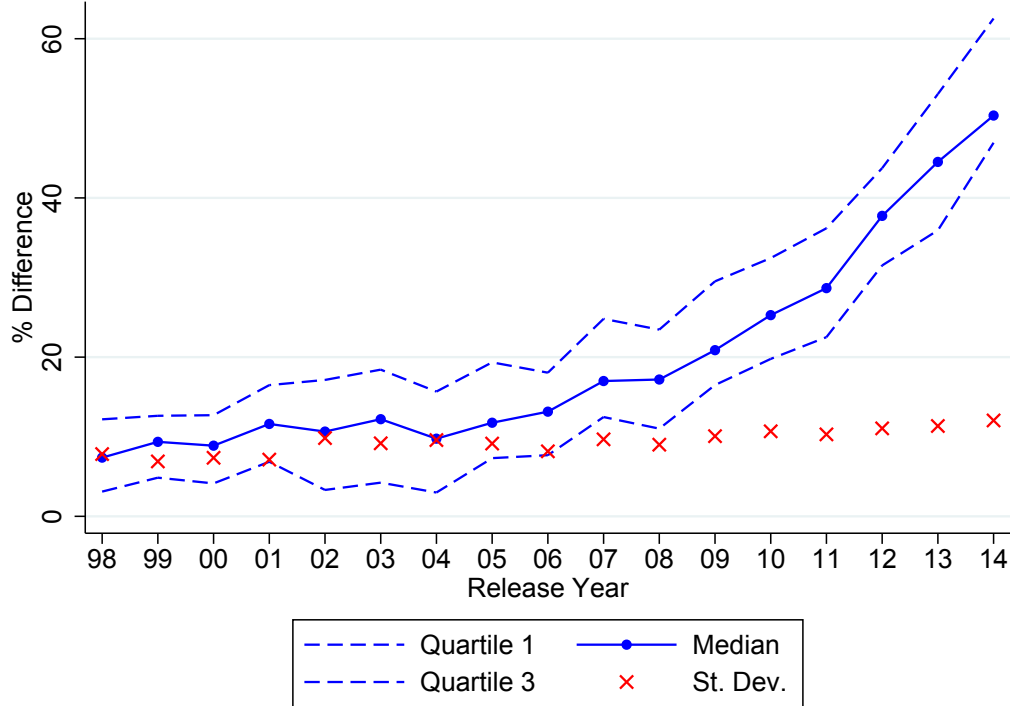
To study potential heterogeneity in gaming at the vehicle level we need vehicle level estimates of the performance gap. In obtaining these estimates we want to take into account the underlying data quality for each vehicle controlling for the large variance in reported odometer readings, and the variation in the number of drivers and visits observed for different types of cars. Therefore, we follow the teacher value added literature, specifically Chetty, Friedman, and Rockoff (2014) and Kane and Staiger (2008), to estimate on-road fuel consumption  $r_j$  using precision weights and an empirical Bayes correction. We construct a precision-weighted mean of  $r_j$  and shrink it according to the reliability of the observations for  $j$ . Because of the measurement error in  $r_{nij}$  it is optimal, from a prediction standpoint, to use a biased but more precise estimate of each vehicle gap.<sup>16</sup> The empirical Bayes procedure and accompanying variance decomposition is discussed in Appendix A.3.<sup>17</sup>

Figure 4 plots the quartiles of the estimated Bayes corrected gaps for all release years as well as the standard deviation of the estimated fixed effects. Again we see a clear increase in the estimated gap over time, but perhaps surprisingly, this increase in the gap is not associated with a large change in the dispersion of the gap. If firms intensively game only some of their models, we would expect a rise in average gaming to correspond to a greater interquartile spread or larger standard deviation, across model types within a vintage. Figure 4 shows, however, that these measures of dispersion did not rise over time. Instead, the data suggest that we should think about

<sup>16</sup>Note that this approach does not allow us to control for selection issues where drivers with different performance gaps select into different type of vehicles. Correlations between driving style and vehicle selection could influence our results, but again we emphasize our focus on changes in the performance gap over time, which mitigates these concerns.

<sup>17</sup>The empirical results are reported in Table A.4. Table A.4 also includes estimates of an alternative two stage approach. The first stage is a regression of the gap on 266,000 driver by vehicle fixed effects and year and month controls. The second stage regresses the driver by vehicle effects on release year dummies, resulting in the same strong upward trend.

**Figure 4:** Distribution of vehicle level fixed effects



gaming in the EU market as an industry-wide shift, rather than a shift among select vehicles. This pattern is consistent with the underlying economics of gaming. Because the regulation is based on fleet averages, automakers have an incentive to game all of their vehicles, not just those with high consumption ratings. Moreover, to the extent that gaming consists largely of modifying test vehicles, the costs of gaming are largely a fixed cost, and it is plausible that techniques developed to game the test vehicle for one model line can be used for most other models.

## 2.8 Summary of reduced-form evidence

We conclude from this empirical section that there is very strong evidence of substantial gaming of fuel consumption ratings. Depending on our estimation method, the performance gap between on-road consumption and official test ratings varies between 10% to 15% before regulation. By 2012 this gap had increased to 30%, and all specifications estimate a mean gap of more than 45% in 2014. We have not found a single robustness check where this spectacular increase in the gap does not occur. Our findings regarding the time path of the performance gap is consistent with analysis performed by the International Council for Clean Transportation ([The International Council on Clean Transportation 2014, 2015](#)).<sup>18</sup> Relative to those reports, we demonstrate robustness to

<sup>18</sup>The ICCT reports results from various sources including the data from Travelcard NV we obtained. The other sources include consumer self-reports uploaded to websites and consumer car testing magazines. The data from Travelcard are by far the most extensive in terms of number of vehicles and drivers included, but it is reassuring that

the inclusion of various controls (most notably time period fixed effects), include standard errors, develop an empirical Bayes correction for sampling error, and calculate the distribution of the vehicle-level performance gap. Regarding the latter, this yields the striking and perhaps surprising finding that the gap is increasing for all vehicles and brands in the sample, so that the distribution of the gap is shifting, but the shape is relatively constant.

We see the decrease in the correlation between official and on-road ratings at a time when stakes on official ratings increase as a manifestation of Goodhart’s law—once policy incentives were loaded onto the fuel economy test results, those test results became a less reliable measure of actual fuel economy performance. In the remainder of the paper, we interpret this time trend as evidence that automakers gamed official ratings in response to new regulations. We recognize that this evidence is suggestive, but we cannot think of any plausible alternative explanation.<sup>19</sup> While the EU commission has released an external evaluation of the emission standard, the official ratings are still the basis for the evaluation of the program and external communications about the program’s impact.<sup>20</sup> Analysts who perform cost-benefit analyses of the regulation may significantly mis-evaluate the program if they calculate savings based on official ratings.

We now turn our attention to our main economic question—who benefits from this gaming? Carmakers game the official rating to avoid stringent regulation, lowering compliance costs of the carmakers. Buyers will benefit from gaming to the extent these cost savings are passed through. But, these official ratings are also used by buyers when deciding what vehicle to buy. We explore these competing forces first with a model and then with a calibrated market simulation.

### 3 A model of gaming

The goal of our model is to describe the impact of gaming on buyer welfare, taking into account how prices will depend on gaming, the presence of a policy, and the degree to which buyers are fooled by the gaming.

#### 3.1 Setup

We model the market for a good that has a mutable attribute  $x$  that generates private value to buyers and also creates a social externality. In the automobile example,  $x$  is the vehicle’s fuel consumption rating. The full cost of the good is denoted  $f = p + \beta x$ , where  $p$  is the up-front purchase price and  $\beta$  is a coefficient that translates fuel consumption ratings into dollars. Note that  $x$  is a bad. For a vehicle,  $x$  is the L/100km rating and  $\beta$  is the price of fuel per liter, inclusive of any taxes  $\tau$ , times the number of present discounted lifetime kilometers driven  $k$ :  $\beta = (p_f + \tau) \times k$ . This

---

all data in the ICCT reports point to very similar results.

<sup>19</sup>Tanaka (2017) shows quasi-experimental evidence on fuel consumption gaming in Japan. The Japanese standard specifies discrete notches, rather than a sales-weighted average as in our case, allowing for a comparison in gaps below and above a notch. Similarly, the first paper to show manipulation of fuel economy ratings in response to incentives uses tax notches to identify manipulation (Sallee and Slemrod 2012).

<sup>20</sup>See for example <http://ec.europa.eu/clima/policies/transport/vehicles/cars>, read on 09/13/2016.

tax inclusive interpretation allows us to later compare the effects of a fuel tax versus a regulation without complicating notation.

We suppose that the good is sold by a monopoly seller. For simplicity, we assume that all attributes of a good are fixed other than  $x$  and  $p$ ; this can be interpreted as a short-run assumption.

Buyers do not observe  $x$  directly, but instead receive a message  $m$  from the seller (i.e., the official fuel consumption rating that appears on labels). Gaming occurs when the seller sends a message  $m$  that differs from  $x$ , where gaming  $g$  is defined as the difference between the true attribute and the message ( $g = x - m$ ). Under full information, buyer demand  $D$  for the product depends on its various attributes (held fixed and thus suppressed in notation) and lifetime cost, which depends on  $x$  and  $p$ :  $D(f) = D(p + \beta x)$ . Buyers, however, observe  $m$  not  $x$ , so they use the observed signal to form beliefs, labeled  $\tilde{x}$ , where  $\tilde{f} = p + \beta \tilde{x}$ . We assume risk neutrality.

Buyer beliefs  $\tilde{x}$  are assumed to be a weighted average of the truth and the signal:

$$\tilde{x}_t = \alpha x + (1 - \alpha)(x - g) = x - (1 - \alpha)g.$$

This is a tractable form of beliefs that encompasses a variety of possibilities. When  $\alpha = 1$ , buyers can see through gaming completely. When  $\alpha = 0$ , buyers are completely fooled. For intermediate cases,  $(1 - \alpha)$  represents the extent to which the buyer is fooled by gaming.<sup>21</sup> The key for our purpose is that we want to allow the *possibility* that firms might fool buyers, at least to some degree. This specification is appealing in part because, for all of our results, we will be able to see how effects differ when buyers are not fooled just by setting  $\alpha = 1$ . When buyers are not fooled at all, firms may nevertheless still choose to game in response to the policy incentives we describe below.

Changing  $x$  or  $g$  incurs a cost. We assume that these costs are separable, and that changing  $x$  determines marginal cost, while changing  $g$  incurs a fixed cost. These assumptions about cost, which are convenient but not essential, are made in line with the market we study. Gaming is reported to be achieved largely through manipulation of the test vehicle, while real fuel consumption improvements require marginal costs.

We write the marginal cost of production for a product as  $c(x)$ , which is decreasing ( $x$  is a bad) and convex ( $c' < 0$ ,  $c'' > 0$ ). The cost of gaming is denoted  $h(g)$ , which is increasing and convex ( $h' > 0$ ,  $h'' > 0$ ). The cost of gaming can include both real resource costs (e.g., engineering costs to game the test) and the risk of regulatory penalty.

Policy intervention is motivated by a negative externality associated with  $x$ , which we assume is linear and equal to  $\phi x$ . The regulator observes  $m$ , so policy must be based on  $m$ . In line with our empirical application, we focus on a fuel economy regulation that requires that the (reported) attribute be below a threshold:  $\sigma: m = x - g \leq \sigma$ . We will use  $\lambda$  to denote the shadow price of the regulation *per unit*; i.e., the constraint on the firm's profit function is written  $\lambda \times (\sigma - x + g) \times D$ . We later compare this to a tax on fuel, which can be modeled as just a shift in  $\beta$ , which is tax-inclusive.

---

<sup>21</sup>Below we discuss in detail the information about fuel consumption and gaming available to buyers in our empirical setting, see Section 4.1.

### 3.2 Summary of seller behavior

The monopoly seller chooses  $p$ ,  $x$  and  $g$  to maximize profits, given the demand system, the degree of buyer sophistication, and policy interventions. The firm’s Lagrangean is:

$$\mathcal{L} = (p - c(x) + \lambda(\sigma - x + g))D(p + \beta x - (1 - \alpha)\beta g) - h(g). \quad (3)$$

What values will the seller choose? Price will be chosen according to the standard optimal markup formula for a monopolist that equates marginal revenue and marginal cost per unit sold (given the profit-maximizing  $x$  and  $g$ ).

The choice of  $x$  depends directly on policy. Absent policy ( $\lambda = 0$ ), the seller will choose  $x$  so that  $-c'(x) = \beta$ , which is privately efficient.<sup>22</sup> The seller lowers  $x$  until the marginal cost of doing so equals the marginal value of increased energy efficiency to the buyer. The seller makes this privately efficient investment because any improvement in energy efficiency that is cost effective allows them to raise prices.

When there is a binding policy, the seller will choose  $x$  to solve  $-c'(x) = \beta + \lambda$ . This is the standard “internalization” of the externality suggested by the Pigouvian tradition. If there were perfect compliance and there was no distortion to quantities sold due to market power, then the first-best solution would be obtained by setting  $\sigma$  so that  $\lambda = \phi$ .

The seller’s choice of gaming will depend on the cost function  $h(g)$ . Specifically, the seller’s choice will satisfy  $h'(g) = \lambda D - (p - c + \lambda)(1 - \alpha)\beta D'$ . This equates the marginal cost of gaming to the marginal benefit, which includes a regulatory benefit ( $\lambda D$ ) and a benefit from shifting out demand, which is equal to the markup  $(p - c + \lambda)$  times the shift in demand due to gaming  $(1 - \alpha)\beta D'$ .

In sum, firms have two benefits from gaming. If consumers are fooled (at least partly), then gaming increases demand. Even if consumers are completely aware of gaming, if there is a policy, firms may game in order to minimize compliance costs.<sup>23</sup>

### 3.3 Note on model assumptions

We model a single-product setting for tractability and clarity. The main drawback to this is that it does not directly represent compliance that occurs through mix shifting—that is, shifting sales towards high efficiency models in order to meet a fleetwide average rating—which is the focus of our simulations. The impact of mix shifting will manifest as a set of price wedges for vehicles that operates similarly to the cost wedge modeled here.

Our model assumes that consumers properly value fuel economy  $x$ , even if they are unaware of gaming  $\alpha$ . A literature on the energy efficiency gap debates whether consumers might undervalue fuel economy. If they do, then the regulation fixes an “externality” and improves *private* buyer

<sup>22</sup>The first-order condition for  $x$  is  $0 = -c'(x)D + (p - c(x))\beta D'$ . Substituting in the optimal markup from the first-order condition for price ( $p - c(x) = -D/D'$ ) yields the result.

<sup>23</sup>The performance gap in early years of our sample is consistent with gaming to fool consumers, but it is also possible that the test is simply not well calibrated to measure typical driving patterns in the Netherlands.

surplus. In this scenario, gaming will not have a beneficial price effect. We demonstrate this in our simulation results below.

We model a representative buyer, but if buyers differ in their  $\alpha$ , additional strategic considerations might come into play. For example, if relatively naïve buyers are also more price sensitive, then gaming could draw more price sensitive buyers into the market, which might put downward pressure on prices, thereby benefitting inframarginal, sophisticated buyers. On the other hand, with buyer heterogeneity of this sort and multiple products, a seller might use gaming to differentiate products to screen buyer types and achieve second degree price discrimination.

In addition, we assume that the costs of gaming are strictly fixed costs, but it is possible that they scale with sales, either because of production costs or regulatory risk (i.e., because legal penalties scale with sales). In such a model, the same key forces that we identify will still be at work, but the formula would presumably be different.

Finally, regulatory penalties might be depend on whether an individual firm games significantly more than the rest of industry, as a regulator may find it impractical to levy punitive damages on the whole industry simultaneously. For example, in the Volkswagen diesel case in the US, regulators forced quick removal of non-compliant vehicles from the market. Had the problem affected a majority of new vehicles, it is not plausible to imagine the same penalty. We suspect that part of the evolution of gaming was enabled by firm’s feeling insulated from penalty by their competitors behavior, but this interaction is not present in our monopoly model.

### 3.4 Absent policy, gaming lowers buyer surplus

Our main interest is in establishing how gaming affects buyer welfare in this setting. We show how the firm alters its choice of  $p$  and  $x$  for any given level of  $g$ , which allows us to then characterize buyer surplus. The optimal degree of gaming will depend on the regulatory punishment function, so this exercise can be thought of as examining how shifts in the punishment function that alter the equilibrium  $g$  will influence buyer surplus.

In our analysis we focus on buyer surplus, rather than full social welfare, because we believe it is the incidence of gaming about which there is an interesting theoretical question. If the policy is set optimally, gaming obviously lowers total social welfare, once buyer surplus, seller surplus and the externality are taken into account. This is true by definition. Accordingly, we use the term “buyer surplus” to denote the private benefits, abstracting from the externality, rather than the more encompassing term of consumer surplus.

Our first proposition shows that, when there are no corrective policies, an increase in gaming hurts buyers. (All proofs are in Appendix A.4.)

**Proposition 1.** *In the absence of a binding policy ( $\lambda = 0$ ), buyer surplus falls with the level of gaming. Specifically:*

$$\frac{dBS}{dg} \approx \underbrace{-\rho(1-\alpha)\beta D}_{\text{price effect}} + \underbrace{(1-\alpha)^2\beta^2 D'g}_{\text{choice distortion}} \leq 0.$$

The effect of gaming on buyer surplus has two components. The first is a *price effect*. Gaming lowers *perceived* fuel costs of the product, which shifts out demand by the perceived reduction in full price, equal to  $-(1 - \alpha)\beta$ .<sup>24</sup> The seller responds to this demand shift by raising prices. How much prices rise depends on the pass-through rate, which we denote  $\rho$ .<sup>25</sup> The pass through rate is typically used to describe how the burden of a tax is shared between buyers and sellers. An increase in gaming affects a shift in perceived demand that functions exactly like a tax to buyers, hence the analogy. This price effect is unambiguously bad for buyer surplus because the shift out in demand was a mistake—the buyer still has to pay the true operating cost of the vehicle, which is unchanged, but the seller has taken advantage of the mistake to raise prices.

The second effect on buyers of gaming is what we call the *choice distortion*. The buyer misperceives the true full cost  $f$  of the product due to gaming and thus purchases too much of the good given its true ownership cost. This misoptimization creates deadweight loss that is directly analogous to a Harberger triangle. Its width is the difference in demand, at the final price, induced by the gaming:  $D(p + \beta x - (1 - \alpha)\beta g) - D(p + \beta x)$ . Its height is the perceived gap in fuel cost induced by gaming:  $(1 - \alpha)\beta g$ . The term that appears in Proposition 1 is the derivative of this distortionary triangle. It is zero when  $g = 0$ ; it is rising with the square of the distortionary wedge; and it is larger when demand is more elastic.

Both the price effect and the choice distortion unambiguously hurt buyers who are fooled by gaming. The perceived improvement in the product causes sellers to raise prices, and the buyers buy too much of the good. When buyers are fully sophisticated ( $\alpha = 1$ ), gaming causes no change in buyer surplus; both terms go to zero. Both effects are larger when the degree of misperception  $(1 - \alpha)$  is larger. The price effect scales with the level of demand. The choice distortion scales with the slope of demand.

### 3.5 With a regulation, gaming can raise buyer surplus

Now suppose that there is a binding regulation ( $\lambda > 0$ ). Proposition 2 shows that the regulation (potentially) flips the sign of the price effect, so that buyers may in fact benefit from gaming, even when they are fooled by it.

**Proposition 2.** *In the presence of a binding standard ( $\lambda > 0$ ), a change in gaming affects buyer surplus as follows:*

$$\frac{dBS}{dg} \approx \underbrace{(-\rho(c' - \alpha\beta) - \beta)D}_{\text{price effect}} + \underbrace{(1 - \alpha)^2 \beta^2 D' g}_{\text{choice distortion}}.$$

As in the case with no regulation, the effect of gaming on buyer surplus includes a price effect and a choice distortion. The choice distortion formula is not affected by the policy, it is the same as in

<sup>24</sup>Note that gaming does not affect the choice of  $x$ , which satisfies  $-c'(x) = \beta$ .

<sup>25</sup>For a monopoly seller,  $\rho = (1 + (\epsilon_D - 1)/\epsilon_S + 1/\epsilon_{ms})^{-1}$ , where  $\epsilon$  is the elasticity of demand  $D$ , supply  $S$  and the inverse of the marginal surplus curve  $ms$ . Under another form of competition, the price effect will have the same formula but the pass-through rate will have a different functional form (Weyl and Fabinger 2013).



Proposition 1 and is unambiguously negative (lowers buyer surplus).

The price effect, however, is significantly altered by the externality-correcting regulation and may be positive. Gaming causes three effects that determine the full price ( $p + \beta x$ ). First, the seller will use an increase in  $g$  to reduce  $x$  under a binding standard (i.e.,  $dx = dg$  because  $\sigma = x - g$  when the standard binds). This raises the true full cost of ownership by  $-\beta$  per unit sold. Second, true cost of production falls by amount  $c'(x)$ . Some of this cost savings—equal to  $\rho c'$ —will be passed through to buyers via a reduction in  $p$ . Third, buyers will perceive a change in operating cost. This shifts in demand, and the seller will adjust  $p$  accordingly in reaction to the perceived operating cost increase. Specifically, with  $dx = dg$ , the buyer will perceive an operating cost change of  $\alpha\beta > 0$ , of which  $\rho\alpha\beta$  will be finally born by buyers.

The sign of the full price effect is ambiguous. The term  $-\rho(c' - \alpha\beta)$  is positive (it raises buyer surplus), but it may or may not be larger than the rise in true costs (equal to  $\beta$ ), which lowers buyer surplus. When the price effect is positive (buyers gain), it may be larger than the negative choice distortion, so that gaming benefits buyers. Intuitively, the benefit to buyers of gaming comes because the seller experiences a cost reduction when they avoid the costly regulation. By definition of a binding standard, the standard forces  $x$  to move away from the privately efficient level, so gaming provides a private welfare gain by allowing the seller to move back towards that level. Even if buyers are unaware of the firm's gaming, they will enjoy some of the cost savings due to standard pass through forces.

Several key comparative statics that we return to in our simulation analysis are apparent from Proposition 2. Empirically, we are particularly interested in establishing conditions under which buyers are likely to benefit from gaming. Factors that influence this outcome include the following:

1. More sophisticated (higher  $\alpha$ ) buyers gain more (or lose less) from gaming.
2. Tighter standards (more negative  $c'$ ) imply larger buyer gains (or smaller losses) from price effects.
3. More gaming (higher  $g$ ) implies smaller gains (or larger losses).

On the first point, more sophisticated buyers experience smaller choice distortions, and they will also get more favorable price effects.

On the second point, for a given amount of gaming, price effects get more positive (favorable to buyers) under a tighter standard because gaming creates a larger cost savings. By assumption of the convexity of  $c$ , a tighter standard will push the marginal cost  $c'$  to be larger (more negative). Gaming, which effectively relaxes the standard, therefore has a larger cost savings and subsequently a larger buyer price effect. For a sufficiently tight standard, the price effect will always be positive. This is true even if buyers are fully unaware of gaming.

On the third point, for a given standard, the shadow price of the regulation is falling as gaming increases by concavity of the profit function. In contrast, the choice distortion is rising in gaming, because the total choice distortion is a function of the square of  $g$ .

To recap, under a binding standard, gaming induces a price effect of ambiguous sign, and a choice distortion that lowers buyer surplus. When choice distortions are sufficiently small, positive price effects can dominate, so that increases in gaming raise buyer surplus, even when the buyers are fooled by the gaming. In the final part of the paper, we conduct calibrated simulations of the European vehicle market to show that it is indeed likely that car buyers experienced net benefits from gaming, even if they were completely fooled by it.

### 3.6 With a fuel tax, gaming lowers buyer surplus

A large literature compares the efficiency of fuel economy regulations to taxes on fuel (for a review, see [Anderson and Saltee 2016](#)). Because  $\beta$  is tax inclusive, if the planner uses a fuel tax instead of a fuel economy regulation to address the externalities related to  $x$ , the impact of gaming on buyer surplus is identical to the expression in Proposition 1—changing the tax just shifts  $\beta$ . Gaming does not help the buyer avoid the fuel tax  $\tau$ , so it provides them no benefits.

Normally, the statutory incidence of a tax is irrelevant to its ultimate economic incidence; i.e., it does not matter if a retail sales tax is levied on buyers or sellers. But, this irrelevance need not hold when there are opportunities for avoiding or evading the tax ([Slemrod 2008](#), [Kopczuk, Marion, Muehlegger, and Slemrod 2016](#)). Our result here is a manifestation of this same result. In our model, a regulation  $\sigma$  creates a shadow price  $\lambda$  that corrects the externality associated with  $x$ . A tax  $\tau$  does exactly the same thing. If there were no evasion, the two policies would be identical whenever  $\tau = \lambda$ , but gaming offers an opportunity to avoid a regulation but not a fuel tax, which breaks the symmetry and provides an enforcement rationale for preferring the tax over a regulation.

### 3.7 The environmental consequences of gaming

Theoretically, we are focused on the question of the incidence of gaming, but it is worth returning now to a discussion of the externality. The environmental impact of gaming depends upon the degree to which the planner adjusts policy to account for gaming, as well as the magnitude of what we call the mechanical and market size effects. For a given standard, gaming has a mechanical effect on pollution—gaming substitutes one for one for emissions reductions. Less obvious is the potential market size effect. If buyers are fooled by gaming, then gaming may lower the perceived price and the total market size will increase, which also increases pollution.<sup>26</sup>

This market size effect means that, on net, the introduction of a policy could *increase* pollution. For example, let  $x^b$  be the equilibrium attribute absent policy. A planner introduces a binding standard  $\sigma < x^b$ . Suppose that the firm complies exclusively through gaming. If buyers are partially naïve, they will perceive an improvement in the product. The net effect will be a perceived price reduction compared to the baseline, which leads to larger total demand than in the pre-policy baseline. Because the true attribute is unchanged from the baseline, the market expansion implies an unambiguous total increase in pollution caused by the policy. The final impact on emissions

---

<sup>26</sup>Note that, in the presence of market power, pollution increases due to an expansion of the market may not be welfare decreasing, in the spirit of [Buchanan \(1969\)](#).

depends on general equilibrium factors that are hard to predict. In a mature automobile market, an increase in new sales may lead primarily to a reduction in use of old vehicles. Vehicle retirement, however, might involve old cars being exported to lower-income markets.

Moreover, it is highly concerning for international climate negotiations that are structured like the Paris Accord, under which there may be considerable pressure for individual nations to deliver carbon reductions “on paper” that they do not actually wish to enforce. If lax enforcement leads to inflate energy efficiency statistics that partially fool buyers into perceiving cheaper durable goods, then this could act to raise total emissions.

Suppose instead that a regulator really does wish to lower emissions and is aware of gaming. If both the planner and buyers are aware of the degree of gaming, then the planner can simply ratchet down the standard to achieve the true desired level of emissions, accounting for gaming. In this case, gaming will still occur, but there is no choice distortion and the only inefficiency comes from resources spent to game, embodied in  $h(g)$ . To the extent that these are real resource costs, this is a social waste, but much of  $h(g)$  may represent regulatory risk.

More interesting then is the case when the planner is aware of gaming, but buyers have less than perfect sophistication ( $\alpha < 1$ ). In this case, the regulator can ratchet down the standard in order to achieve the desired level of pollution. But, following the logic of Goodhart’s Law, where tighter standards induce more gaming, tighter policies will induce a larger choice distortion. The choice distortion raises the social marginal cost of abating pollution, which will lead the regulator to attenuate the second-best policy away from the Pigouvian benchmark, under which the standard is tightened until  $\lambda$  equals marginal damages per unit of  $x$ .

## 4 Calibration of the incidence of gaming

To quantify the effects of gaming, we estimate a discrete choice model for new vehicles in order to obtain the taste parameters of consumers. We show that there is a direct analog to the choice distortion and price effects identified by our theory that can be decomposed from the standard log-sum welfare formula for discrete choice models. Given this decomposition and the estimated taste parameters, we calibrate outcomes under various assumptions about consumer awareness, policy stringency, and the degree of gaming. This allows us to show comparative statics that illustrate how the incidence of gaming depends on key parameters. We use parameters informed by the actual market context, but we stress that we show results for a range of parameters to illustrate how the incidence of gaming depends on these values.

Our theory describes a monopolist, but our simulation is broader and considers a multi-firm, multi-product setting. The main insights regarding the price effects are similar in either case, but with multiple products the choice distortion comes both from switches between products and from changes in the overall size of the market; that is, from the outside good. The final step of our simulation relates to competition and gaming. We show that regulation increases pressure on firms to game in this setting.

## 4.1 Setup and parameter choice

**Buyer preferences:** We begin by modeling a consumer making a discrete choice about which vehicle to buy. Each consumer  $i$  chooses the vehicle that maximizes her indirect utility, which we write as:

$$u_{ij} = \Delta_j \gamma_i - \eta p_j + \beta_i (\kappa_j * x_j) + \xi_j + \varepsilon_{ij},$$

where  $\Delta_j$  is a vector of vehicle characteristics,  $p_j$  is price divided by income and  $\kappa_j * x_j$  is the operating cost of the vehicle for driver  $i$ , measured as fuel consumption transformed into euros per kilometer using fuel prices  $\kappa_j$ .<sup>27</sup> Utility from the outside good (not buying a new vehicle) is normalized to zero,  $u_{i0} = \varepsilon_{i0}$ . We estimate a random coefficient logit model and assume that  $\beta_i$  and  $\gamma_i$  are independently normally distributed. We estimate the mean and standard deviation of  $\beta_i$  and  $\gamma_i$ . All remaining consumer heterogeneity is contained in the additive idiosyncratic error term  $\varepsilon_{ij}$ .<sup>28</sup> The  $\xi_j$  term represents the value of product attributes unobserved by the researcher but observed by firms and consumers. Following [Berry, Levinsohn, and Pakes \(1995\)](#), we integrate out  $\varepsilon_{ij}$  to construct choice probabilities. After inverting the choice probabilities to obtain  $\xi_j$ , we use a GMM estimator to estimate the vector of taste parameters.

Our dataset, a panel containing sales, prices and characteristics for all new vehicle sales in seven European countries between 1998-2011 allows us to estimate a rich demand model with flexible cross price elasticities. We only use data from before 2008 to estimate demand, so that our estimates come from a period in which the performance gap was stable. We thus assume variation in  $x_j$  is informative about actual fuel cost differences. Note that we assume that consumers have on average a correct belief about future fuel costs before gaming is introduced.<sup>29</sup> Estimated taste parameters and standard errors are reported in appendix Table A.7. Price and fuel costs have the expected negative effect on utility. The mean of the estimated own price elasticities is -5.45, in line with the previous literature. We find that cross price elasticities are higher for vehicles with similar fuel consumption and that the outside good is a stronger substitute for low fuel consumption vehicles. In the appendix we detail the estimation and we summarize substitution patterns in Table A.8.

**Profits:** In the empirical analysis we consider the impact of gaming with and without policy. The Lagrangean that the firm faces is equivalent to the one defined in (3) but firms optimize profits taking the products in their portfolio as given. Firms thus maximize profits by choosing price (for

<sup>27</sup>Notice there is a slight change of notation relative to the theory: here  $\beta$  captures the willingness to pay for fuel costs while in the theory  $\beta$  captures both the willingness to pay for fuel cost as well as the rescaling of fuel consumption into money. In the demand estimation we rescale fuel costs using observed fuel prices  $\kappa_j$ .

<sup>28</sup>We estimate the model with aggregated country level market shares and have no information on individual household incomes. Our model therefore does not allow us to study the distributional impacts of gaming. Related, we only estimate diagonal elements in the preference variance co-variance matrix and assume that there is no correlation between preferences for different attributes. We thus do not capture any correlation between price elasticities, fuel economy valuation or awareness of gaming ( $\alpha$ ), which means that strategic behavior based on heterogeneity in those parameters is absent from our simulations.

<sup>29</sup>If there were undervaluation of fuel costs, gaming could mitigate internality costs as described in [Allcott and Wozny \(2014\)](#), and as modeled in [Allcott and Taubinsky \(2015\)](#).

the vehicles they produce), taking levels of  $x$  as given:

$$\max_p \sum_{j \in F} [(p_j - c_j) q_j(p, g)] \text{ s.t. } \sum_{j \in F} q_j \cdot (x_j - g_j) / \sum_{j \in F} q_j = \sigma$$

in which  $F$  is the set of products a firm produces and  $q_j$  is quantity demanded, which depends on the full vectors of prices and gaming. Marginal costs  $c_j$  are not observed in the data. Because we estimate demand in a period where there is no regulation we are able to back out marginal costs from the unconstrained first order conditions. Given these marginal costs and estimated preferences we solve for new equilibrium prices with different levels of  $g$  and  $\sigma$  as we explain now.<sup>30</sup>

The policy, as stated in the constraint, is a requirement for each firm to decrease its sales-weighted emissions and mimics the EU policy.<sup>31</sup> When there is a binding policy, selling polluting products makes it harder for the firm to comply. Firms thus must adjust their price schedules or make costly adjustments to vehicle fuel consumption in order to reach the target. When we introduce gaming, these changes will be lowered because gaming makes the constraint less binding (it lowers the policy wedge  $\lambda$ ). Gaming also affects demand and optimal prices directly when consumers are at least partly fooled because demand shifts out in response to perceived higher quality.

**Policy target and compliance:** We model the policy as a requirement for each firm to decrease its sales-weighted emissions by 5%. This 5% is a limited decrease but it mimics a situation in which firms can potentially comply without changing characteristics of vehicles. We model honest compliance in two ways. First, we consider sales mixing: firms comply by changing prices to tilt demand towards efficient vehicles so as to raise their fleet average until compliant. This compliance channel raises median prices by €1860, a substantial amount. To justify a policy with these compliance costs, the social cost of carbon would need to be €286 per ton, which exceeds conventional estimates.

Second, we consider the addition of fuel saving technologies that raise marginal costs. In this version, we do not solve a compliance optimization problem, but instead simply add technologies to all models for a firm until they achieve compliance. Technology increases marginal costs, which raises prices in equilibrium. We take the marginal cost of fuel consumption improvements from a regulatory document based on engineering estimates conducted by the EU in studying the policy [TNO \(2011\)](#). Technology costs from this study are lower than private net present benefits from fuel savings. This is typical of engineering estimates of energy efficiency, which has given rise to a contentious debate around what is known as the energy efficiency gap ([Jaffe and Stavins 1994](#), [Allcott and Greenstone 2012](#)).

<sup>30</sup>Note that prices and estimated marginal costs vary across markets because we assume each consumer to shop in her own country. In the simulation we will use a single market: the Netherlands in 2007.

<sup>31</sup>We use the actual formula used by the EU in its emission standard, so that the policy is based on vehicle weight. For each firm we compute a sales weighted average emission rate, with emissions for heavier (lighter) cars receiving a bonus (penalty) in the weighted sum. Specifically, the emissions rating of each vehicle  $j$  in the weighted sum receives a penalty of  $0.046 * (weight_j - 1370)$  where weight is measured in kilograms.

We use the engineering costs in three scenarios inspired by interpretations of the energy efficiency gap. In the first scenario we take the engineering cost function at face value, even though the marginal costs changes needed to obtain a 5% improvements are very low. In this case the policy will be welfare enhancing: fuel consumption improvements are cheaper than buyer willingness to pay for these improvements. This is consistent with an interpretation of the energy efficiency gap that says there is an additional market failure and the policy corrects this failure. We stress that, in contrast to our model, gaming will not improve buyer welfare in this case because the regulation is raising *private* buyer surplus. Many argue that the energy efficiency gap is an artifact of hidden costs and optimistic engineering estimates. In our second scenario, we shift up the intercept of the cost function until marginal improvements in fuel consumption are equal to buyer willingness to pay for the improvements. In this scenario there is thus no scope for the policy to increase private surplus. In our third scenario, rather than shifting the intercept, we pick the point on the cost curve that equates marginal cost and marginal private benefits and assume that additional improvements must move up the curve from there.

We model separate scenarios that feature these two compliance strategies, which is consistent with our goal of exploring the quantitative significance of the channels impacting the incidence of gaming identified in our theory. A full welfare accounting would account for the optimal mix of those strategies, as well as attend to changes in fixed costs and vehicle redesign that we cannot identify in our model.<sup>32</sup> In terms of establishing how gaming affects buyers, as long as compliance strategies raise the net *private* cost of vehicles, either directly or through implicit price wedges, the basic insights of our model will follow, though the quantitative values will generally depend on what compliance options are available to firms. As we show below, if the regulation is privately beneficial because of another market failure, then our results change qualitatively.

**Level of gaming:** We model 5% gaming as a way to describe what the effects are of gaming when it completely offsets the policy. When we introduce gaming, firms that game can meet the policy requirement without a sales shift. In our simulations we also show what happens to buyer welfare when gaming is lower or higher than the policy target. Consistent with media coverage of the EU case, our procedure assumes that gaming is a fixed cost: it does not shift the marginal costs of the products. Buyers will perceive fuel costs as lower so that demand shifts out and prices increase. The market equilibrium will thus change to the extent that consumer demand, and hence prices and quantities, shift in response to gaming. Our procedure estimates the equilibrium prices and market shares for different levels of gaming, but it does not attempt to describe the optimal amount of gaming (this would require information on the costs of gaming). Rather, our goal is to use our demand system estimates to determine whether or not the price effects from gaming could plausibly dominate the choice distortion, leading to increases in consumer welfare even when consumers are naïve, and how to show how results hinge on the existence of a policy. This quantifies the main insights from our theory.

---

<sup>32</sup>For research indicating the importance of attribute changes, see [Klier and Linn \(2012\)](#), [Leard, Linn, and Springel \(2019\)](#). See [Reynaert \(2017\)](#) for a comparison of the welfare effects of different abatement strategies.

**Consumer awareness:** A crucial parameter is consumer awareness  $\alpha$ . When consumers are fully aware of gaming our model collapses to a simple setting where regulatory avoidance benefits both buyers and sellers. In our empirical setting it is unlikely that car buyers were aware of the full extent of gaming, but it is also unlikely that buyers did not update their beliefs over time.

We are not aware of any direct evidence relating to buyer knowledge of test gaming. There is, however, a substantial literature on fuel economy valuation that demonstrates that buyers respond to fuel price changes in ways that are consistent with them being attentive to test ratings and to the value of fuel savings.<sup>33</sup> That literature suggests that buyers pay attention to government test ratings, but it does not necessarily imply that buyers are aware of gaming or are able to correctly deflate government ratings in the presence of rapid declines in test accuracy.

Might consumers have been informed by the media? A news search suggests that the first reports about gaming appeared in the fall of 2013 in Germany. In other countries we find that the first newspaper articles appear only in 2014, by which time we estimate a fuel consumption gap in excess of 40%.<sup>34</sup> In 2015 there is a major news episode related to the Volkswagen crisis. This suggests that there was not EU-wide media coverage of the performance gap until 2015.

One other source of information are websites where consumers upload information about their experienced fuel consumption. Several such website exist during this time period. We note, however, that our empirical results demonstrate that there is significant heterogeneity across trips and drivers in experienced fuel consumption due to driving conditions. This makes it difficult for consumers to distinguish legitimate variation from gaming without a large sample. Thus, we conclude that consumers might have had some information, but the available information was far from perfect.

We have tried to directly estimate consumer awareness of gaming with our data, but did not find statistically meaningful results. The challenge is that our estimate of on-road consumption is a function of gaming and selection, and the lack of heterogeneity in gaming across models within a year (as shown in Figure 4) makes it difficult to isolate meaningful variation between vehicles that can be used to explain market shares.<sup>35</sup>

In what follows we show how the price effect and choice distortion vary when consumers range

---

<sup>33</sup>A recent literature (see Busse, Knittel, and Zettelmeyer 2013, Allcott and Wozny 2014, Sallee, West, and Fan 2016, Grigolon, Reynaert, and Verboven Forthcoming) has shown that buyers do take fuel consumption ratings into account when purchasing a vehicle. In a field experiment Allcott and Knittel (2017) show that treating buyers with detailed information about fuel consumption has no impact on their choice.

<sup>34</sup>In particular, we perform a Google news search with the words “ICCT” and “gap” in German, English and Dutch. The ICCT reports were discussed in the German media in 2013 and 2014, but we find no references in other languages, except for two automobile websites in English. Other word combinations gave no additional results.

<sup>35</sup>In particular, we did the following. We regressed our vehicle level gap estimates on a set of explanatory variables and predicted gaps for vehicles in the choice set for which we have no on-road estimates. We then tested if differences in gaps explain differences in market shares by introducing the estimated gaps in our discrete choice framework. This resulted in very noisy estimates with a wide confidence interval for the effect of the gap on sales. In a second stage, we tried to control for selection, using the assumption that gaming is absent prior to 2007. This allows us to decompose the gaps post 2007 into gaps from driver selection and gaming. When we introduce this measure in the discrete choice framework we find no statistically (or economically) meaningful estimates. One key limitation is that our sales data end in 2011, which excludes the period of greatest interest. A related exercise that tests for changes in pricing due to gaming similarly yielded only noisy results. We thus cannot make any statistical statement about consumer awareness of gaming based on our data.



from completely naïve to fully sophisticated ( $\alpha$  ranges from 0 to 1). We argue that in the long term buyers will update their beliefs, but in the short term consumers may be distorted in their choices by gaming.

**Buyer surplus** To compute buyer welfare when consumers are affected by gaming ( $\alpha < 1$ , in terms of our theory) we follow [Dubois, Griffith, and O’Connell \(2018\)](#), who describe the welfare effects of persuasive advertising. As described in Section 3, gaming changes prices, as firms will exploit higher demand from lower perceived fuel costs. We label the equilibrium price without gaming  $p_j^0$ , and the prices with gaming as  $p_j^1$ . The misperception of gamed fuel costs also distorts consumer choice. To separate the two effects we make a distinction between decision (at the moment of purchase) and experience utility (at the moment of utilization).<sup>36</sup> When there is gaming, a naïve consumer will perceive fuel costs, following her belief, as  $\tilde{x}_j = x_j - g_j$ , in which  $g_j$  is the increase in gaming. The consumer will make her choice based on  $\tilde{x}_j$  and will perceive her decision as yielding a utility of:

$$\tilde{V}_{ij}(g, p^1) = \Delta_j \gamma_i - \eta p_j^1 + \beta_i \tilde{x}_j + \xi_j + \varepsilon_{ij}.$$

After purchasing the vehicle, true fuel costs are revealed, and the consumer has experience utility:

$$V_{ij}(g, p^1) = \Delta_j \gamma_i - \eta p_j^1 + \beta_i x_j + \xi_j + \varepsilon_{ij}.$$

The difference between decision and experience utility is the optimization error, which in this case is the value of the additional fuel costs for the consumer  $\beta_i(x_j - \tilde{x}_j) = \beta_i g_j$ . Buyer surplus with gaming can then be written as:

$$\begin{aligned} \tilde{W}_i(g, p^1) &= \mathbb{E}_\varepsilon[\tilde{V}_{ij}] - \mathbb{E}_\varepsilon[\beta g_j] \\ &= W_i(g, p^1) - \sum_j [s_{ij} \beta g_j], \end{aligned}$$

where  $s_{ij}$  are the choice probabilities obtained from maximizing the decision utility. We compute  $W_i(g, p^1)$  by applying the log-sum formula of [Small and Rosen \(1981\)](#). We can then decompose the change from the equilibrium with honesty ( $0, p^0$ ) to the equilibrium with gaming ( $g, p^1$ ) into a price effect and choice distortion:

$$W_i(0, p^0) - \tilde{W}_i(g, p^1) = \underbrace{W_i(0, p^1) - \tilde{W}_i(g, p^1)}_{\text{Choice Distortion}} + \underbrace{W_i(0, p^0) - W_i(0, p^1)}_{\text{Price Effect}}.$$

When there is no policy the reference prices  $p^0$  are the observed prices in the market. When we introduce gaming without policy, prices  $p^1$  will be higher than  $p^0$  because of market power. When we introduce the policy, the reference prices  $p^0$  become the prices obtained from honest compliance

---

<sup>36</sup>This is conceptually the same as the choice distortion described in [Allcott \(2013\)](#) and [Sallee \(2014\)](#) who study misperception of fuel costs, but our approach here is more general because we allow firms to change prices in response to gaming. [Leggett \(2002\)](#) also models a similar distortion in a discrete choice setting.

to the policy. These prices will differ from the no-policy reference prices because of the tax wedge that the policy introduces ( $\lambda > 0$ ). When we introduce gaming, the prices  $p^1$  will change for two reasons: gaming lowers the tax wedge and firms with market power increase prices when gaming. In contrast, the choice distortion will be similar with and without policy as it is not a function of the tax wedge.

Note that consumer choice is distorted by gaming even when all products are gamed an identical amount. Intuitively, a uniform reduction in fuel costs is akin to a reduction in the coefficient on fuel costs in the utility function. Thus, when all firms game and perceived fuel costs drop, consumers will, to some degree, reorder their favored products because they effectively put more weight on other vehicle characteristics.<sup>37</sup>

**Procedure:** In all simulations, we start with a base market, which is the observed market in the Netherlands in 2007. That is, we use several markets over several years to estimate the demand system, but we need to choose one particular market as our baseline from which to calculate changes. We compute consumer surplus from this observed market and our parameter estimates. To simulate gaming, we first lower the perceived fuel cost of the products made by firms that game. Second, with these new fuel costs, we solve for demand and prices using the first-order conditions for profit maximization. Third, we compute changes in welfare relative to the reference market with prices  $p^0$ . The obtained changes in welfare are the changes in yearly utility, profits and emissions from new vehicle sales.

## 4.2 Incidence of gaming

In this section we simulate the effect of gaming on buyer surplus and pollution. We start from a base scenario in which we impose a policy of 5%, with honest compliance coming from mix shifting, and gaming equal to 5%. We focus first on the role of consumer awareness of gaming before discussing policy stringency, alternative compliance, and the level of gaming.

**Awareness:** In Figure 5 we plot buyer welfare changes from 5% gaming when there is no binding policy and when there is a 5% policy for various values of consumer awareness of gaming  $\alpha$ . When there is no policy both the choice distortion and the price effect are negative: gaming hurts buyers. Both the choice distortion and the price effect go to zero when consumer awareness  $\alpha$  increases from 0 to 1. When we introduce the policy, the choice distortion remains very similar and again goes to zero when consumer awareness increases. The price effect, however, switches sign and becomes positive and an order of magnitude larger than the choice distortion. This causes overall buyer surplus to *increase* with gaming. The intuition for the switch in the price effect is that cost avoidance of regulatory compliance is passed through from seller to buyer. Notice that this

---

<sup>37</sup>Note that we limit ourselves to reporting changes in consumer welfare from purchasing a new vehicle. We do not consider effects on the second hand market or on the amount that consumers drive. Consumers could potentially react to gaming by changing the amount they choose to drive in the purchased vehicle, a potential effect that we also abstract from.

positive price effect is very similar regardless of consumer awareness; pass through occurs whether or not consumers are fooled. The negative effects from gaming via choice distortion increase when consumers are fooled, but even when consumers are fully unaware of the gaming, the price effect strongly dominates.

**Main result:** Having established how results depend on  $\alpha$ , we now explore our main results assuming that consumers are unaware ( $\alpha = 0$ ), which is meant to be a worst-case scenario for buyers and thus conservative against our findings of buyer gains. Table 2 shows the effect of gaming on buyer surplus and pollution when consumers are unaware ( $\alpha=0$ ). Column I is the no policy case. We show how gaming changes buyer surplus, decomposed into both a choice distortion and a price effect, both of which are negative when there is no policy. Column II shows the case where there is a policy and compliance comes from mix shifting. As predicted by theory, the sign of the price effect flips, and the price effect is estimated to be an order of magnitude larger than the choice distortion. On net, gaming increases buyer surplus when there is a policy.

Table 2 also shows the effect of gaming on pollution. The first line shows the percentage-point change in total sales, which captures how the size of the market changes as a result of gaming. The second line shows the percentage change in average fuel consumption, which captures to what extent buyers substitute toward more polluting vehicles as a result of gaming. When there is no policy (Column I), we see that gaming increases the market size by 3% and that consumers buy vehicles that are slightly more polluting (0.44%). When a policy is in place (Column II), these effects increase considerably to a 4.6% increase in market size and a 6.3% increase in average fuel consumption. This is because the reference point is honest compliance with the policy with a target of 5% reduction in average fuel consumption. Not only does gaming undo the target (the mechanical effect), but consumers also buy more vehicles and shift toward more polluting vehicles when they are fooled by the gaming (the market size effect).

Proposition 2 in our theory showed that it was possible that consumers benefit from gaming, even when they are fooled by it. Our simulation results suggest that this is indeed the case in the European auto market. In line with Proposition 1, we find that consumers unambiguously lose from gaming when there is no policy. But, this result is overturned with the introduction of a policy. Consumers no longer benefit from honesty. Moreover, when choice distortions are significant, emissions actually *rise* in response to the policy, due to a combination of weak enforcement and significant shifts in the market due to choice distortions. We believe this raises an important concern for policies, like current international climate negotiations, that rely on voluntary on-paper emissions pledges rather than effective and credible measurement.

**Compliance mechanisms:** In the second panel of Table 2 we change the compliance mechanism from sales mixing to technological change that increases marginal costs. There are three different scenarios consistent with our discussion above. In Column VI we show the case where compliance costs are lower than consumer willingness to pay for fuel consumption. In this case the policy is understood to solve a market failure related to the provision of energy efficiency (e.g., consumer

**Figure 5:** Consumer awareness and incidence of gaming




Figure plots buyer surplus, choice distortion and price effect in Euro when consumers range from unaware ( $\alpha = 0$ ) to aware ( $\alpha = 1$ ). In panel (a) we model 5% gaming and no policy, in panel (b) we model 5% gaming with a 5% policy.

**Figure 6:** Level of Gaming




Figure plots buyer surplus, choice distortion and price effect in Euro when gaming ranges from 0% to 10% in response to a 5% policy.

**Table 2:** Effects of 5% Gaming on Buyer Surplus and Pollution

Scenario: Policy:	Sales Mixing			
	I 0%	II 5%	III 5%	IV 5%
Buyer Surplus (€per household)				
Total BS	-24.05 [-24.21,-15.32]	215.03 [214.08,245.51]	193.08 [192.01,223.23]	629.40 [628.58,725.13]
Choice Dist.	-17.07 [-17.17,-13.34]	-16.70 [-16.81,-13.07]	-30.68 [-30.82,-24.84]	-8.81 [-8.85,-7.25]
Price Effect	-6.98 [-7.09,0.69]	231.73 [231.21,262.82]	223.76 [223.18,253.65]	638.20 [636.86,734.80]
Pollution (% change)				
Sales	3.05 [3.05,3.54]	4.61 [4.60,5.31]	4.55 [4.54,5.26]	0.51 [0.51,0.58]
Emissions	0.44 [0.44,0.49]	6.27 [6.27,6.32]	6.64 [6.63,6.72]	5.68 [5.68,5.701]
Fuel Efficiency Improvements				
Scenario: Policy:	V 0%	VI 5%	VII 5%	VIII 5%
	Buyer Surplus (€per household)			
Total BS	-24.05 [-24.21,-15.32]	-375.29 [-376.46,-324.96]	72.74 [72.17,230.62]	126.09 [125.56,149.97]
Choice Dist.	-17.07 [-17.17,-13.34]	-16.90 [-17.02,-13.21]	-16.78 [-16.85,-13.14]	-15.59 [-15.65,-12.26]
Price Effect	-6.98 [-7.09,0.69]	-358.39 [-359.05,-310.03]	89.51 [89.37,244.79]	141.66 [141.46,167.80]
Pollution (% change)				
Sales	3.05 [3.05,3.54]	0.24 [0.24,0.27]	3.82 [3.78,4.95]	4.28 [4.27,5.05]
Emissions	0.44 [0.44,0.49]	5.26 [5.26,5.26]	5.70 [5.69,6.28]	5.40 [5.40,5.43]
Increasing Policy Stringency				
Scenario: Policy:	IX 0%	X 3%	XI 5%	XII 10%
	Buyer Surplus (€per household)			
Total BS	-24.05 [-24.21,-15.32]	16.18 [16.02,31.32]	215.03 [214.08,245.51]	614.59 [613.22,697.70]
Choice Dist.	-17.07 [-17.17,-13.34]	-17.75 [-17.88,-13.86]	-16.70 [-16.81,-13.07]	-13.46 [-13.49,-10.57]
Price Effect	-6.98 [-7.09,0.69]	33.93 [33.78,48.29]	231.73 [231.21,262.82]	628.05 [626.91,712.60]
Pollution (% change)				
Sales	3.05 [3.05,3.54]	3.12 [3.10,3.62]	4.61 [4.60,5.31]	8.29 [8.28,9.40]
Emissions	0.44 [0.44,0.49]	6.25 [6.25,6.30]	6.72 [6.27,6.32]	5.84 [5.83,5.90]

Table gives changes in buyer surplus and pollution from gaming for unsophisticated consumers ( $\alpha = 0$ ). Bootstrapped 95% CI reported between brackets. Top panel: (I) 5% gaming when no policy is in place; (II) 5% gaming when a 5% policy is in place; (III) uniformly distributed 0-10% gaming when a 5% policy is in place; (VI) 5% gaming when a 5% policy is in place and when outside good share is 10% instead of 61%. Middle panel shows results when firms respond to policy by increasing fuel efficiency when (VI) technology is low cost, (VII) technology cost function is flat and (VIII) technology cost function is convex, Bottom panel increases the policy stringency from 0% to 10% while keeping gaming constant at 5%. (Note that I=V=IX and VII=II.)

undervaluation), and the policy raises private buyer surplus. As a result, gaming hurts consumers in this scenario, not only because of the choice distortion, but also because buyers miss out on the positive price effects of the policy (the price effect of gaming is thus negative). Column VII and VIII show the case where compliance costs are high enough so that the policy does not increase buyer surplus. In both columns we find price effects to be larger than choice distortions, and the more convex compliance costs in VIII increase the price effect. These results are qualitatively similar to the results based on mix shifting, though they are somewhat muted.

**Stringency of policy and level of gaming:** In the third panel of Table 2 we vary the stringency of the policy from 3% to 5% to 10%, while keeping the degree of gaming constant at 5%. In all cases, the price effect is positive and larger than the choice distortion. The price effect is only marginally larger than the choice distortion for the 3% target, but the price effect outweighs the choice distortion by hundreds of euros when the standard becomes more stringent. This is in line with the theory; the price effect is increasing with the stringency of the regulation. This is relevant for assessing the EU standard, which is arguably the most demanding fuel consumption standard in the world.

Figure 6 shows buyer surplus and its decomposition when we increase the level of gaming from 0% to 10% in response to a 5% policy. The choice distortion hurts consumers more and more as we increase gaming. The price effect is a parabola because of two off-setting effects. The curve is steepest with the first units of gaming as these soften the burden of the policy the most. With the following units of gaming, up until 5%, more and more of the policy is avoided but this gives smaller incremental increases. At the same time higher gaming increases prices because of market power. This effect becomes dominant once the policy is fully avoided and thus the price effect decreases once more gaming is done than required to comply with the policy. The combination of the decreasing choice distortion and parabolic price effect gives an overall buyer surplus that is also parabolic. Note that gaming has positive effects on buyer surplus until 9%, almost twice the level required to comply with policy.

In appendix Table A.10 we also show results when the policy is a fuel tax instead of a standard.<sup>38</sup> This confirms that gaming with a fuel tax is harmful to buyers, consistent with our discussion in section 3.6. The table shows results from fuel taxes that were found to produce the same change in fuel consumption as the standards. We find that gaming decreases buyer surplus in very similar magnitudes to our base specification when firms game in the absence of a policy. Buyers will not benefit from gaming in response to a fuel tax because gaming will distort their choice; the price effect will lead them to pay higher, not lower, prices for the vehicle; and gaming does nothing to help them avoid the fuel tax.

**Heterogeneity in Gaming and the Outside Good:** The choice distortion drives the negative effects of gaming on buyer surplus and pollution. Because fuel consumption ratings are shaded

---

<sup>38</sup>We also show results for a flat emission standard that is not attribute based. This flat standard is less demanding for most firms and so the price effects are smaller, but still larger than the choice distortion

by the gaming, buyers purchase the wrong amount and the wrong type of vehicle. In Table 2 columns I-II and IX-XII, we see large increases in sales (between 3.1% and 8% points) and limited increases in average fuel consumption per vehicle (between 0.8%-1.7% on top of the mechanical 5%). Gaming thus leads to considerable changes in the size of the market and in the type of products purchased. The size of these two effects, however, depends partly on two underlying assumptions in our simulations: the lack of heterogeneity in gaming and the amount of substitution to the outside good.

In Column III of Table 2 we relax the assumption of homogenous 5% gaming on all products and instead draw gaming from the uniform distribution  $[0\%, 10\%]$ , so that mean gaming is 5% as in the main scenario. Now gaming not only reduces the importance of fuel consumption to choice relative to other attributes, but also scrambles the ranking of preferred products more directly by providing a bigger improvement to some products than to others. This has two implications relative to the homogeneous 5% gaming in Column II. First, we see that the choice distortion almost doubles from €17 to €31. Because gaming is heterogeneous, buyers buy the wrong type of cars more often and end up further from their preferred vehicle. Nevertheless, despite the doubling of the choice distortion, the price effect still dominates, and buyers are estimated to gain on net. Second, the environmental damage due to gaming is larger and average emissions now increase by 1.6%. This is because consumers substitute more toward high quality and high pollution vehicles that were gamed relatively more.

In Column IV we change the size of the outside good in our estimation and in our simulations. The estimated demand parameters do not change significantly because we include market fixed effects, but the results of the simulation do change noticeably. In Column II we obtained numbers from a model in which we assume that one out of seven households are potentially interested in purchasing a new car each year such that the outside good share is 61%. This assumption determines the price elasticity of the industry. When we lower the price elasticity of industry by decreasing the outside good share to 10%, we find significant changes. First, the positive price effects increase significantly because buyers do not stop purchasing polluting vehicles even at much higher prices. The price changes needed for firms to comply thus increase and the benefits of avoiding the regulation are larger. Second, the choice distortion shrinks. Because the price elasticity of industry is much lower in this scenario, gaming changes the market size by less. This is also important for pollution. The size of the market increases by 0.5% instead of 4.5% points, though buyers still substitute to more polluting vehicles on average. In sum, the positive price effect increases and the choice distortion decreases when the outside good share shrinks.

### 4.3 Competition and gaming

So far we have assumed that the whole industry games, in line with the empirical evidence shown in Section 2.5. Here, we study what happens if not all firms game. From the point of view of a single firm, gaming confers a competitive advantage. This is much the same intuition as for the monopolist; the residual demand curve faced by the firm can be shifted out by gaming. Conversely,



**Table 3:** Market changes from asymmetric gaming

	Gaming Alone		Not Gaming Along	
	No Policy	Policy	No Policy	Policy
	Consumer Surplus (€per household)			
Total	-3.83	15.79	-24.31	189.94
Choice Distortion	-3.17	-3.24	-17.23	-16.79
Price Effect	-0.66	19.20	-7.08	206.7
	Changes in pollution			
Total Sales (% point)	0.27	0.42	2.81	4.24
Emissions per vehicle %	0.04	0.62	0.41	5.45
	Changes in Profits			
Gamer/non gamer %	20.41	37.94	-9.86	-15.83

Column (1) and (3) give changes in market equilibrium when one firm games. We present the average of letting each firm game (we run a separate simulation per firm). Changes in profits are the average for the firms that game (not for the industry). Column (2) and (4) give changes in market equilibrium when all but one firm games. We run a separate simulations for each firm not gaming along when all others game. Changes in profits are the average for the firms that do not game along.

when other firms are gaming, an honest firm would be at a competitive disadvantage.

In Table 3 we compute market outcomes with asymmetry in gaming. In the left panel only one firm games, while all other firms remain honest. In the right panel, labeled “not gaming along,” all but one firm games. We simulate these asymmetric cases separately for each firm and present the average outcomes over these cases. In all scenarios we find results that are qualitatively similar to our industry-wide symmetric gaming case in that gaming hurts buyers when there is no policy, while buyers benefit from gaming under the policy. Buyer surplus and pollution effects are much smaller when only one of the carmakers games. Effects are almost equal to symmetric gaming when all but one firm games.

Table 3 also reports how profit changes with gaming. In the left panel we give the gains in profits for the single firm that games (the gamer), in the right panel we present the losses in profits for the single firm that does not game along (the non-gamer). The average gain from gaming unilaterally is large: a 20% increase in profits, even in the absence of policy. But the profitability from gaming rises still further, to 38%, with policy. This shows that firms have a very strong incentive to unilaterally deviate from honesty, and this incentive rises with policy. The profit losses from not “gaming along” are also large. Being the sole honest firm leads to a 10% decrease in profits without policy and a 16% decrease with policy. Thus, if competitors start to game, it is very costly not to follow along, and this incentive is enhanced by the policy.

We conclude that the policy thus clearly increases the incentive to game in a multi-firm context: the unilateral profits from deviating from honesty increase and the costs from not deviating when

others do increases.<sup>39</sup> This is related to an existing literature that asks whether competition induces gaming (e.g., [Becker and Milbourn 2011](#), [Bennett, Pierce, Snyder, and Toffel 2013](#)), but we add to that the suggestion that corrective policy amplifies these competitive effects.

## 5 Conclusion

This paper does three things. First, it documents a remarkable and growing divergence between official and on-road fuel consumption rates, which we call the performance gap, in a reduced-form empirical analysis using novel data. This performance gap is coincident with the introduction of strict standards based on the official ratings, and we interpret this pattern as evidence of Goodhart’s Law, whereby policy-induced stakes lead to the strategic manipulation of a measure. Combined with the still-growing Volkswagen scandal regarding local air pollution, the performance gap in carbon emissions points to a veritable crisis in the administration of environmental regulations for automobiles. Moreover, combined with other findings that call into question the efficacy of other energy efficiency policies ([Davis, Fuchs, and Gertler 2014](#), [Fowlie, Greenstone, and Wolfram Forthcoming](#), [Levinson 2016](#)), we believe our findings point to a serious challenge for the future of climate negotiations, which are now built upon bottom-up pledges based on unreliable ex ante estimates of program conservation.

Second, we develop a theoretical framework in order to ask how such policy-induced gaming of energy-efficiency ratings will impact buyers. The model points out how the existence of an externality-correcting policy fundamentally alters the incidence of gaming on buyers. Absent policy, gaming hurts consumers through both the price effect and choice distortion channels. But, when gaming is done to avoid a regulation, consistent with Goodhart’s Law, then buyers can benefit via pass through of cost savings. We show that beneficial price effects can dominate so that private buyer surplus rises with gaming.

Third, we conduct a structural analysis and calibrated welfare simulation in order to quantify the price effects and choice distortions identified by theory. For a wide range of parameter values, we find that price effects do indeed dominate so that the theoretical result that consumers might benefit from gaming is not just a curious possibility, but in fact appears to be a likely probability.

Our analysis emphasizes private buyer surplus, abstracting from the externality. Assuming that the policies are optimally designed and the externality benefits accrue to buyers, the buyers would be better off when the seller complies honestly, once the externality is taken into account. But, in our empirical context, neither condition is likely to hold. Analysis of the fuel economy standards in the US typically finds that the standards do not pass a cost-benefit test unless buyers, contrary to recent empirical evidence, have large biases in their valuation of fuel economy. Moreover, the climate benefits of reduced emissions largely accrue to future generations and to people other than

---

<sup>39</sup>Note that this setting is not necessarily a prisoner’s dilemma. The total industry profit from the price equilibrium with gaming relative to the price equilibrium with honesty will depend on the total legal and regulatory costs of gaming relative to the total costs of compliance. The evidence presented here shows that honesty becomes harder to support when the policy is in place.

European car buyers. However, co-benefits from local air pollution in Europe could be substantial given the large fraction of diesel vehicles, which might improve the cost-benefit analysis of the fuel-consumption standards there.

We anticipate that the same concerns we highlight here are present in a variety of regulatory settings, and we expect that both our theoretical and empirical approach can be applied elsewhere. Critical to our setup is that the measure which can be gamed is used by the regulator *and* is instrumental to consumer choice. For example, in education, test scores of current students might be gamed to satisfy policy mandates that affect resource transfers. The test scores also affect the demand for admission from future waves of students. Consumers (students) can benefit from gaming when gaming frees up resources at a school that would otherwise be expended to satisfy policy and obtain transfers. Or, in finance, capital requirements can be used to minimize an individual bank's role in creating systemic risk. Customers can privately benefit if a bank is able to rebalance its portfolio more profitably after gaming the regulated measures.

## References

- Allcott, Hunt. 2013. "The Welfare Effects of Misperceived Product Costs: Data and Calibrations from the Automobile Market." American Economic Journal: Economic Policy 5 (3):30–66. URL <http://www.aeaweb.org/articles.php?doi=10.1257/pol.5.3.30>.
- Allcott, Hunt and Michael Greenstone. 2012. "Is There an Energy Efficiency Gap?" Journal of Economic Perspectives 26 (1):3–28.
- Allcott, Hunt and Christopher Knittel. 2017. "Are Consumers Poorly Informed about Fuel Economy? Evidence from Two Experiments." NBER Working Paper 23076.
- Allcott, Hunt and Dmitry Taubinsky. 2015. "Evaluating Behaviorally Motivated Policy: Experimental Evidence from the Lightbulb Market." American Economic Review 105 (8):2501–2538.
- Allcott, Hunt and Nathan Wozny. 2014. "Gasoline Prices, Fuel Economy, and the Energy Paradox." Review of Economics and Statistics 96 (5):779–795.
- Anderson, Soren T. and James M. Sallee. 2011. "Using Loopholes to Reveal the Marginal Cost of Regulation: The Case of Fuel-Economy Standards." American Economic Review 101 (4):1375–1409.
- . 2016. "Designing Policies to Make Cars Greener: A Review of the Literature." Annual Review of Resource Economics 8:157–80.
- Austin, David and Terry Dinan. 2005. "Clearing the Air: The Costs and Consequences of Higher CAFE Standards and Increased Gasoline Taxes." Journal of Environmental Economics and Management 50 (3):562–582.
- Becker, Bo and Todd Milbourn. 2011. "How Did Increased Competition Affect Credit Ratings?" Journal of Financial Economics 101:493–514.
- Ben-David, Itzhak. 2011. "Financial Constraints and Inflated Home Prices during the Real Estate Boom." American Economic Journal: Applied Economics 3:55–87.

- Bennett, Victor Manuel, Lamar Pierce, Jason A. Snyder, and Michael W. Toffel. 2013. "Customer-Driven Misconduct: How Competition Corrupts Business Practices." Management Science 59 (8):1725–1742.
- Bento, Antonio M., Kenneth Gillingham, and Kevin Roth. 2017. "The Effect of Fuel Economy Standards on Vehicle Weight Dispersion and Accident Fatalities." Working Paper: Yale University.
- Bento, Antonio M., Lawrence H. Goulder, Mark R. Jacobsen, and Roger H. von Haefen. 2009. "Distributional and Efficiency Impacts of Increased US Gasoline Taxes." American Economic Review 99 (3):667–699.
- Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." Econometrica 64 (4):841–890.
- Booz & Company. 2012. "Company Cars and Fring Benefit Tax: Understanding the Impacts on Strategic Transport Targets." NZ Transport Agency research report 474.
- Buchanan, James M. 1969. "External Diseconomies, Corrective Taxes, and Market Structure." American Economic Review 59 (1):174–177.
- Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer. 2013. "Are Consumers Myopic? Evidence from New and Used Car Purchases." American Economic Review 103 (1):220–256.
- Card, David and Thomas Lemieux. 2001. "Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis." The Quarterly Journal of Economics 116 (2):705–746. URL <https://ideas.repec.org/a/oup/qjecon/v116y2001i2p705-746..html>.
- Chetty, Raj, John N. Friedman, and Johan E. Rockoff. 2014. "Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates." American Economic Review 104 (9):2593–2632.
- Davis, Lucas W., Alan Fuchs, and Paul Gertler. 2014. "Cash for Coolers: Evaluating a Large-Scale Appliance Replacement Program in Mexico." American Economic Journal: Economic Policy 6 (4):207–238.
- Dee, Thomas S., Will Dobbie, Brian A. Jacob, and Johan E. Rockoff. 2016. "The Causes and Consequences of Test Score Manipulation: Evidence from the New York Regents Examinations." NBER Working Paper 22165.
- Drugov, Mikhail and Marta Troya-Martinez. 2015. "Vague Lies and Lax Standards of Proof: On the Law and Economics of Advice." SSRN Working Paper.
- Dubois, Pierre, Rachel Griffith, and Martin O’Connell. 2018. "The Effects of Banning Advertising in Junk Food Markets." The Review of Economic Studies 1 (1):396–436. URL <https://dl.dropboxusercontent.com/u/21832052/site/dgo.pdf>.
- Durrmeyer, Isis and Mario Samano. Forthcoming. "To Rebate or not to Rebate: Fuel Economy Standards vs. Feebates." Economic Journal URL [http://docs.wixstatic.com/ugd/23918f\\_a6dbf5dcc28046cabd57be5b8528b91e.pdf](http://docs.wixstatic.com/ugd/23918f_a6dbf5dcc28046cabd57be5b8528b91e.pdf).

- Fabra, Natalia and Mar Reguant. 2014. "Pass-Through of Emissions Costs in Electricity Markets." American Economic Review 104 (9):2872–2899.
- Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram. Forthcoming. "Do Energy Efficiency Investments Deliver? Evidence from the Weatherization Assistance Program." Quarterly Journal of Economics .
- Frankel, Alex and Navin Kartik. Forthcoming. "Muddled Information." Journal of Political Economy .
- Gabaix, Xavier and David Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." Quarterly Journal of Economics 121 (2):505–540.
- Gerlagh, Reyer, Inge van den Bijgaart, Hans Nijland, and Thomas Michielsen. Forthcoming. "Fiscal policy and CO2 emissions of new passenger cars in the EU." Environmental and Resource Economics .
- Gillingham, Kenneth. 2013. "The Economics of Fuel Economy Standards versus Feebates." Working Paper: National Energy Policy Institute.
- Goldberg, Pinelopi Koujianou. 1998. "The Effects of the Corporate Average Fuel Efficiency Standards in the U.S." The Journal of Industrial Economics 46 (1):1–33.
- Goodhart, Charles. 1981. Problems of Monetary Management: The U.K. Experience. Springer.
- Grigolon, Laura, Mathias Reynaert, and Frank Verboven. Forthcoming. "Consumer Valuation of Fuel Costs and the Effectiveness of Tax Policy: Evidence from the European Car Market." American Economic Journal: Economic Policy URL <http://docs.google.com/viewer?a=v&pid=sites&srcid=ZGVmYXVsdGRvbWVfbnxbmcmFua3ZlcmJvfGd40jZiMzQ4YzUwYjhiNjdjOTk>.
- Heidhues, Paul, Botond Köszegi, and Takeshi Murooka. 2017. "Inferior Products and Profitable Deception." Review of Economic Studies 84:323–356.
- Ito, Koichiro and James M. Sallee. 2018. "The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel-Economy Standards." Review of Economics and Statistics 100:319–336.
- Jacobsen, Mark R. 2013. "Fuel Economy and Safety: The Influences of Vehicle Class and Driver Behavior." American Economic Journal: Applied Economics 5 (3).
- Jacobsen, Mark R., Christopher R. Knittel, James M. Sallee, and Arthur A. van Benthem. 2016. "Sufficient Statistics for Imperfect Externality-Correcting Policies." Manuscript: University of California at Berkeley.
- Jacobsen, Mark R., James M. Sallee, Joseph S. Shapiro, and Arthur A. van Benthem. 2019. "Are New Vehicle Emissions Standards Effective and Efficient?" Manuscript: University of California, Berkeley.
- Jacobsen, Mark R. and Arthur A. van Benthem. 2015. "Vehicle Scrappage and Gasoline Policy." American Economic Review 105 (3):1312–1338.
- Jaffe, Adam B. and Robert N. Stavins. 1994. "The Energy-Efficiency Gap." Energy Policy 22 (10):804–810.

- Jiang, Wei, Ashlyn Aiko Nelso, and Edward Vytlačil. 2014. "Liar's Loan? Effects of Origination Channel and Information Falsification on Mortgage Delinquency." The Review of Economics and Statistics 96 (1):1–18.
- Kane, Thomas J. and Douglas O. Staiger. 2008. "Estimating Teacher Impacts on Student Achievement: An Experimental Evaluation." NBER Working Paper 14607.
- Klier, Thomas and Joshua Linn. 2012. "New-vehicle Characteristics and the Cost of the Corporate Average Fuel Economy Standard." RAND Journal of Economics 43 (1):186–213.
- Kopczuk, Wojciech, Justin Marion, Erich Muehlegger, and Joel Slemrod. 2016. "Does Tax-Collection Invariance Hold? Evasion and the Pass-through of State Diesel Taxes." American Economic Journal: Economic Policy 8 (2):1–36.
- Langer, Ashley and Shaun McRae. 2014. "Step on It: Evidence on the Variation in On-Road Fuel Economy." Manuscript: University of Arizona.
- Leard, Benjamin, Joshua Linn, and Katalin Springel. 2019. "Pass-Through and Welfare Effects of Regulations that Affect Product Attributes." Resources for the Future Discussion Paper 19-07.
- Leggett, Christopher G. 2002. "Environmental Valuation with Imperfect Information." Environmental and Resource Economics 23:343–355.
- Levinson, Arik. 2016. "How Much Energy Do Building Energy Codes Save? Evidence from California Houses." American Economic Review 106 (10):2867–2894.
- Parry, Ian W.H., David Evans, and Wallace E. Oates. 2014. "Are Energy Efficiency Standards Justified?" Journal of Environmental Economics and Management 67:104–125.
- Reynaert, Mathias. 2017. "Abatement Strategies and the Cost of Environmental Regulation: Emission Standards on the European Car Market." Manuscript: Toulouse School of Economics.
- Reynaert, Mathias and Frank Verboven. 2014. "Improving the performance of random coefficients demand models: The role of optimal instruments." Journal of Econometrics 179 (1):83–98. URL <https://ideas.repec.org/a/eee/econom/v179y2014i1p83-98.html>.
- Rhodes, Andrew and Chris M. Wilson. Forthcoming. "False Advertising." RAND Journal of Economics .
- Roychowdhury, Sugata. 2006. "Earnings Management through Real Activities Manipulation." Journal of Accounting and Economics 42:335–370.
- Sallee, James M. 2014. "Rational Inattention and Energy Efficiency." Journal of Law and Economics 57 (3):781–820.
- Sallee, James M. and Joel Slemrod. 2012. "Car Notches: Strategic Automaker Responses to Fuel Economy Policy." Journal of Public Economics 96 (11-12):981–999.
- Sallee, James M., Sarah E. West, and Wei Fan. 2016. "Do Consumers Recognize the Value of Fuel Economy? Evidence from Used Car Prices and Gasoline Price Fluctuations." Journal of Public Economics 135:61–73.
- Slemrod, Joel. 2008. "Does It Matter Who Writes the Check to the Government? The Economics of Tax Remittance." National Tax Journal 61 (2):251–275.

- Small, Kenneth A. and Harvey S. Rosen. 1981. “Applied welfare economics of discrete choice models.” Econometrica .
- Tanaka, Shinsuke. 2017. “Mind the Gap! Tax Incentives and Incentives for Manipulating Fuel Efficiency in the Automobile Industry.” Working Paper: Tufts University.
- The International Council on Clean Transportation. 2014. “From Laboratory to Road: A 2014 Update of Official and “Real-World” Fuel Consumption and CO<sub>2</sub> Values for Passenger Vehicles in Europe.” ICCT White Paper.
- . 2015. “From Laboratory to Road: A 2015 Update of Official and “Real-World” Fuel Consumption and CO<sub>2</sub> Values for Passenger Cars in Europe.” ICCT White Paper.
- TNO. 2011. “Support for the revision of Regulation (EC) No 443/2009 on CO<sub>2</sub> emissions from cars.” URL [https://ec.europa.eu/clima/sites/clima/files/transport/vehicles/cars/docs/study\\_car\\_2011\\_en.pdf](https://ec.europa.eu/clima/sites/clima/files/transport/vehicles/cars/docs/study_car_2011_en.pdf). Framework Contract No ENV.C.3./FRA/2009/0043.
- Walker, Reed W. 2013. “The Transitional Costs of Sectoral Reallocation: Evidence from the Clean Air Act and the Workforce.” Quarterly Journal of Economics 128 (4):1787–1835.
- Weyl, E. Glen and Michael Fabinger. 2013. “Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition.” Journal of Political Economy 121 (3):528–583.
- Whitefoot, Kate S., Meredith L. Fowlie, and Steven J. Skerlos. 2017. “Compliance by Design: Influence of Acceleration Trade-offs on CO<sub>2</sub> Emissions and Costs of Fuel Economy and Greenhouse Gas Regulations.” Environmental Science & Technology 51 (18):10307–10315.
- Whitefoot, Katie and Steven Skerlos. 2012. “Design Incentives to Increase Vehicle Size Created from the U.S. Footprint-based Fuel Economy Standards.” Energy Policy 41:402–411.



## A APPENDIX MATERIAL

### A.1 Sample selection process:

Self-reported odometer readings are sometimes missing or clearly incorrect. We eliminate unreliable data through a data selection process detailed here.

Our data selection process first eliminates a number of drivers (accounts) that have unreliable information or too few observations. With the accounts that remain, we then consider several ways to account for mismeasurement in odometer readings for individual transactions. Specifically, we first limit the sample to gasoline- or diesel-fueled vehicles, which eliminates 6.7 million transactions. Second, we drop vehicles that use the wrong type of fuel for their engine in more than 1% of the visits, e.g., putting diesel fuel in a vehicle that is labeled as gasoline in our data. Inconsistencies might be in the data because drivers use their card for a different vehicle, or these observations might be mistakes in the assignment of vehicle type. This drops 7.5 million transactions. Third, we pose some minimum requirements on the driving patterns of the drivers that produce the transactions. We drop drivers that never report an increase of more than 150km in their odometer reading (2.5 million transactions).<sup>40</sup> We drop car models with fewer than 10 drivers, and drivers with fewer than 10 fuel station visits (1.3 million transactions). We drop drivers that did not report driving more than 5,000km in total or reported driving more than 500,000km in total (11.3 million). Having isolated a set of drivers (accounts) with ample data, within those accounts we drop individual transactions in two steps. First, we drop transactions where the odometer difference is lower than 100km or higher than 3000km (7.6 million). Second, we drop transactions that result in a fuel consumption that is outside 1.25 times the interquartile range of estimated fuel consumption for each car model in the data (5.1 million). This results in the final dataset of 24 million observations.

### A.2 The EU emission standard and member state taxes

The discussion in this section is based on [Reynaert \(2017\)](#). The European regulation on emission standards for new passenger cars, Regulation (EC) No. 443/2009, sets a mandatory fleet average of  $\kappa = 130$  grams CO<sub>2</sub> per kilometer. Denoting the sales of each product  $j$  by  $q_j$  and the emissions of each product by  $e_j$ , the target for a firm is as follows:

$$\frac{\sum_{j \in \text{fleet}} q_j (e_j - f(w_j))}{\sum_{j \in \text{fleet}} q_j} \leq 130.$$

The attribute basing  $f(w_{jm}) = a(w_j - w_0)$  adjusts the emissions of each vehicle by the distance in the vehicle weight  $w_j$  from a shifting point  $w_0$  (the pivotal weight point). The shifting point  $w_0$  is a mass of 1370 kg and the difference in weight from that point is multiplied by  $a = 0.046$ . The target is set for each producer's fleet of new vehicles sold in a calendar year and the trading of excess emissions between producers is not allowed. Manufacturers can obtain lower average emissions by

---

<sup>40</sup>Note that the range of a combustion engine is easily more than 800km.

gathering super credits. These credits are given for vehicles that emit less than 50 g/km. There are also separate standards for small manufacturers making less than 30,000 vehicles per year. Both of these exceptions count for a very small share of the total market. There is no banking system for excess emissions over time.

When producers exceed the standard they have to pay premiums for excess emissions. The premium is €5 per unit sold for the first excess g/km and increases to €95 per unit above 134 g/km. A manufacturer obtaining a sales weighted emission of 146 g/km, the average in 2007, would face a significant penalty of €1,280 per vehicle (against an average sales price of €22,250). The regulation was proposed by the European Commission in 2007 and became a European law in 2009. In 2012, 65% of manufacturer's sales had to comply with the emission standard. This rose to 75% in 2013 and 80% in 2014, and the standard was fully binding from 2015 onward. Every firm succeeded in reaching the full target by 2014.

On top of the EU wide emission standard, all European member states have fiscal policies in the vehicle market. The national policies are a combination of value added tax, registration taxes, annual taxes, fuel taxes, road taxes and emission zones. These taxes differ substantially between member states and there are also changes over time. A full discription of these taxes is out of scope, we recommend ACEA tax reports as a source for more information on these tax systems (the mere description of these taxes takes several pages per country). Using the ACEA data source, [Gerlagh et al. \(Forthcoming\)](#) show that several countries have increased the extent to which these taxes are base on CO<sub>2</sub> emissions. Several examples are a bonus malus system in France announced in 2008, Spain included CO<sub>2</sub> emissions in a complicated formula that determines registration taxes in 2008, Germany changed the formula of annual taxes in 2009, and the Netherlands changed tax formulas in 2010.

### A.3 Empirical Bayes Correction of on-road ratings

We start by decomposing the total variance in the sample  $Var(r_{nij}) = \sigma_r^2$  into three components: variance in performance of vehicles  $\sigma_j^2$ , drivers  $\sigma_i^2$ , and pump visits  $\sigma_n^2$ . We estimate the variance between pump visits of the same driver as:

$$\sigma_n^2 = \frac{1}{N - I} \sum_n (r_{nij} - \bar{r}_{ij})^2,$$

in which  $\bar{r}_{ij}$  is the mean fuel consumption of driver  $i$ ,  $N$  is the total number of observations and  $I$  is the total number of drivers. Next, we estimate the covariance between drivers of the same vehicles as:

$$\sigma_j^2 = cov(\bar{r}_{ij}, \bar{r}_{kj}).$$

The estimated covariance is obtained as a weighted average of covariances between randomly sorted pairs  $(i, k)$  of drivers of the same car. We weigh each pair of drivers  $(i, k)$  by the sum of their visits.

Finally, we obtain  $\sigma_i^2$  as the remaining variance:  $\sigma_i^2 = \sigma_r^2 - \sigma_n^2 - \sigma_j^2$ . The precision of the estimated gap for each driver is then defined as:

$$h_i = 1/(\sigma_i^2 + \sigma_n^2/n_i),$$

so that drivers with a high number of visits have a higher precision. We obtain precision weighted means per car as the weighted average of  $\bar{r}_{ij}$  with  $h_i$  as weights. Second we shrink these precision weighted means with an estimate of their reliability:

$$\psi_j = \sigma_j^2/(\sigma_j^2 + 1/\sum_i h_{ji}),$$

where the reliability is defined as the signal  $\sigma_j$  over the total variance. We use the per vehicle shrunk on-road estimates  $\hat{r}_j$  to construct an alternative estimate of the gap defined in (2) and to inform us about the distribution of the gap between vehicles.

Table A.3 describes the variation in the on-road fuel consumption ( $\sigma_d^2$ ) across release years, as well as its decomposition across three components: variation across refueling transactions for the same driver ( $\sigma_n^2$ ), variation across drivers of the same vehicle ( $\sigma_i^2$ ) and variation across vehicles ( $\sigma_j^2$ ). We decompose the variation separately for each release year and describe the mean and standard deviation across release years in the table.

More than 25% of the variance is attributable to within driver variance. This variance is due to driving conditions, stockpiling effects and errors in odometer reporting. We find that the variance across drivers of the same car  $\sigma_i^2$  is 0.21. This is an economically large number; it means that the on-road fuel consumption is estimated to be 0.28 liter/100km higher at the third quartile than at the first quartile of drivers in the same car.<sup>41</sup> A policy that would shift a driver from the third quartile of the fuel consumption gap to the first quartile would decrease fuel consumption by 3%. These numbers are interesting from a policy perspective as they give an indication of the extent to which fuel consumption and emissions can be reduced by teaching and incentivizing drivers to drive a vehicle more efficiently.<sup>42</sup> The remaining part of the variance  $\sigma_j^2$  is the co-variance between drivers of the same car and can be seen as the information available to estimate the car specific component of on-road fuel consumption. We estimate this to be 1.35, which is more than 60% of the total variance. Table A.3 also shows that the variance components are relatively stable over time; each component has a low standard deviation across release years. There is variation in the size of the fuel consumption gap between cars and between drivers, but this variation is stable over time. Given this variance decomposition we turn next to the estimates of the distribution of  $r_j$  and  $d_j$  for each release year. Table A.4 reports the unweighted mean estimate of  $r_j$  and  $\hat{r}_j$ , obtained with the empirical Bayes correction. The mean value of both  $r_j$  and  $\hat{r}_j$  are decreasing

<sup>41</sup>If we assume that conditional on car  $j$ ,  $r$  has a normal distribution, the interquartile distance is  $1.349 \cdot \sigma_i^2$ .

<sup>42</sup>Significant variation across drivers of identical cars is consistent with results reported in [Langer and McRae \(2014\)](#), who analyze extremely detailed driving data from a few dozen drivers of an identical car, the Honda Accord. In contrast, our data come from a large sample and cover many models.

over the release years. In all years the corrected means are lower than the raw means, because on average vehicles with high  $r_j$  have less precise underlying data, but overall shrinkage and precision weighting has small effects. The resulting gap  $\hat{d}_j$  is estimated to be an imprecise 10% up until 2006. From 2007 onwards we see a significant increase in the performance gap, consistent with the previous estimates.

## A.4 Proofs

**Proposition 1.** *In the absence of policy ( $\lambda = 0$ ), buyer surplus falls with the level of gaming. Specifically:*

$$\frac{dBS}{dg} \approx \underbrace{-\rho(1-\alpha)\beta D}_{\text{price effect}} + \underbrace{(1-\alpha)^2\beta^2 D'g}_{\text{choice distortion}} \leq 0.$$

We derive this by starting with the definition of buyer surplus as the integral under the inverse demand curve. Because the true attribute of the good  $x$  is unaffected by gaming, we can analyze buyer surplus using the true demand curve. The standard portion of buyer surplus is the integral from the final price, denoted  $p^*$  up to infinity. Denote by  $\tilde{p}$  the upfront purchase price that would induce a sophisticated consumer to purchase the amount of the good that is in fact purchased at price  $p^*$  by the consumer with perception  $(1-\alpha)$ . The choice distortion can be written as the difference between the revenue generated between  $p^*$  and  $\tilde{p}$  and the consumer value generated between those points.

$$BS = \underbrace{\int_{p^*}^{\infty} D(z + \beta x) dz}_{\text{BS of correct quantity}} + \underbrace{\int_{\tilde{p}}^{p^*} D(z + \beta x) dz}_{\text{Value of excess quantity}} - \underbrace{\int_{\tilde{p}}^{p^*} D(\tilde{p} + \beta x) dz}_{\text{Cost of excess quantity}}. \quad (\text{A.1})$$

Differentiation of equation A.5, in which  $p^*$  and  $\tilde{p}$  are endogenously determined by  $g$ , yields the result. Note that the inside of the third integral is a constant with respect to the variable of integration, so it can be pulled out of the integral, leaving only the constant 1 inside. Specifically:

$$\begin{aligned} \frac{dBS}{dg} = & -D(p^* + \beta x) \frac{dp^*}{dg} + \left\{ D(p^* + \beta x) \frac{dp^*}{dg} - D(\tilde{p} + \beta x) \frac{d\tilde{p}}{dg} \right\} \\ & - \left\{ (p^* - \tilde{p}) D'(\tilde{p} + \beta x) \frac{d\tilde{p}}{dg} + D(\tilde{p} + \beta x) \left( \frac{dp^*}{dg} - \frac{d\tilde{p}}{dg} \right) \right\}. \end{aligned} \quad (\text{A.2})$$

Using the pass through coefficient  $\rho$ , a change in  $g$  scales to a change in tax by  $(1-\alpha)\beta$ , so  $dp^*/dg = \rho(1-\alpha)\beta$ . This simplifies the first term to yield the result.

For the second term in A.2, note that  $\tilde{p} = p^* - (1-\alpha)\beta g$  by definition. Then, substitute a first-

order Taylor approximation to write demand at  $\tilde{p}$  as a function of demand at  $p^*$  and  $D' \times (\tilde{p} - p^*)$ :

$$\begin{aligned}
& D(p^* + \beta x) \frac{dp^*}{dg} - D(\tilde{p} + \beta x) \frac{d\tilde{p}}{dg} \\
&= D(p^* + \beta x) \frac{dp^*}{dg} - D(p^* + \beta x - (1 - \alpha)\beta g) \frac{d\tilde{p}}{dg} \\
&\approx \{D(p^* + \beta x - (1 - \alpha)\beta g) + D'(p^* + \beta x - (1 - \alpha)\beta g)(1 - \alpha)\beta\} \frac{dp^*}{dg} - D(p^* + \beta x - (1 - \alpha)\beta g) \frac{d\tilde{p}}{dg} \\
&= D(p^* + \beta x - (1 - \alpha)\beta g) \left( \frac{dp^*}{dg} - \frac{d\tilde{p}}{dg} \right) + D'(p^* + \beta x - (1 - \alpha)\beta g)(1 - \alpha)\beta \frac{dp^*}{dg} \tag{A.3}
\end{aligned}$$

Now consider the third term in A.2. After substituting  $\tilde{p} = p^* - (1 - \alpha)\beta g$ , we see that the term that multiplies the difference in derivatives will cancel in the third term of A.2 and the second-term, defined using A.3. This means that A.2 can be written:

$$\begin{aligned}
\frac{dBS}{dg} &= -D(p^* + \beta x)\rho(1 - \alpha)\beta + D'(p^* + \beta x - (1 - \alpha)\beta g)(1 - \alpha)\beta \frac{dp^*}{dg} \\
&\quad - \left\{ (p^* - \tilde{p})D'(\tilde{p} + \beta x) \frac{d\tilde{p}}{dg} \right\}. \tag{A.4}
\end{aligned}$$

Substitute  $p^* - \tilde{p} = (1 - \alpha)\beta g$  and  $d\tilde{p}/dg = (\rho - 1)(1 - \alpha)\beta$ , which follows from differentiating the definition of  $\tilde{p}$  and using the pass through result for  $p^*$ . Then, simplification yields the final result:

$$\begin{aligned}
\frac{dBS}{dg} &= -D(p^* + \beta x)\rho(1 - \alpha)\beta + D'(p^* + \beta x - (1 - \alpha)\beta g)(1 - \alpha)\beta \rho(1 - \alpha)\beta \\
&\quad - \{ (1 - \alpha)\beta g D'(\tilde{p} + \beta x)(\rho - 1)(1 - \alpha)\beta \} \\
&= -D(p^* + \beta x)\rho(1 - \alpha)\beta + D'(p^* + \beta x - (1 - \alpha)\beta g)(1 - \alpha)^2 \beta^2 \rho \\
&\quad - (1 - \alpha)^2 \beta^2 g D'(\tilde{p} + \beta x)(\rho - 1) \\
&= -D(p^* + \beta x)\rho(1 - \alpha)\beta + (1 - \alpha)^2 \beta^2 g D'(\tilde{p} + \beta x).
\end{aligned}$$

Note that, with the local linear demand assumption, the derivative of  $D$  evaluated at either  $\tilde{p}$  or  $p^*$  is the same. ■

**Proposition 2.** *In the presence of a binding standard ( $\lambda > 0$ ), a change in gaming affects buyer surplus as follows:*

$$\frac{dBS}{dg} \approx \underbrace{(-\rho(c' - \alpha\beta) - \beta)D}_{\text{price effect}} + \underbrace{(1 - \alpha)^2 \beta^2 D'g}_{\text{choice distortion}}.$$

Proposition 2 is derived in a similar way to Proposition 1, but we define surplus using integrals over the demand function starting with full prices  $f$ . (This same could have been done in the prior proof, yielding the same result.) Recall that  $f = p + \beta x$  and  $\tilde{f} = p + \beta x - (1 - \alpha)\beta g$ .

$$BS = \underbrace{\int_{f^*}^{\infty} D(z)dz}_{\text{BS of correct quantity}} + \underbrace{\int_{\tilde{f}}^{f^*} D(z)dz}_{\text{Value of excess quantity}} - \underbrace{\int_{\tilde{f}}^{f^*} D(\tilde{f})dz}_{\text{Cost of excess quantity}}. \quad (\text{A.5})$$

Differentiating, using the same Taylor approximation as above to simplify, yields:

$$\frac{dBS}{dg} \approx -D(f^*) \frac{df^*}{dg} + D'(f^*)(f^* - \tilde{f}) \left( \frac{df^*}{dg} - \frac{d\tilde{f}}{dg} \right).$$

Substitute the pass through result described in the text:  $df^*/dg = \rho(c'(x) + \alpha\beta) + \beta$ . Substitute the definition of  $\tilde{f}$ , which shows that  $f^* - \tilde{f} = (1 - \alpha)\beta g$ . And substitute  $\frac{df^*}{dg} - \frac{d\tilde{f}}{dg} = (1 - \alpha)\beta$ :

$$\begin{aligned} \frac{dBS}{dg} &\approx -D(f^*)(\rho(c'(x) + \alpha\beta) + \beta) + D'(f^*) \{(1 - \alpha)\beta g\} \{(1 - \alpha)\beta\} \\ &= -D(f^*)(\rho(c'(x) + \alpha\beta) + \beta) + D'(f^*)(1 - \alpha)^2 \beta^2 g. \end{aligned} \quad (\text{A.6})$$

This yields the result. ■

## A.5 Estimation Details

We estimate the parameters  $\theta$  using the following GMM objective:

$$\min_{\theta} \xi_j' Z' \omega Z \xi_j$$

in which  $Z$  is a matrix of instruments and  $\omega$  is a weighting matrix.

We use a panel containing sales, prices and characteristics for all new vehicle sales in seven European countries between 1998 and 2007.<sup>43</sup> We only use data from before 2008 to estimate demand, so that our estimates come from a period in which the performance gap was stable. We thus assume variation in  $x_j$  is informative about actual fuel cost differences. The vector  $\Delta_j$  contains information on horsepower, weight, footprint (a measure of vehicle size), height and a dummy specifying if the car is of a foreign brand (e.g., Fiat in France). Additionally, we include fuel type by market dummies, dummies for the number of months a vehicle was on sale in a country-year, country fixed effects, a linear time trend, body type fixed effects, vehicle class fixed effects and brand fixed effects. We divide prices by income per capita in each country-year, so that price sensitivity varies with income in the market. We need instruments for price and for the standard deviations of the random coefficients. We instrument for prices using both cost shifters and sums of characteristics instruments, which follows [Berry, Levinsohn, and Pakes \(1995\)](#). The sums of characteristics instruments are the sum of fuel costs, horsepower, weight, footprint and

<sup>43</sup>The seven countries are: Belgium, France, Germany, Great Britain, Italy, Netherlands and Spain.

height across all other products in the market and across all other products within the same firm in the market. We also include the number of competing products and the number of products in the same firm. The cost shifters are the log of labor costs in the country of production and a dummy specifying if the vehicle is sold in the country of production. For the standard deviations, we use approximately optimal instruments that are constructed using a two-step procedure as described in [Reynaert and Verboven \(2014\)](#). We estimate considerable heterogeneity in the taste for horsepower, weight and footprint. Vehicles perceived as foreign are less attractive for consumers. Cars with four doors are preferred over cars with two doors.

## A.6 Additional Tables

**Table A.1:** Summary statistics

	Mean	St. Dev.
<b>Car Characteristics</b>		
Fuel Consumption (L/100km)	6.65	1.73
Vehicle Weight in kg	1,354	245
Diesel Engines	0.46	0.50
Drivers per car	107	219
<b>Driver Characteristics</b>		
Pump visits	134	80
Total liters purchased	6,015	3,666
Total distance (km)	111,726	53,942
<b>Pump Visit Characteristics</b>		
Liters per visit	45.3	10.8
Odometer increase per visit	671	192

The table gives summary statistics for the 2,696 vehicles, 266,616 drivers and 23,989,576 pump visits in the sample.

**Table A.2:** Summary statistics: Netherlands and Travelcard

	Mean	St. Dev.	Mean	St. Dev.
	<b>TravelCard</b>		<b>Netherlands</b>	
Price (euro)	31,672	13,367	40,767	29,676
Fuel Consumption (L/100km)	6.74	1.60	7.89	2.46
Vehicle Weight in kg	1,344	230	1,409	308
Diesel Engines	0.45	0.50	0.36	0.48

Summary statistics for the TravelCard sample and the full dutch market between 1998 and 2011.



**Table A.3:** Variance decomposition

	$\sigma_d^2$	$\sigma_i^2$	$\sigma_j^2$	$\sigma_n^2$
Mean	2.11	0.21	1.35	0.56
Standard deviation	0.57	0.03	0.45	0.11
Variance decomposition (%)	100	10.36	62.34	27.30
Standard deviation		2.54	6.76	4.40

$\sigma_d^2$  is the total variance in  $r_{nij}$ ,  $\sigma_i^2$  is the variation attributable to differences across individuals driving the same vehicle,  $\sigma_j^2$  is the covariance between drivers in the same vehicle,  $\sigma_n^2$  is the variation across refueling visits of the same driver in the same vehicle. The variance decomposition is performed separately for each release year, and the mean and standard deviation across years are reported in table.

**Table A.4:** Estimated vehicle gaps by release year, with and without empirical Bayes correction

Release	Unweighted $r_j$	$\hat{r}_j$	Shrinkage	$\hat{d}_j$	TS $\hat{d}_j$
1998	7.91 (1.41)	7.87 (1.39)	0.99 (0.00)	7.85 (7.82)	
1999	7.90 (1.49)	7.88 (1.46)	0.99 (0.01)	8.70 (6.90)	1.16 (1.29)
2000	7.96 (1.44)	7.94 (1.43)	0.99 (0.01)	8.60 (7.36)	-0.17 (0.88)
2001	7.73 (1.44)	7.72 (1.41)	0.99 (0.00)	11.15 (7.12)	2.10 (0.98)
2002	7.63 (1.58)	7.60 (1.57)	0.99 (0.00)	11.02 (9.17)	1.86 (1.58)
2003	7.83 (1.61)	7.80 (1.56)	0.99 (0.00)	11.01 (9.18)	3.21 (1.47)
2004	8.31 (1.68)	8.29 (1.64)	0.99 (0.00)	9.52 (9.58)	2.90 (1.47)
2005	7.87 (1.71)	7.82 (1.59)	0.99 (0.01)	12.56 (9.14)	6.37 (1.57)
2006	8.10 (1.57)	8.06 (1.53)	0.99 (0.00)	12.55 (8.17)	11.57 (4.99)
2007	7.71 (1.37)	7.68 (1.34)	0.99 (0.00)	17.66 (9.67)	10.04 (1.48)
2008	7.46 (1.36)	7.45 (1.34)	0.99 (0.00)	17.57 (9.00)	12.14 (1.28)
2009	7.24 (1.30)	7.21 (1.28)	0.99 (0.00)	22.34 (10.07)	21.73 (2.01)
2010	7.15 (1.33)	7.13 (1.32)	0.99 (0.00)	26.38 (10.68)	27.24 (2.23)
2011	6.92 (1.24)	6.90 (1.21)	0.99 (0.01)	29.74 (10.29)	27.16 (1.35)
2012	6.64 (1.08)	6.62 (1.06)	0.99 (0.01)	37.46 (11.05)	38.92 (2.84)
2013	6.23 (0.92)	6.23 (0.89)	0.99 (0.01)	44.23 (11.34)	45.72 (1.79)
2014	6.18 (1.20)	6.15 (1.13)	0.98 (0.02)	52.90 (12.04)	48.68 (2.46)

Table reports mean and standard deviations for the distribution of estimated on-road consumption  $r_j$ , Bayes corrected on-road consumption  $\hat{r}_j$ , shrinkage factor and Bayes corrected efficiency gap  $\hat{d}_j$  by release year. The final column gives the efficiency gap estimated from a two stage procedure. The first step regresses the gap on driver by vehicle fixed effects with time controls, the second step regresses the driver by vehicle effects on release year.

**Table A.5:** Regression of performance gap on controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Base	Control	Gasoline	Diesel	MonthYear	Age1	Age2	Age3	Location	First Year
Release:										
1999	1.996 (1.454)	0.526 (0.881)	-0.103 (1.042)	-0.749 (1.146)	0.524 (0.880)		0.798 (0.886)	0.691 (0.878)	0.523 (0.871)	
2000	0.319 (0.884)	0.745 (0.724)	1.216 (1.001)	0.582 (1.137)	0.745 (0.724)		0.875 (0.700)	0.898 (0.705)	0.727 (0.716)	
2001	2.329 (0.976)	-0.179 (0.949)	1.367 (0.869)	-2.423 (1.183)	-0.178 (0.949)		-0.0845 (0.886)	-0.0199 (0.904)	-0.163 (0.941)	
2002	1.766 (1.561)	1.271 (1.067)	2.150 (1.060)	-0.152 (1.552)	1.278 (1.066)		1.161 (1.039)	1.532 (1.044)	1.279 (1.061)	
2003	2.463 (1.319)	1.307 (0.801)	1.661 (1.048)	-0.498 (1.144)	1.313 (0.801)		1.032 (0.787)	1.410 (0.782)	1.304 (0.793)	
2004	3.300 (1.408)	2.566 (0.783)	2.256 (0.940)	1.221 (1.233)	2.570 (0.782)		2.279 (0.763)	2.570 (0.763)	2.506 (0.778)	
2005	4.709 (1.432)	4.200 (0.966)	6.193 (1.531)	1.942 (1.254)	4.201 (0.966)		3.565 (0.923)	3.863 (0.928)	4.189 (0.956)	2.497 (2.391)
2006	10.59 (5.085)	3.842 (1.085)	5.559 (1.236)	2.758 (1.482)	3.834 (1.085)	0.284 (1.029)	3.308 (1.151)	3.395 (1.077)	3.808 (1.075)	5.803 (1.733)
2007	7.381 (1.385)	8.460 (1.036)	8.619 (1.344)	8.057 (1.347)	8.444 (1.035)	4.404 (1.071)	7.643 (0.999)	7.985 (0.998)	8.392 (1.030)	2.790 (2.087)
2008	9.533 (1.370)	8.857 (1.003)	9.230 (1.060)	8.292 (1.476)	8.833 (1.002)	4.271 (1.080)	7.536 (0.988)	7.983 (0.984)	8.796 (0.992)	12.60 (2.474)
2009	18.50 (2.238)	13.74 (1.411)	10.37 (1.098)	16.22 (2.444)	13.72 (1.409)	8.715 (1.717)	11.97 (1.337)	12.38 (1.316)	13.66 (1.397)	15.66 (2.207)

**Table A.5:** Regression of performance gap on controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Base	Control	Gasoline	Diesel	MonthYear	Age1	Age2	Age3	Location	First Year
2010	22.57 (2.100)	17.93 (1.442)	15.54 (1.312)	19.60 (1.815)	17.91 (1.442)	12.47 (1.544)	15.52 (1.383)	15.99 (1.384)	17.82 (1.427)	19.85 (2.078)
2011	23.43 (1.570)	18.34 (1.128)	15.82 (1.425)	20.41 (1.428)	18.33 (1.129)	12.34 (1.386)	15.47 (1.111)	16.10 (1.130)	18.27 (1.116)	23.28 (2.095)
2012	35.16 (2.948)	28.62 (2.086)	19.72 (1.601)	35.46 (2.095)	28.65 (2.085)	22.09 (2.311)	24.80 (2.022)	25.61 (2.068)	28.60 (2.057)	36.96 (4.378)
2013	41.24 (1.887)	35.86 (1.655)	31.80 (2.233)	39.04 (1.987)	35.93 (1.651)	28.87 (1.955)	30.82 (1.624)	31.69 (1.683)	35.84 (1.632)	43.16 (2.482)
2014	42.68 (2.466)	40.72 (3.501)	32.26 (3.982)	48.83 (2.309)	40.71 (3.495)	33.68 (3.660)	35.96 (3.155)	35.48 (3.450)	40.54 (3.470)	41.50 (2.521)
Year:										
2005		0.121 (0.133)	0.537 (0.135)	-0.165 (0.180)						
2006		0.560 (0.222)	1.608 (0.259)	-0.107 (0.293)						
2007		0.710 (0.285)	1.681 (0.348)	0.0657 (0.378)						
2008		1.068 (0.330)	1.984 (0.398)	0.390 (0.448)						
2009		1.378 (0.377)	2.000 (0.467)	0.661 (0.520)						
2010		1.966 (0.422)	2.823 (0.526)	0.831 (0.588)						

**Table A.5:** Regression of performance gap on controls

	(1) Base	(2) Control	(3) Gasoline	(4) Diesel	(5) MonthYear	(6) Age1	(7) Age2	(8) Age3	(9) Location	(10) First Year
2011		2.363 (0.489)	2.899 (0.609)	1.209 (0.693)						
2012		3.330 (0.601)	3.251 (0.674)	2.436 (0.870)						
2013		3.919 (0.689)	3.423 (0.729)	3.115 (1.017)						
2014		3.139 (0.762)	2.240 (0.800)	2.464 (1.125)						
2015		4.743 (0.817)	3.555 (0.870)	3.842 (1.202)						
Month:										
2		-0.914 (0.0205)	-0.812 (0.0308)	-1.001 (0.0267)						
3		-2.049 (0.0616)	-2.155 (0.123)	-1.988 (0.0554)						
4		-3.605 (0.106)	-4.170 (0.210)	-3.219 (0.0723)						
5		-4.197 (0.138)	-5.303 (0.253)	-3.429 (0.0984)						
6		-4.479 (0.157)	-5.567 (0.299)	-3.738 (0.107)						
7		-3.841 (0.164)	-5.135 (0.302)	-2.949 (0.112)						

**Table A.5:** Regression of performance gap on controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Base	Control	Gasoline	Diesel	MonthYear	Age1	Age2	Age3	Location	First Year
8		-4.292 (0.161)	-5.650 (0.293)	-3.348 (0.104)						
9		-4.923 (0.155)	-5.834 (0.297)	-4.318 (0.121)						
10		-4.129 (0.115)	-4.566 (0.224)	-3.889 (0.105)						
11		-2.400 (0.0751)	-2.338 (0.146)	-2.537 (0.0795)						
12		0.120 (0.0425)	0.360 (0.0489)	-0.191 (0.0532)						
Diesel		2.083 (0.336)			2.089 (0.336)	2.085 (0.341)	3.273 (0.329)	2.449 (0.327)	2.182 (0.334)	
DistPerYear		-1.48e-05 (4.40e-06)	-3.98e-05 (5.52e-06)	-5.30e-06 (5.67e-06)	-1.47e-05 (4.39e-06)	-1.54e-05 (4.41e-06)	-6.47e-06 (4.53e-06)	-4.87e-05 (3.23e-06)	-1.21e-05 (4.33e-06)	
DistPerYear <sup>2</sup>		2.09e-10 (0)	4.05e-10 (0)	1.29e-10 (0)	2.06e-10 (0)	2.06e-10 (0)	-0 (0)	2.78e-10 (0)	1.91e-10 (0)	
DistPerYear <sup>3</sup>		-0 (0)	-0 (0)	-0 (0)	-0 (0)	-0 (0)	0 (0)	-0 (0)	-0 (0)	
Age:										
1						0.599 (0.380)		-1.871 (0.114)		
2						-0.210 (0.458)		-3.366 (0.204)		

**Table A.5:** Regression of performance gap on controls

	(1) Base	(2) Control	(3) Gasoline	(4) Diesel	(5) MonthYear	(6) Age1	(7) Age2	(8) Age3	(9) Location	(10) First Year
3						-0.885 (0.545)		-4.627 (0.308)		
4						-1.532 (0.640)		-5.876 (0.422)		
5						-1.932 (0.737)		-7.277 (0.523)		
6						-2.272 (0.788)		-8.915 (0.638)		
7						-2.599 (0.860)		-10.60 (0.774)		
8						-3.127 (0.944)		-12.55 (0.930)		
9						-3.734 (1.034)		-12.92 (1.213)		
10						-4.628 (1.147)		-14.77 (1.447)		
11						-5.327 (1.242)		-18.40 (2.009)		
12						-5.852 (1.366)				
13						-5.803 (1.542)				
14						-5.579				



**Table A.5:** Regression of performance gap on controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Base	Control	Gasoline	Diesel	MonthYear	Age1	Age2	Age3	Location	First Year
15						(1.718)				
						-4.740				
16						(1.893)				
						-3.471				
17						(2.068)				
						-3.851				
						(2.272)				
Odometer							-4.24e-05			
							(2.11e-06)			
Constant	11.69 (0.660)	13.52 (0.654)	16.15 (0.658)	15.02 (0.887)						10.28 (1.385)
Model F.E.	-	569	569	569	569	569	569	596	596	-
MonthYear F.E.	-	-	-	-	135	135	135	135	135	-
Location F.E.	-	-	-	-	-	-	-	-	19954	-
Observations	23,989,576	23,989,576	10,173,983	13,815,593	23,989,576	23,989,576	23,989,576	23,989,576	23,751,895	779,761
R-squared	0.281	0.400	0.361	0.459	0.401	0.401	0.409	0.405	0.413	0.469

Table reports coefficients of regressions of performance gap on controls. (1) includes only release year fixed effects (2) controls for year, month, fuel type, distance traveled per year and model fixed effects (absorbed), (3) equals (2) but limits the sample to gasoline engines, (4) equals (2) but limits the sample to diesel engines, (5) includes month by year fixed effects and model fixed effects (both absorbed), (6) equals (5) but includes age dummies, (7) equals (5) but includes odometer reading, (8) equals (5) but includes alternative age dummies, (9) equals (5) but adds fuel station fixed effects and (10) limits the sample to first year of driving.

**Table A.6:** Regression of performance gap in levels on controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Base	Control	Gasoline	Diesel	MonthYear	Age1	Age2	Age3	Location	First Year
Release:										
1999	0.0813 (0.0724)	0.0627 (0.0468)	-0.00764 (0.0608)	-0.0143 (0.0509)		0.0626 (0.0468)	0.0776 (0.0472)	0.0709 (0.0469)	0.0626 (0.0463)	
2000	0.0305 (0.0495)	0.0649 (0.0391)	0.0821 (0.0591)	0.0486 (0.0573)	0.0649 (0.0391)		0.0720 (0.0378)	0.0725 (0.0382)	0.0637 (0.0386)	
2001	0.118 (0.0512)	0.0252 (0.0446)	0.109 (0.0517)	-0.0945 (0.0491)	0.0252 (0.0446)		0.0304 (0.0419)	0.0330 (0.0429)	0.0261 (0.0442)	
2002	0.0241 (0.0721)	0.0725 (0.0504)	0.123 (0.0641)	-0.0458 (0.0660)	0.0730 (0.0503)		0.0666 (0.0487)	0.0855 (0.0493)	0.0730 (0.0499)	
2003	0.137 (0.0769)	0.0692 (0.0412)	0.0937 (0.0577)	-0.0139 (0.0570)	0.0696 (0.0411)		0.0542 (0.0408)	0.0744 (0.0405)	0.0690 (0.0407)	
2004	0.191 (0.0603)	0.149 (0.0390)	0.177 (0.0533)	0.0746 (0.0531)	0.149 (0.0390)		0.133 (0.0380)	0.149 (0.0383)	0.145 (0.0386)	
2005	0.214 (0.0713)	0.238 (0.0494)	0.365 (0.0843)	0.0951 (0.0579)	0.238 (0.0494)		0.203 (0.0466)	0.221 (0.0473)	0.237 (0.0488)	0.0286 (0.135)
2006	0.514 (0.165)	0.231 (0.0569)	0.326 (0.0688)	0.179 (0.0837)	0.230 (0.0569)	0.0281 (0.0533)	0.202 (0.0642)	0.209 (0.0568)	0.228 (0.0563)	0.376 (0.128)
2007	0.400 (0.0776)	0.413 (0.0493)	0.495 (0.0680)	0.345 (0.0583)	0.413 (0.0493)	0.184 (0.0520)	0.369 (0.0478)	0.390 (0.0478)	0.409 (0.0489)	0.0848 (0.142)
2008	0.458 (0.0629)	0.429 (0.0474)	0.507 (0.0578)	0.366 (0.0672)	0.428 (0.0474)	0.171 (0.0518)	0.357 (0.0469)	0.386 (0.0468)	0.425 (0.0468)	0.535 (0.132)
2009	0.715 (0.0600)	0.556 (0.0526)	0.505 (0.0535)	0.597 (0.0823)	0.555 (0.0525)	0.274 (0.0630)	0.459 (0.0500)	0.489 (0.0494)	0.551 (0.0519)	0.587 (0.123)

**Table A.6:** Regression of performance gap in levels on controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Base	Control	Gasoline	Diesel	MonthYear	Age1	Age2	Age3	Location	First Year
2010	0.800 (0.0549)	0.700 (0.0472)	0.746 (0.0707)	0.708 (0.0577)	0.699 (0.0472)	0.393 (0.0622)	0.568 (0.0453)	0.604 (0.0455)	0.693 (0.0466)	0.768 (0.119)
2011	0.912 (0.0543)	0.762 (0.0489)	0.769 (0.0764)	0.782 (0.0585)	0.762 (0.0490)	0.422 (0.0663)	0.605 (0.0481)	0.653 (0.0483)	0.758 (0.0483)	0.887 (0.116)
2012	1.153 (0.0585)	0.977 (0.0584)	0.843 (0.0703)	1.093 (0.0695)	0.979 (0.0583)	0.604 (0.0783)	0.767 (0.0579)	0.829 (0.0582)	0.975 (0.0573)	1.220 (0.119)
2013	1.241 (0.0459)	1.142 (0.0518)	1.215 (0.0868)	1.173 (0.0639)	1.145 (0.0517)	0.739 (0.0786)	0.865 (0.0528)	0.935 (0.0532)	1.140 (0.0509)	1.266 (0.111)
2014	1.319 (0.0542)	1.251 (0.0892)	1.220 (0.134)	1.413 (0.0685)	1.251 (0.0891)	0.835 (0.110)	0.991 (0.0760)	0.991 (0.0883)	1.242 (0.0887)	1.259 (0.114)
Year:										
2005		0.00828 (0.00750)	0.0314 (0.00895)	-0.0126 (0.00927)						
2006		0.0388 (0.0130)	0.109 (0.0167)	-0.0142 (0.0153)						
2007		0.0451 (0.0165)	0.115 (0.0221)	-0.00926 (0.0199)						
2008		0.0622 (0.0190)	0.134 (0.0253)	0.00328 (0.0237)						
2009		0.0708 (0.0213)	0.128 (0.0290)	0.0128 (0.0278)						
2010		0.0981 (0.0232)	0.170 (0.0321)	0.0165 (0.0309)						

**Table A.6:** Regression of performance gap in levels on controls

	(1) Base	(2) Control	(3) Gasoline	(4) Diesel	(5) MonthYear	(6) Age1	(7) Age2	(8) Age3	(9) Location	(10) First Year
2011		0.105 (0.0254)	0.160 (0.0360)	0.0203 (0.0341)						
2012		0.132 (0.0283)	0.160 (0.0393)	0.0527 (0.0379)						
2013		0.145 (0.0310)	0.156 (0.0423)	0.0667 (0.0419)						
2014		0.0948 (0.0336)	0.0804 (0.0458)	0.0288 (0.0455)						
2015		0.127 (0.0357)	0.105 (0.0489)	0.0541 (0.0484)						
Month:										
2		-0.0526 (0.00110)	-0.0534 (0.00175)	-0.0526 (0.00139)						
3		-0.114 (0.00225)	-0.133 (0.00392)	-0.101 (0.00166)						
4		-0.208 (0.00401)	-0.265 (0.00569)	-0.168 (0.00246)						
5		-0.245 (0.00553)	-0.337 (0.00648)	-0.178 (0.00361)						
6		-0.259 (0.00617)	-0.350 (0.00830)	-0.195 (0.00412)						
7		-0.224 (0.00674)	-0.321 (0.00892)	-0.153 (0.00480)						

**Table A.6:** Regression of performance gap in levels on controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Base	Control	Gasoline	Diesel	MonthYear	Age1	Age2	Age3	Location	First Year
8		-0.252 (0.00662)	-0.358 (0.00803)	-0.175 (0.00438)						
9		-0.282 (0.00577)	-0.366 (0.00800)	-0.224 (0.00411)						
10		-0.237 (0.00408)	-0.288 (0.00599)	-0.202 (0.00358)						
11		-0.137 (0.00266)	-0.145 (0.00486)	-0.133 (0.00281)						
12		0.00161 (0.00204)	0.0203 (0.00298)	-0.0175 (0.00231)						
Diesel		-0.206 (0.0167)			-0.206 (0.0167)	-0.203 (0.0172)	-0.141 (0.0159)	-0.188 (0.0161)	-0.199 (0.0165)	
DistPerYear		-1.57e-06 (2.13e-07)	-3.26e-06 (3.49e-07)	-6.16e-07 (2.55e-07)	-1.56e-06 (2.13e-07)	-1.59e-06 (2.14e-07)	-1.12e-06 (2.20e-07)	-3.21e-06 (1.85e-07)	-1.35e-06 (2.05e-07)	
DistPerYear <sup>2</sup>		0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	
DistPerYear <sup>3</sup>		-0 (0)	-0 (0)	-0 (0)	-0 (0)	-0 (0)	-0 (0)	-0 (0)	-0 (0)	
Age:										
1						0.0115 (0.0205)		-0.0970 (0.00487)		
2						-0.0330 (0.0251)		-0.172 (0.00852)		

**Table A.6:** Regression of performance gap in levels on controls

	(1) Base	(2) Control	(3) Gasoline	(4) Diesel	(5) MonthYear	(6) Age1	(7) Age2	(8) Age3	(9) Location	(10) First Year
3						-0.0740 (0.0297)		-0.232 (0.0122)		
4						-0.111 (0.0346)		-0.291 (0.0165)		
5						-0.130 (0.0399)		-0.359 (0.0213)		
6						-0.151 (0.0431)		-0.446 (0.0290)		
7						-0.167 (0.0465)		-0.526 (0.0397)		
8						-0.188 (0.0501)		-0.638 (0.0463)		
9						-0.216 (0.0541)		-0.651 (0.0645)		
10						-0.262 (0.0596)		-0.770 (0.0904)		
11						-0.295 (0.0645)		-0.940 (0.136)		
12						-0.325 (0.0703)				
13						-0.331 (0.0784)				
14						-0.328				

**Table A.6:** Regression of performance gap in levels on controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Base	Control	Gasoline	Diesel	MonthYear	Age1	Age2	Age3	Location	First Year
15						(0.0862)				
						-0.299				
16						(0.0928)				
						-0.262				
17						(0.100)				
						-0.264				
						(0.107)				
Odometer							-2.32e-06			
							(1.15e-07)			
Constant	0.723	1.014	1.062	0.838						0.722
	(0.0356)	(0.0326)	(0.0390)	(0.0407)						(0.105)
Model F.E.	-	569	569	569	569	569	569	596	596	-
MonthYear F.E.	-	-	-	-	135	135	135	135	135	-
Location F.E.	-	-	-	-	-	-	-	-	19954	-
Observations	23,989,576	23,989,576	10,173,983	13,815,593	23,989,576	23,989,576	23,989,576	23,989,576	23,751,895	779,761
R-squared	0.117	0.211	0.182	0.259	0.212	0.212	0.220	0.215	0.225	0.191

Table reports coefficients of regressions equivalent to those reported in Table A.5 but the dependent variable is the level gap not %-gap. (1) includes only release year fixed effects (2) controls for year, month, fuel type, distance traveled per year and model fixed effects (absorbed), (3) equals (2) but limits the sample to gasoline engines, (4) equals (2) but limits the sample to diesel engines, (5) includes month by year fixed effects and model fixed effects (both absorbed), (6) equals (5) but includes age dummies, (7) equals (5) but includes odometer reading, (8) equals (5) but includes alternative age dummies, (9) equals (5) but adds fuel station fixed effects and (10) limits the sample to first year of driving.

**Table A.7:** Estimation Results

	Mean Taste		St. Dev.	
	Coeff.	St. Error	Coeff.	St. Error
Price/Inc.	-6.51	(0.45)		
Fuel Cost	-0.53	(0.01)	0.05	(0.03)
Horsepower	1.48	(0.21)	1.78	(0.10)
Weight	0.22	(0.21)	4.32	(0.16)
Footprint	0.88	(0.05)	0.58	(0.04)
Foreign	-0.92	(0.03)	0.02	(0.16)
Height	0.02	(0.02)		
Doors	0.50	(0.11)		

The table shows estimated taste parameters from a random coefficient logit estimation on the the car market for seven EU countries using data from 1998 to 2007. Taste distributions are assumed to be normal, and mean and standard deviations are estimated for selected characteristics. Additional controls are fuel type by market dummies, months for sale if less than 12, country fixed effects, linear time trend, body type fixed effects, vehicle class fixed effects and brand fixed effects. Model is estimated using a two-step GMM using approximate optimal instruments with sum of characteristics and cost shifter instruments for prices.

**Table A.8:** Diversion Ratios

	Own Price Elast.	Outside Good	FC Q1	FC Q2	FC Q3	FC Q4
Fuel Cons Q1	-4.47	0.45	0.15	0.21	0.12	0.08
Fuel Cons Q2	-4.78	0.41	0.15	0.22	0.13	0.09
Fuel Cons Q3	-5.30	0.37	0.16	0.23	0.13	0.10
Fuel Cons Q4	-7.36	0.29	0.17	0.26	0.16	0.13

The table shows the average own price elasticity in each of the fuel consumption quartiles in the vehicles included in the counterfactual (Netherlands, year 2007). Column II gives the average diversion ratio to the outside good, while Column III-VI give diversion ratios from the row fuel consumption quartile to the column fuel consumption quartile. The diversion ratio is defined as  $(\partial s_k / \partial p_j) / |\partial s_j / \partial p_j|$  so that Coluln II-VI sum up to 1.

## A.7 Additional Figures



**Table A.9:** Diversion Ratios specification IV Table 2

	Own Price Elast.	Outside Good	FC Q1	FC Q2	FC Q3	FC Q4
Fuel Cons Q1	-4.24	0.03	0.28	0.38	0.20	0.11
Fuel Cons Q2	-4.55	0.02	0.27	0.37	0.20	0.14
Fuel Cons Q3	-5.04	0.01	0.25	0.37	0.22	0.15
Fuel Cons Q4	-7.01	0.00	0.21	0.35	0.23	0.20

The table shows the average own price elasticity in each of the fuel consumption quartiles in the vehicles included in the counterfactual (Netherlands, year 2007) when the demand is estimated with a low 10% outside good share. Column II gives the average diversion ratio to the outside good, while Column III-VI give diversion ratios from the row fuel consumption quartile to the column fuel consumption quartile. The diversion ratio is defined as  $(\partial s_k / \partial p_j) / |\partial s_j / \partial p_j|$  so that Column II-VI sum up to 1.

**Table A.10:** Consumer surplus changes from gaming under alternative policies

	Policy Target		
	3%	5%	10%
Attribute-Based Standard			
Total	33.19	208.27	608.74
Choice Distortion	-17.77	-16.72	-13.47
Price Effect	50.96	224.99	622.22
Flat Standard			
Total	3.92	107.93	398.69
Choice Distortion	-17.58	-16.93	-15.01
Price Effect	21.50	124.86	413.70
Fuel Tax			
Total	-24.59	-24.94	-25.80
Choice Distortion	-17.42	-17.66	-18.23
Price Effect	-7.16	-7.29	-7.57

This Table presents consumer surplus changes when there is industry wide 5% gaming and policies with varying stringency: in Column (1) the target is a 3% reduction in emissions, Column (2) a 5% reduction (as in the previous Tables) and Column (3) a 10% reduction. The first panel shows the consumer surplus changes when firms game and there is an attribute-based standard, the second panel shows the effect of gaming when there is a flat standard, the third panel shows the effect of gaming when the policy is a fuel tax rather than a standard.

**Figure A.1:** Release year coefficients from fixed effect regressions

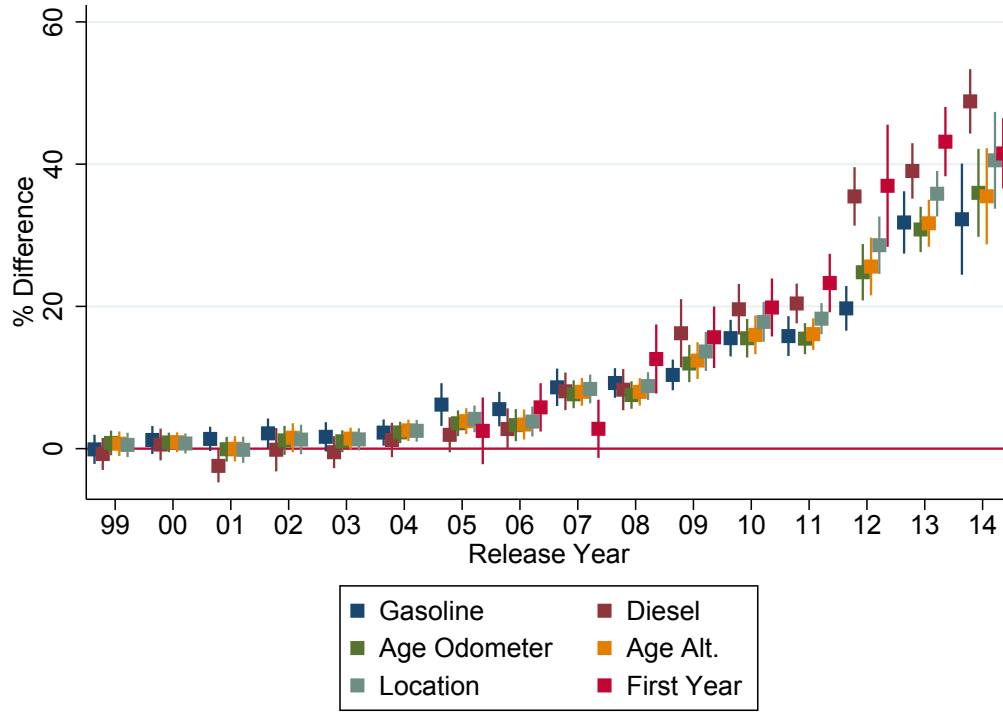


Figure plots coefficients from a regression of the performance gap ( $d_{nij}$ ) on release year fixed effects. Coefficients correspond to regressions (3),(4),(7),(8),(9) and (10) and from Table A.5.

**Figure A.2:** Release year coefficients from fixed effect regressions

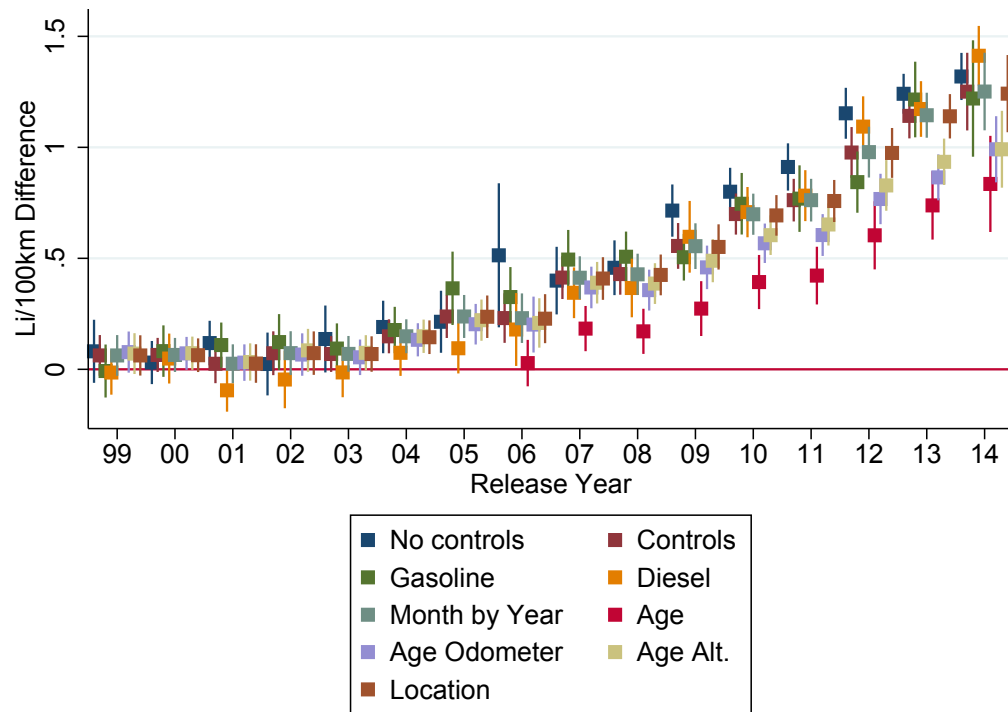


Figure plots coefficients from a regression of the performance gap in li/100km, on release year fixed effects. Coefficients correspond to regressions from Table A.6.

**Figure A.3:** Gap between on-road and official fuel consumption per firm in levels




Figure shows estimated coefficients and standard errors from regressing the performance gap in levels on three sets of model release years (early, middle and late) per automaker, the data is restricted to the ten brands with most observations.

**Figure A.4:** Gap between on-road and official fuel consumption per weight quartile




Figure shows estimated coefficients and standard errors from regressing the performance gap on three sets of model release years (early, middle and late) per weight quartile.

**Figure A.5:** Gap between on-road and official fuel consumption per weight quartile




Figure shows estimated coefficients and standard errors from regressing the performance gap in levels on three sets of model release years (early, middle and late) per weight quartile.

## Response to Editor and Overview

Thank you very much for the opportunity to revise our paper. All three referees provided many valuable suggestions. We have made substantial modifications of the paper to address these excellent comments. We believe that the revisions have substantially improved the paper. The major improvements are as follows:

1. Reduced form (Section 2): we added a variety of robustness checks to address several concerns in both the main text and appendix. We also present more results graphically. In brief, our headline finding is robust to all of the additional checks.
2. Theory (Section 3): we have consolidated this section and added new material that makes clear where simplifying assumptions might impact the theoretical results.
3. Simulation (Section 4): we added a parallel set of simulations in which firms comply by adding technology rather than “mix shifting.” These results align with our theory and the prior results.

Next we provide point by point responses to your comments. We have replicated your comments in shaded boxes, numbered as E.x with your original numbered points. Our responses follow.

**E.1** Please compress the literature review at the end of the Introduction. It is almost two full pages right now. One option would be to include more detailed literature review in the Online Appendix and refer to it in the main text. In any case, please limit this part of the Introduction to at most a single page.

We have sharply condensed this material and it is now right around 1 page of main text.

**E.2** The abstract says “on net, consumers benefit from gaming,” but I would like you to consider using “buyers” instead, since the thoughtful distinction between consumer surplus and buyer surplus runs throughout the paper.

Thank you. We made that change.

**E.3** R2 comment 2(f) discusses “lack of manipulation in pre-period,” but my understanding of model’s predictions is that firm can be expected to game a little bit in pre-period (depending on cost of gaming); but, when fuel economy standard are introduced or tightened, then firms are expected to respond by gaming more. So to me, this seems completely consistent with stylized model and suggests the paper could be clearer on this dimension (in terms of exposition).

We agree with your interpretation. In responding to the referee we have added clarifying text at the end of our discussion of the firm's first-order condition for gaming, where we clarify the two incentives for gaming (p. 19):

In sum, firms have two benefits from gaming. If consumers are fooled (at least partly), then gaming increases demand. Even if consumers are completely aware of gaming, if there is a policy, firms may game in order to minimize compliance costs.<sup>44</sup>

**E.4** Please follow R2 comment 2(g) on combining Propositions and 2(h) on consolidating simulation section.

We have eliminated Proposition 3, which was nested in Proposition 1 as R2.2g indicated. We did this by pointing out that the cost parameter  $\beta$  can be interpreted as tax inclusive, so that the desired points can be made with direct reference to Proposition 1 without additional notation. This allowed us to shorten the related section 3.6 (p. 23) by half (about half of a page).

In terms of the simulation section, Referee comment 2.2h suggests that we “consolidate the simulation section” without suggesting what to cut, but instead suggesting that we add additional caveats. We looked long at this material and do not see any section that can be removed (to an appendix) without creating additional confusion or striking what we believe is a substantive result. Indeed, Referee 1 and 3 ask us to add both additional results and more detailed explanation of the demand estimation in the main paper. What we have done is add a substantial new set of parallel results based on allowing compliance to come through technology adoption, as well as addressing a number of smaller comments, while adding 1 page net length to the section after trimming where we could.

**E.5** On the third part of the paper, two referees mention that firms can respond in practice by both changing actual fuel economy as well as changing the amount of “gaming.” I would like this comment addressed in a meaningful way. R3 mentions this in final section of report, and this is also one of main comments of R1 in letter to me. I do not want to “micro-manage” this part of the revision, but this is one of the most important areas of improvement that I will focus on when evaluating the revised paper. I realize that you cannot make much progress estimating the “cost of gaming,” but the entire third part is really a suggestive calibration, anyway, so you could explore sensitivity to different assumptions on shape of this cost function. I agree with referees that this is important since the time series shows change in actual fuel economy within a car model (across release years). So this seems important to capture in your analysis, both conceptually and empirically.

---

<sup>44</sup>The performance gap in early years of our sample is consistent with gaming to fool consumers, but it is also possible that the test is simply not well calibrated to measure typical driving patterns in the Netherlands.



We have added a parallel set of simulations that allow firms to comply with a regulation by adding costly technological improvements to raise the fuel economy of their vehicles, rather than complying through a mix shifting strategy. To do so, we need to calibrate the cost of increasing fuel economy. We use a cost curve for technological improvements that was estimated by regulators as part of the evaluation of the policy. (The same approach has been pursued by several papers in the US that study fuel economy standards.) We discuss our approach more fully in section 4.1.

In terms of results, we have revised the main simulation table, Table 2, to reflect various changes. The middle panel of that table now reports results from four simulations that assume firms comply by adding technology instead of by mix shifting. When we introduce gaming, firms need to spend less on technology improvements. This creates cost savings pass through that has the same dynamics as in our original version where firms comply by mix shifting.

In using an engineering cost curve, we have to confront an issue that is at the core of a substantial literature around what is called the “energy efficiency gap.” It is quite common that engineering cost estimates imply that energy efficiency can be improved at costs that are a modest fraction of the implied benefits of improved energy efficiency. That is, energy efficiency has implied private benefits. This is true in our data—the engineering cost estimates are below the private savings from reduced fuel consumption.

The interpretation of the energy efficiency paradox is hotly contested and long debated. One interpretation is that it is evidence of a market failure—consumers, for example, undervalue energy efficiency and forcing greater energy efficiency into the market solves this “externality” and boosts welfare. A second interpretation is that the engineering estimates are badly biased or fail to take account of other attributes (e.g., performance).

We do not have anything to add to this debate, but we must nevertheless decide whether or not to take the engineering estimates at face value. We thus report two versions of our results. The first takes the cost curve at face value. This leads to quite different welfare results in our model. Intuitively, in this scenario, there is an additional market failure and forcing technology into the vehicle has a social benefit related to closing the energy efficiency gap. There is no price effect benefit from gaming because the regulation is not costly to the buyer. We report this version of our result in Table 2 (column VI) and discuss it in section 4.2 (p. 31).

Then we present another version that shifts the engineering cost curve up so that the market behavior before the policy is break even. Essentially we assume that there is no externality market failure and that the technology costs are equal to net present private benefits from fuel savings before policy is introduced. We then simulate outcomes when policy forces technology adoption that is privately costly (but may be justified by reduced emissions). We show these results in Table 2 (columns VII and VIII). In that case, we get back the qualitative results that we had in our original estimates based on mix shifting—price effects are good for buyers and they substantially dominate choice distortions. These results are discussed in more detail in section 4.2 (p. 31).

**E.6** In addition to the previous comment, the referees mention several other limitations of theoretical model and simulation exercises (e.g., multiple firms, multiproduct firms, etc.). I would like you to treat these other comments (other than endogenous actual fuel economy) as largely expositional comments, rather than a call for much more development of the model and/or simulations. I will be satisfied if paper is just clearer about the assumptions and limitations of stylized model, and I would like you to read these referee comments as asking for more caveats. Future research can generalize the model and investigate sensitivity to different assumptions; I will be satisfied as long as the analysis of model makes clear that a richer model could very well deliver different results.

We have added a short subsection, titled “Note on model assumptions” (section 3.3, p. 19) that does just this. It highlights again our key restrictive assumptions and says that alternative versions could yield different results, along the lines of referee comments related to multiple products, fixed costs and compliance strategies.

**E.7** R1 (comment 2) suggests a potential correlation between sensitivity to price and sensitivity to gaming. I would like you to discuss this briefly (and intuitively) in main text, but you should feel free to treat a rigorous treatment of this as outside the scope of the paper.

In the new subsection on model assumptions mentioned in the prior item, we include a specific discussion of this issue (p. 20):

We model a representative buyer, but if buyers differ in their  $\alpha$ , additional strategic considerations might come into play. For example, if relatively naïve buyers are also more price sensitive, then gaming could draw more price sensitive buyers into the market, which might put downward pressure on prices, thereby benefitting inframarginal, sophisticated buyers. On the other hand, with buyer heterogeneity of this sort and multiple products, a seller might use gaming to differentiate products to screen buyer types and achieve second degree price discrimination.

We point to this again when explaining our simulations in a new footnote 28 (p. 25).

**E.8** As I read through the first part of the paper, I thought that a natural research design would be to compare across automakers that were differentially affected by the new policy based on sales-weighted fuel economy prior to the policy change. For example, if there was automaker with sales-weighted average already “in compliance,” then couldn’t that automaker serve as a control group? Is this feasible? This research design is conceptually similar to Amy Finkelstein’s QJE paper on the aggregate effects of Medicare. I would like a discussion of this design and/or other possible alternative research designs (either in main text or footnote), since I think many AEJ: Policy readers will find the pure time series analysis much too suggestive. I’m guessing that, in

reality, all automakers reacted similarly (perhaps because technology of gaming can be easily spread across models and automakers), but in any case I think this merits some more discussion.

Thank you for this comment. We long sought alternative identification strategies ourselves, but we face limitations of data and policy. We summarize the limitations of research designs that we explored in a new footnote 10 (p. 11):

Ideally one would like to leverage variation in the regulatory incentives across vehicles or firms to aid causal identification, but the features of the policy preclude several approaches. First, the regulation was announced in 2007 and was made fully binding in 2015, which means that there is not a sudden surprise policy shock to examine. Second, the regulation binds for the fleet wide average for an automaker, so automakers have common incentives to improve all vehicles. Moreover, the program lacks the non-linear notches used for identification in related contexts (Ito and Sallee 2018, Sallee and Slemrod 2012, Tanaka 2017). Third, the regulation was explicitly designed with a weight-based standard to reduce variation in compliance burdens across automakers, which sharply reduces the cross-firm variation in policy impacts. Moreover, the policy allows for “joint compliance” between automakers, which further limits cross-firm variation. As a result, no firms were already in compliance, and all firms had to make similar improvements.

**E.9** Please present Table 1 graphically so it is easier to compare across columns (you already do this for the “with controls vs. without controls,” but I would like to see more of these type of figures in both main text and Online Appendix). In Online Appendix, it would be particularly useful to see similar figures for comparisons across different combinations of controls.

Figure 3 now shows the five regressions included in Table 1. We report a series of additional robustness checks in Table A.5. Regressions not included in Figure 3 are plotted in Figure A.7. All regression coefficients for covariates are reported in Table A.5. Additionally, we now report parallel specifications with the fuel economy gap in levels (rather than percentages) in Appendix Table A.6 and with release year coefficients plotted in Appendix Figure A.2.

**E.10** Please include self-contained Online Appendix section with more details of EU policy. This can draw from some of the work currently cited in the paper.

We have added a section with a detailed description of the EU policy in appendix A.2. Footnote 6 refers the reader to this section.

**E.11** In reality (and I think in model, as well), firms have incentive to engage in “gaming” even if consumers are sophisticated, since this allows firm to avoid paying penalty. I would like discussion of model to be clearer about this case. Even though you don’t measure the cost of gaming empirically, when going through model you can describe how these costs translate into actual “gaming” response and how the magnitude of this response depends on consumer sophistication.

Your interpretation is correct. Firms may game when consumers are sophisticated if there is a policy incentive. We have modified our text and added a sentence in our description of the model setup (p. 18) to make sure this is clear:

“The key for our purpose is that we want to allow the *possibility* that firms might fool buyers, at least to some degree. This specification is appealing in part because, for all of our results, we will be able to see how effects differ when buyers are not fooled just by setting  $\alpha = 1$ . When buyers are not fooled at all, firms may nevertheless still choose to game in response to the policy incentives we describe below.”

**E.12** All of the referees and I are satisfied with how you are currently modeling consumers, since “gaming” can be partly a response by firms to the fact that they are able to “fool” at least some of the consumers some of the time. But I would like Section 2 to add a little more background on what consumers likely understood about actual fuel economy prior to the 2007 law. I do not consider myself a car expert, but whenever I have bought a car I have looked at the MPG on the sticker and assumed that it was an over-estimate of what I would actually experience when driving. I also recall reading consumer reports that carried out their own tests that I think were consistent with this “folk wisdom.” Were these kind of reports available during this time period? Ultimately, the paper does not have much concrete evidence about how much gaming actually erodes information, so I would like more narrative/qualitative discussion of this.

We discuss this point in Section 4.1 (p. 28) where we give a detailed discription of what information was available to buyers in our empirical setting. Given that this is an empirical discussion related to parameters in the buyer utility function, we decided to keep this in Section 4.1, rather than to move it to Section 2 as you suggest. We have also added a footnote that points to this discussion when we bring up buyer awareness in the theory, that is Footnote 21.

**E.13** I am interested in more sensitivity analysis for the fixed effects regressions in the first part of the paper. Ultimately, I see this as a conceptually similar to the standard age-cohort-time effects problem. With fully flexible age effects and time effects, you are not able to identify a full set of cohort effects without a normalization (same problem as in Card and Lemieux QJE paper <http://davidcard.berkeley.edu/papers/men-return-college.pdf>). Suppose you imposed all pre-1998

cohorts to be the equal to each other, and you included a full set of age effects and time effects, what would be effect on the cohort trends in “release year” cohort effect estimates? This feels like a more rigorous way to address potential confounding age effects that what is currently being done in the paper. In general, I would like identification of cohort effects to be described more rigorously in revision given potential confounding and under-identification. I also would like to see sensitivity to including full set of year x month fixed effects (in addition to current set of controls).

Regarding the last point—on year by month fixed effects—we added a specification (column 3) in Table 1 that does just this. It has almost no impact on our release coefficient estimates.

In terms of the age-period-cohort problem, thank you for drawing our attention to this. We had been largely unconcerned by age effects given conventional wisdom that the actual fuel consumption of vehicles does not decline noticeably with age. One of us is an author on a working paper that shows precisely this. But, in thinking about your comment, we realized that it is quite plausible that there are selection effects, by which we mean that older vehicles are driven in systematically different ways (i.e., less city driving, more aggressive stopping, etc.) that could yield differences in the measured performance gap. We thus undertook a variety of strategies to try to deal with this issue, leading with the one you suggest.

Column (4) in Table 1 is now a specification that includes year by month fixed effects in addition to controls and age dummies. In order to identify the age dummies we follow your suggestion and group all pre-2006 cohorts to be equal to each other. The cohort (release year) trend is robust. We show three additional specifications in Appendix Table A.5 that deal with age effects in different ways. First, we specify age based on the first year observed in the sample. Because cohorts enter the data over several years (and a vehicle doesn’t necessarily age before it is sold) we can identify both age and cohort effects. Second, if we think that the age effects are captured by the odometer of the car rather than calendar age, then we can introduce odometer controls and still achieve identification because there is variation within release year by calendar year cells in odometer. This is an imperfect solution, however, because the identification is coming off of drivers who drive vehicles different distances per year, which is probably correlated with how much city or highway driving they do, which can also impact the performance gap. Third, we include only data from the first year of a vehicle’s life, which eliminates age concerns but does not allow for time effects. (This version was included in the first version of the paper.)

All these results are summarized in the revised Section 2.6. We discuss our priors on why age effects are unlikely to be a major factor and then discuss the results as follows (p. 12):

Our specifications above identify release year effects controlling for time period shocks, but an additional concern arises if fuel consumption changes with the age of a vehicle. In that case, we face a version of the well known age-cohort-time identification problem (see, e.g., Card and Lemieux 2001). Emissions of local air pollution are known to degrade substantially with vehicle age, but fuel economy is typically assumed to be

stable. Even so, there could be a form of selection through which older vehicles are driven in different patterns or styles on average, which will manifest as an age effect in our on-road performance data. Regardless, we would like to be able to control for age effects, but age fixed effects are perfectly collinear with release year and time fixed effects. We cannot separately identify all effects without some additional restriction. One approach is to assume that release year effects are completely flat prior to 2006, so that a single release year dummy for that period is included. A full set of age effects can be estimated based on that period, and age and time effects can both be included while still identifying the remaining release year effects. Column (4) reports results from a specification that does that. Results are qualitatively similar, though the magnitude of our estimated performance gap does change somewhat, falling from 41% to 34% as compared to column (3), which has the same time period controls. As discussed below, we provide alternative approaches of addressing age effects in the appendix, which yield qualitatively similar results to the results reported here.

Then, in discussing the additional age controls that are in the appendix, we say (p. 15):

We also introduce three additional methods of controlling for age. First, we define age as the year a license plate (combination of vehicle and driver) is first observed in our data set, which differs from the release year. This controls for the length of time a driver has had the vehicle. In this way we can include a full set of release, age and time fixed effects because there is variation in age within release year and time, based on when a driver obtained a vehicle. Second, we define age in terms of odometer instead of calendar years and introduce the level of the odometer as a control.<sup>45</sup> Third, we keep only observations of  $d_{nij}$  that take place in the release year of the vehicle, so that we are capturing fuel consumption gaps only among the newest cars. This eliminates concern about age as a confounder, but it does not allow us to control for time period effects. As a final robustness check, we add almost twenty thousand fixed effects that control for the location of the fuel station for each transaction. Figure plots the release year coefficients for all specifications not included in Table 1. In every specification we find a qualitatively similar trend in the performance gap by release year, with a sharp increase in the middle of the sample that ends with a large gap on the order of 30 to 50% by 2014.

**E.14** R3 provides five numbered comments on the “second part” of paper (model). I see 1, 3, and 4 as outside of the scope of the paper, and so you should feel free to respond that way. However, I would like to see more thorough responses to both 2 and 5.

<sup>45</sup>Note that the age dummies have surprising negative magnitudes in this regression. These can only be due to compositional changes within model groups because we find age trends to be flat within engine versions.

(We have renumbered R3’s comments, so the 5 comments you refer to are now numbered R3.3-R3.7 and your instructions suggest we should respond thoroughly to R3.4 and R3.7. But, for clarity of addressing your point, we keep the original numbering here.)

Point 2 is about our comparative statics, which we improved. This was really just a poor choice of notation on our part and we have changed it and updated the discussion.

Point 5 asks us to pursue a complete welfare analysis of the policy, shifting our focus from buyer surplus to overall social welfare. We consider this to be a very large undertaking, as it requires us to take a stand on fixed costs, another stand on the costs of gaming, and to build a new supply side model that optimizes compliance. We think that is outside the scope of the paper given your request to de-emphasize and condense this part of the paper in favor of the reduced form section and theory.

In contrast, Point 4, which you identified as “outside of scope”, is just a point a request that we clarify language delineating two types of “consumer awareness”, one regarding gaming and the other regarding the proper valuation of fuel economy. We thus speculate that you might have intended to tell us to thoroughly address Points 2 and 4, while pointing to 1, 3 and 5 as outside of scope.

Based on that interpretation, we have corrected the language issues implied in Points 2 and 4, and we have included substantial discussion and clarification about the modeling points in 1, 3 and 5. Details of our responses are in the referee letter and are not reproduced fully here.

## Responses to Referee 1

Thank you very much for your insightful comments and suggestions. In this letter, we have included your full set of comments and quoted as much material as possible in the hopes of minimizing the need for you to reference back to the paper in order to understand our responses.

We believe that the revisions have substantially improved the paper. The major improvements are as follows:

1. Reduced form (Section 2): we added a variety of robustness checks to address several concerns in both the main text and appendix. We also present more results graphically. In brief, our headline finding is robust to all of the additional checks.
2. Theory (Section 3): we have consolidated this section and added new material that makes clear where simplifying assumptions might impact the theoretical results.
3. Simulation (Section 4): we added a parallel set of simulations in which firms comply by adding technology rather than “mix shifting.” These results align with our theory and the prior results.

To explain our responses to your detailed comments, we have replicated your comments here in shaded boxes, numbered as R1.x, with x representing the item number from your report. Our responses follow.

**R1.1** Figure 1 suggests that once the emission caps were adopted there was an improvement in on-road fuel efficiency (though this was smaller than the improvement in the official rating). In the third part of the paper you do not allow for firms to reoptimize the true fuel efficiency, meaning, conditional on a given level of gaming, the only way to achieve a binding emissions standard is through price adjustments that skew sales towards more fuel efficient vehicles. It seems important to provide evidence that this happened in practice. If in fact firms did not adjust prices in this way, one of the main effects in the simulations would be misleading and purely an artefact of how you set the problem up, rather than something that was an important phenomenon in reality.

Thank you for this comment. In the revised version, we have added a second set of simulations that suppose that firms comply with the regulation not by mix shifting, but by adding technology to their products. In a scenario in which the added technology is privately costly, we get results that are qualitatively similar to our original exercise. The new results are reported in a second panel of our results in Table 2, alongside the original mix shifting scenario.

To do this, we use engineering estimates of the cost of adding technology to vehicles from regulatory documents. This raises a set of issues related to what is known as the energy efficiency gap. The engineering cost estimates imply that fuel savings from added technology are worth more than the estimated cost. That is, there is a *private* benefit of technology forcing. There is a



substantial literature on the energy efficiency gap that argues over the interpretation of technologies with engineering costs that are “negative cost” but enjoy low take up in the market. This may represent a market failure (most often described as an “externality” in which buyers undervalue energy efficiency due to some limited rationality). If we use the original low cost curve, this, not surprisingly, changes our results. A policy that forces technology is not costly, but rather is cost saving, and so gaming to avoid it is bad for buyers. We report results from this interpretation in Table 2 (column VI), discuss the setup in section 4.1 (p. 26), and discuss the results in section 4.2 (p. 31).

As an alternative, if one interprets the energy efficiency gap as evidence that engineering costs are biased downwards, then there is not an externality to be corrected. To capture that situation, we shift upwards the engineering cost curve until marginal costs are equal to benefits for the technology deployed prior to policy. Policy is then forcing privately costly (but possibly socially beneficial) fuel economy improvements. Those results are also contained in Table 2 (column VII and VIII), with results discussed in section 4.2 (p. 31). In this case, as alluded to above, the same qualitative results emerge as in our original scenario—gaming creates a price effect that benefits consumers and significantly outweighs choice distortions.

**R1.2** The analysis in the second part of the paper relies on there being no correlation in preferences over price and fuel efficiency, and this kind of correlation is ruled out in the third part of the paper. Yet it seems plausible (likely?) that more price sensitive consumers will have stronger preferences over fuel efficiency. This means that gaming will tend to attract more price elastic consumers into the market which (in the absence of an emissions cap) may mean the equilibrium price falls rather than rises. I can see that formally modelling this in part two of the paper would significantly complicate the analysis. However, I would like to see a discussion of this restriction and the likely impacts of it not holding (which I think probably would reinforce your results). In part three of the paper it would, in principle, be possible to allow for a price random coefficient correlated with the fuel efficiency random coefficient, though it may be the variation in the data is not sufficiently rich to convincingly estimate this.

In the theory section, we have added a new subsection, titled “Note on model assumptions” (section 3.3, p. 19), that makes clearer how relaxing some assumptions might lead to alternative results. We mention the point you raise about correlated preferences there, as follows (p. 20):

We model a representative buyer, but if buyers differ in their  $\alpha$ , additional strategic considerations might come into play. For example, if relatively naïve buyers are also more price sensitive, then gaming could draw more price sensitive buyers into the market, which might put downward pressure on prices, thereby benefitting inframarginal, sophisticated buyers. On the other hand, with buyer heterogeneity of this sort and multiple products, a seller might use gaming to differentiate products to screen buyer types and achieve second degree price discrimination.

In the simulation section we added Footnote 28 to explicitly explain that we do not estimate correlations between preferences. We agree that it would in principle be possible to estimate covariances between the preferences for fuel consumption and prices. However, in previous work we have found these covariance to be estimated very imprecisely so that we are not able to make statistical claims about these patterns. Data that has richer information on individual buyer characteristics (‘micro choice data’) might be better suited to answer questions related to individual price elasticities and pass through of gaming.

**R1.3** I appreciate that you do not want to devote too much space in the paper to estimation of and results from the demand model. However, I think it would be useful to include some more details than currently are in the paper. In particular, it would be useful to see some cross price elasticities. For the empirical analysis to be credible it is useful to see evidence that the demand estimates are capturing realistic/plausible switching patterns.

In response to comments of the editor and other referees, we have kept discussion of the demand estimation in the paper as concise as possible. To give more detail about the estimation results and resulting substitution patterns we have added Tables A.8 and A.9 to the Appendix. These tables summarize the own price elasticities per fuel consumption quartile. To summarize cross-price elasticities we have added the diversion ratios between the fuel consumption quartiles and the outside good. We refer to these tables in section 4.1. The mean of the estimated own price elasticities is -5.45, in line with the previous literature. We find that cross price elasticities are higher for vehicles with similar fuel consumption and that the outside good is a stronger substitute for low fuel consumption vehicles.

**R1.4** In the supply model, how do you deal with multiple geographical markets? Presumably the price and cost of the same model differs from one country to the next. Do you allow for this in the simulations?

Thank you for this comment. Yes, prices and marginal costs (as backed out from the demand model) vary across markets. In the counterfactual we select the Netherlands 2007 as the market for which to compute counterfactual scenarios. We choose this baseline market both because the Netherlands is quite representative of the EU market overall and because it is the source of our data on gaming. We choose 2007 to represent the most recent market before the major uptick in gaming. We have specified this in Footnote 30 and discuss this again on page 30 (first paragraph of subsection 4.1).

**R1.Smaller Comments** i) In the second half of page 7, the distinction between vehicle types and vehicles/cars was not clear. What precisely is it that distinguishes a vehicle type from a vehicle?

ii) The FOC for  $g$  in the top half of page 18 appears to be missing the regulatory constraint term from the price-cost margin in the second term on the RHS.

i. There is no difference between a vehicle and a vehicle type, there was an inconsistent description of vehicle and vehicle type in two paragraphs on (what was formerly) page 7 and we have corrected that description. We searched for other instances of inconsistency and sought to harmonize our language throughout.

ii. Regarding the first-order condition, this indeed was a typo with a missing term. We have corrected it. Thank you for catching that.

## Responses to Referee 2

Thank you very much for your insightful comments and suggestions. In this letter, we have included your full set of comments and quoted as much material as possible in the hopes of minimizing the need for you to reference back to the paper in order to understand our responses.

We believe that the revisions have substantially improved the paper. The major improvements are as follows:

1. Reduced form (Section 2): we added a variety of robustness checks to address several concerns in both the main text and appendix. We also present more results graphically. In brief, our headline finding is robust to all of the additional checks.
2. Theory (Section 3): we have consolidated this section and added new material that makes clear where simplifying assumptions might impact the theoretical results.
3. Simulation (Section 4): we added a parallel set of simulations in which firms comply by adding technology rather than “mix shifting.” These results align with our theory and the prior results.

To explain our responses to your detailed comments, we have replicated your comments here in shaded boxes, numbered as R2.x, with x representing the item number from your report. Our responses follow.

**R2.1a** The empirical exercise brings novel data to bear and clearly illustrates an increase in the performance gap post 2007. A more thorough investigation would strengthen the authors’ interpretation of the results.

a. If odometer readings are measured with error, the manuscript should note that the mean of  $d_{nij}$  will be positive and increase as measurement error increases. This doesn’t necessarily invalidate the empirical analysis, as one would need measurement error to be correlated with “release year” but should be explicitly noted and complicated interpretation of the “performance gap” pre-2007.

Yes, measurement error is a plausible concern given the nature of our microdata. We think the best check on this is the long difference specification. We include discussion of this issue and describe the long difference specification in the text (p. 12) and footnote 13:

Finally, column (5) reports results from a regression where we use long differences to construct the fuel performance gap, instead of differences between fuel station visits. Odometer readings are self-reported and errors in the readings result in positive measures of the gaps. Resulting attenuation could be responsible for some of the estimated performance gap prior to the increase in gaming, and correlation between reporting errors and release years could bias our findings about the trend over release years. The

analysis with long differences is more robust to individual reporting errors. The drawback is that it does not allow us to control for time of driving effects in the regressions. Regression (5) shows that our findings are robust: we get nearly identical estimates when we use long differences, taken between the last and first fuel station visit for each driver.

**R2.1b** Please plot the month and year fixed effects estimating in column (2) of Table 1 and cite as an appendix table. That would seem the most direct way to address questions related to the traffic, fuel prices and changing driving behavior common to all vehicles.

We have included these coefficients, along with other control variables, for all specifications in a new Appendix Table [A.5](#).

**R2.1c** Please include controls for columns (3)–(4) of table 1. Given the potential sources of bias highlighted on page 11, it’s natural to include the controls here as well.

The specific regressions you refer to in your original comment have been moved to Appendix Table [A.5](#). Nevertheless, we have adopted the spirit of your suggestion by including those controls in our various robustness checks reported in Tables [1](#) and [A.5](#).

The gasoline and diesel specifications you refer to are only reported in the Appendix following other comments. Table [A.5](#) includes the controls for those specifications following your suggestion.

**R2.1d** Is there any concern that a firm might reallocate its vehicle fleet as vehicles that are perceived to have higher fuel efficiency are introduced? One might consider driver\*vehicle fixed effects that are then regressed upon year of release, or plotting mean performance gap over time, by release year.

Our main regressions include model fixed effects. The reallocation you describe thus must be a concern in vehicle models (i.e., more fuel efficient engines are introduced within all the Toyota Camry engines).

We think using driver\*vehicle fixed effects regressed on release year, as you suggest, is very similar to what we already do in Section 2.7. In that section we construct vehicle fixed effects estimates that take into account the underlying precision of the data provided by all drivers of that vehicle using an empirical Bayes approach. We then plot the distribution of these fixed effects over time in Figure [4](#).

In the results table from that exercise, Table [A.4](#), we have followed your suggestion and added results from a two-step procedure. We discuss this procedure in Footnote [17](#). The results of the two-step procedure are very similar to the Bayes correction.

**R2.1e** Plot the ramp up of the standards in Figure 1 as well. The upward trend in figures 1 and 3 suggests a gradual increase in gaming post-2007 driven, presumably, by ratcheting of the benchmark.

The policy was announced in 2007 but did not start to bind until 2012. The policy schedule has not changed since 2007. There was a phase in between 2012 and 2015, but all automakers complied with the full standard by 2014. We do not think the increase in gaming between 2007 and 2015 is necessarily a response to a ratcheting of the benchmark (though they are certainly correlated in the time series). Instead, we think the gradual increase is due to automakers pushing their products forward with the goal of achieving compliance by the end of the period. In addition to the standard, there are several changes in national taxes after the announcement of the standard in 2007 but these are much more difficult to plot. We have added a more detailed discussion of the policy in Appendix A.2 to make all this clearer.

**R2.1f** 2006 seems an outlier in figure 3, potentially worth a mention or further investigation.

The standard error in 2006 is indeed larger than in other years but this disappears when we introduce controls. In 2006 a larger share of observations come from newly introduced models, but we find no other anomaly that might explain the lower precision in that year. We now mention this in footnote 11 (p. 11):

The estimate for release year 2006 is considerably less precise than other years. The Travelcard data contain a larger number of observations from new model releases that year, which might explain the imprecision. Regardless of the cause, the anomaly is eliminated once other controls are added.

**R2.1g** Does each model-year of a vehicle as having a unique fuel consumption rating? If not, how should we think about model-years spanning the introduction of the regulations.

The EU does not have the same strict model-year definition as the US, both in terms of our data and in terms of customer experience. Models are sold with several engine versions and over time these versions are updated. This means that an engine version can be for sale for several years with the same fuel consumption rating. We note this now in footnote 12 (p. 11).

Note that the standard becomes binding only in 2015 and is a fleet wide regulation. Firms thus need to reduce their average fuel economy by a sufficient amount rather than the fuel economy of any specific vehicle. The targets are set specifically over several years so that firms can adjust within their redesign period.

**R2.1h** Are there meaningful differences in regulatory incentives at the member- state level?

Yes, there are meaningful differences in regulatory incentives between member states. We mention this in the main text on page 5, where we cite Gerlagh et al. (Forthcoming), which is a paper that discusses these policies in detail. We describe the policy environment in more detail in the new appendix section A.2. (Note that there are policy differences across states, but we have data on road performance from only the Netherlands.)

**R2.2a** The theory model considers a monopolist that manipulates an attribute attractive to consumers and documents two sources of welfare loss: (1) consumers may overpurchase the good and (2) the firm may be able to increase price for the seemingly higher quality product.

a. The authors note that extending the model to multiple goods or multiple products do not qualitatively change the results, but that would seem depending on a unique outcome to the firm optimization problem. Unless accompanied by formal results, the authors should present the model as it stands and note that it is a simplifying assumption.

We added a new section 3.3 where we emphasize the assumptions of our model and note that alternative assumptions might yield different results. (See next point for more detail.)

In addition, we looked for other places where we may have overstated the extent of our results. We cut the following phrase (p. 18): “but the basic insights of our model do not depend on the form of competition” where we discuss our use of a monopolistic seller.

**R2.2b** As an illustration of the point above, with multiple goods and heterogeneity in sophistication, a profit maximizing monopolist might find it optimal to differentiate their products along the degree of mis-representation to separate naïve consumers from sophisticated consumers.

This is an excellent point. In our new subsection 3.3 (p. 20), we added the following text:

We model a representative buyer, but if buyers differ in their  $\alpha$ , additional strategic considerations might come into play. For example, if relatively naïve buyers are also more price sensitive, then gaming could draw more price sensitive buyers into the market, which might put downward pressure on prices, thereby benefitting inframarginal, sophisticated buyers. On the other hand, with buyer heterogeneity of this sort and multiple products, a seller might use gaming to differentiate products to screen buyer types and achieve second degree price discrimination.

**R2.2c** Section 3.6 oversells the likelihood that gaming could increase emissions by expanding the size of the market, given that the European car market is mature (so overpurchasing of new vehicles

will largely be offset by reducing the intensity of use of older vehicles) and that there is little evidence of a competitive benefit to gaming (and hence a robust consumer response) in the pre-period. That said, the cross-industry implications of illusory emissions reduction would continue to hold.

You raise a good point. We backed off on our statement by deleting the following text: “Given the extensive degree of gaming in the European vehicle market, this scenario is not necessarily far fetched.” We replaced that with (p. 24):

Because the true attribute is unchanged from the baseline, the market expansion implies an unambiguous total increase in pollution caused by the policy. The final impact on emissions depends on general equilibrium factors that are hard to predict. In a mature automobile market, an increase in new sales may lead primarily to a reduction in use of old vehicles. Vehicle retirement, however, might involve old cars being exported to lower-income markets.

**R2.2d** As formalized in the expression for  $h'(g)$  on 19, there are two benefits of gaming: the regulatory benefit and the benefit of making the vehicle more attractive to consumers. The relative absence of gaming in the pre period suggests that the second term in the expression for  $h'(g)$  is small in magnitude.

We agree with this interpretation.  $h'(g)$  is the marginal cost of gaming, but at the optimum it is equal to the benefits. We tried to improve exposition of this by adding language that echoes your own in that section (p. 19):

In sum, firms have two benefits from gaming. If consumers are fooled (at least partly), then gaming increases demand. Even if consumers are completely aware of gaming, if there is a policy, firms may game in order to minimize compliance costs.<sup>46</sup>

**R2.2e** The results of the simulation actually echo this intuition—the “price effect” of manipulation swamps the effect on the distortion of purchases.

**R2.2f** Yet, one of the other conclusions of the structural model, that manipulation would be unilaterally profit maximizing in the absence of the regulatory guideline seems at odds with the relatively lack of manipulation of fuel efficiency measurements by the firms in the pre-period.

We answer points e and f together here. You make a good point. We think that part of the explanation is a safety in numbers issue, whereby firms anticipate a much heavier penalty if they act alone than if they are gaming the same amount as the rest of the industry. Indeed, there has

---

<sup>46</sup>The performance gap in early years of our sample is consistent with gaming to fool consumers, but it is also possible that the test is simply not well calibrated to measure typical driving patterns in the Netherlands.



been little punitive action against the automakers in response to this issue to date. We have added a discussion of this in our section on modeling assumptions (p. 20):

Finally, regulatory penalties might be depend on whether an individual firm games significantly more than the rest of industry, as a regulator may find it impractical to levy punitive damages on the whole industry simultaneously. For example, in the Volkswagen diesel case in the US, regulators forced quick removal of non-compliant vehicles from the market. Had the problem affected a majority of new vehicles, it is not plausible to imagine the same penalty. We suspect that part of the evolution of gaming was enabled by firm’s feeling insulated from penalty by their competitors behavior, but this interaction is not present in our monopoly model.

**R2.2g** Proposition 3 nests proposition 1. Why not just explore the former and examine the case in which there are no fuel taxes as a special case?

In response to this suggestion, we have eliminated what was Proposition 3. We now simply emphasize that the cost parameter  $\beta$  can be interpreted as tax inclusive, so the relevant results can be read directly from Proposition 1 without additional notation. We then shortened the section on the tax case accordingly.

**R2.2h** I would consolidate the simulation section and qualify the results. If the exercise can speak directly to the degree of naivete as captured by alpha, I think the value of the simulation would be greater. But, heterogeneity in the degree of naivete significantly complicates the multiproduct case. In addition, the supply side of the model is uncaptured and clearly relevant. Although significant gaming does occur, the regulations do shift the actual mix of vehicles offered. Absent a framework for thinking about the supply side, it’s unclear how to think about counterfactuals in which regulatory stringency varies.

Your point is well taken, and we have attempted to be very clear about the limitations of our simulation. Even so, we believe it is a useful exercise for exploring the plausible quantitative significance of the effects we describe in the theory.

Regarding changes in product offerings, our revised simulation section now adds a set of alternative simulations that assume that automakers comply by adding costly technology to each vehicle, rather than by mix shifting. Results are presented in an expanded Table 2 that compares results under this scenario to our original. Results are discussed in section 4.2 (p. 31).

Regarding heterogeneity, we highlight the points you make here in the theory discussion (p. 20) and again in the empirical section in a new footnote 28 (p. 25), with the latter quoted here:

We estimate the model with aggregated country level market shares and have no information on individual household incomes. Our model therefore does not allow us to

study the distributional impacts of gaming. Related, we only estimate diagonal elements in the preference variance co-variance matrix and assume that there is no correlation between preferences for different attributes. We thus do not capture any correlation between price elasticities, fuel economy valuation or awareness of gaming ( $\alpha$ ), which means that strategic behavior based on heterogeneity in those parameters is absent from our simulations.

## Responses to Referee 3

Thank you very much for your insightful comments and suggestions. In this letter, we have included your full set of comments and quoted as much material as possible in the hopes of minimizing the need for you to reference back to the paper in order to understand our responses.

We believe that the revisions have substantially improved the paper. The major improvements are as follows:

1. Reduced form (Section 2): we added a variety of robustness checks to address several concerns in both the main text and appendix. We also present more results graphically. In brief, our headline finding is robust to all of the additional checks.
2. Theory (Section 3): we have consolidated this section and added new material that makes clear where simplifying assumptions might impact the theoretical results.
3. Simulation (Section 4): we added a parallel set of simulations in which firms comply by adding technology rather than “mix shifting.” These results align with our theory and the prior results.

To explain our responses to your detailed comments, we have replicated your comments here in shaded boxes, numbered as R3.x. You numbered your items in sections related to each part of the paper (e.g., there are 4 comments numbered 1), but we have renumbered them sequentially. Our responses follow.

**R3.1** Gaming is quantified with a proportional measure. This thus means that there is a lot of heterogeneity in how gaming translates in foregone fuel savings. For very efficient vehicles, 30% in gaming may imply small changes in lifetime fuel savings, but a similar percentage would translate in much larger changes in foregone fuel savings for vehicles with high fuel consumption. This would be an important point to show given that automakers should strategically set prices as a function of fuel savings in absolute values, not proportional changes. Therefore, although gaming appears to be constant in proportional term, it should lead to very heterogeneous impact on markups. In a differentiated market; the pass-through of gaming should also be heterogeneous given that different consumers are attracted to different car segments. I thus think that the most interesting dimension of heterogeneity to investigate is not across make or time, but across segments. In particular, I would like to see if gaming is heterogeneous across different bins of fuel consumption, car size, etc...

Thank you for this detailed comment. We have made several adjustments in response to it.

First, we have replicated our main analysis of the performance gap in levels, and we now report all of our specifications in both levels and percent changes. (Table A.6 shows the level changes.) Level specifications show the same qualitative trends, with a relatively flat performance gap that rapidly rises coincident with the advent of policy.

Second, we have added Figures A.4 and A.5. These Figures split the data in weight quartiles and show that the gap increases in similar amounts in all weight quartiles. We have chosen to show a split by weight as this characteristic is included in the Travelcard data, so that we can use the exact same sample for this as for the other results in Section 2.5. Your intuition is correct that the percentage gap is most pronounced in the lowest weight quartile, but the level gap is statistically indistinguishable between vehicles. We discuss this in Footnote 9 (p. 9). Presenting the gap across fuel consumption bins is harder to interpret as the composition of the bins changes considerably between release years.

Third, you raise interesting points about changes in markups and pass through. Our simulations in the final section of the paper show percentage changes in gaming. In the demand model the percentage change for a low fuel consumption vehicle will indeed result in larger monetary savings for consumers. Our model thus incorporates these diverse effects as the price equilibrium will incorporate these heterogeneous change in fuel costs. We agree that these effects are very interesting to study in detail, but consider this out of scope for this paper as the editor has asked us not to expand the simulation section.

**R3.2** The goal of this section is to quantify the extent that automakers distort ratings. The measure of gaming estimated in section 2.6 provides “sales-weighted” estimates. Would it be more appropriate to focus on “model-weighted” averages instead? For instance, the author could use a two-step estimator, where they estimate a similar model in the first stage, but with model-year fixed effects instead of having RELEASE and MODEL fixed effects. In the second stage, the model-year fixed effects could then be simply regressed on year dummies to obtain a model-weighted average for each year. One other benefit of using this approach would be to plot the distribution of model-year fixed effect for each year, which will provide another approach to show heterogeneity and how it evolves over time.

Thank you for this comment. We believe that section 2.7 already does something very similar to what you describe. In that section, we estimate the distribution of the vehicle-level performance gap using an empirical Bayes approach. This takes into account the quantity and quality of the data available for estimating each vehicle fixed effect. This seems to us the ideal method. Our results are presented in Figure 4, which shows a quite similar overall trend in the performance gap when compared to the results in the figure that plots coefficients from our main regressions (Figure 3). Details of the empirical Bayes correction are in Appendix A.3.

In the result table of that exercise, Table A.4, we have followed your suggestion and added results from a two stage procedure. We discuss this procedure in Footnote 17. The results of the two stage procedure are very consistent with the Bayes correction.

**R3.3** The cost of gaming, i.e., the function  $h()$ , is assumed to be a fixed cost. In my view, it should also scale with demand. I believe that one of the main component of this cost is the settlement

or penalty that should be paid if the firm is caught cheating. This amount is usually function of the quantity sold, especially in the car market. Take for instance the VW scandal, the size of the fine and the settlement was directly determined by the number of faulty VWs sold. A more general/realistic set-up would then consist to treat  $h(g)$  as a variable cost in Equation 3.

You raise a good point. We believe that in a model with marginal costs of gaming, the basic notion that we want to highlight—that cheating engenders cost savings—will also be present, but the mathematics may be somewhat different. We now acknowledge this more clearly in a new section that highlights our model assumptions and notes how some results might change under alternatives (subsection 3.3, p. 19). Regarding fixed costs, there we say:

In addition, we assume that the costs of gaming are strictly fixed costs, but it is possible that they scale with sales, either because of production costs or regulatory risk (i.e., because legal penalties scale with sales). In such a model, the same key forces that we identify will still be at work, but the formula would presumably be different.

**R3.4** I am struggling with some of the comparative statics related to proposition 2. One comparative statics result is related to the parameter  $\lambda$ , i.e., how tight the standard is. The parameter  $\lambda$  is, however, not explicitly expressed in the Proposition 2. I think the readers (the ones that are not the most theoretically incline like me) would need to be walk through the math there (in the main text or the appendix A3). Moreover, I think it will be more intuitive to conduct the comparatic statics with respect to the cost function  $c(x)$  or the level of the standard: the parameter  $\sigma$ , which are true primitives of this market, unlike the shadow cost of the regulation.

Thank you for drawing this to our attention. We have restated the result in terms of  $c'$ , a true primitive, rather than the shadow price  $\lambda$ . The idea is the same. We expanded our discussion in attempt to clarify (p. 22):

On the second point, for a given amount of gaming, price effects get more positive (favorable to buyers) under a tighter standard because gaming creates a larger cost savings. By assumption of the convexity of  $c$ , a tighter standard will push the marginal cost  $c'$  to be larger (more negative). Gaming, which effectively relaxes the standard, therefore has a larger cost savings and subsequently a larger buyer price effect. For a sufficiently tight standard, the price effect will always be positive. This is true even if buyers are fully unaware of gaming.

**R3.5** The demand-side of the model assumes that the consumer pays full attention to  $x$ , and may be fooled only by the amount of gaming. What about the case for which consumers are also (partly) inattentive to the level of  $x$ ? I understand that the authors rule out this case as they have both

showed in separate work that consumers seem to fully internalize fuel costs in the car market. However, I consider that this is still an interesting case to consider. Inattention to fuel costs, or the beliefs that consumers are inattentive, is an important feature of this market. On the firm side, this may rationalize why firms are not shy about gaming—if they believe that consumers do not value this attribute, they might be more likely to cheat. On the policy side, it is one of the main reason why policymakers still favor a standard- type regulation over a tax.

Thank you for raising this issue. We believe that additions to our simulation section (rather than the model) address the spirit of your comments. Specifically, our simulation results table Table 2 includes a new second panel that shows the welfare effects of gaming when firms comply not by mix shifting but by adding technology to vehicles. To model the cost of technology, we use a cost curve from the EU regulation. This cost curve has a feature common among engineering cost curve estimates of energy efficiency: it indicates that modest increases in fuel economy cost less than the net present value of fuel savings benefits. This relates to what is known as the energy efficiency gap. We present two sets of results, consistent with alternative interpretations of the energy efficiency gap.

In one scenario (Column VI of Table 2) we take that engineering cost curve at face value. In this case, the regulation raises private buyer welfare, and as a result gaming is bad for buyers. This interpretation essentially introduces an additional market failure, in which the market delivers too little fuel-saving technology for some reason. The change in results are intuitive: if consumers undervalue fuel economy, then the regulation fixes an externality that benefits buyers, and gaming precludes the welfare gains from correcting the externality. We discuss this interpretation in section 4.2 (p. 31):

In Column VI we show the case where compliance costs are lower than consumer willingness to pay for fuel consumption. In this case the policy is understood to solve a market failure related to the provision of energy efficiency (e.g., consumer undervaluation), and the policy raises private buyer surplus. As a result, gaming hurts consumers in this scenario, not only because of the choice distortion, but also because buyers miss out on the positive price effects of the policy (the price effect of gaming is thus negative).

In our new subsection discussing theoretical assumptions, we note this and foreshadow the new simulations (p. 20):

Our model assumes that consumers properly value fuel economy  $x$ , even if they are unaware of gaming  $\alpha$ . A literature on the energy efficiency gap debates whether consumers might undervalue fuel economy. If they do, then the regulation fixes an “externality” and improves *private* buyer surplus. In this scenario, gaming will not have a beneficial price effect. We demonstrate this in our simulation results below.

**R3.6** A minor, but related point: in Section 4.1, the authors discuss consumer awareness. There, the discussion about awareness to gaming versus the overall level of fuel costs are somewhat conflated, which I think may confused readers that are not well-versed in this literature.

Thank you for pointing this out. In the section you reference, we have rewritten the relevant paragraph to make the distinction sharper (p. 28):

We are not aware of any direct evidence relating to buyer knowledge of test gaming. There is, however, a substantial literature on fuel economy valuation that demonstrates that buyers respond to fuel price changes in ways that are consistent with them being attentive to test ratings and to the value of fuel savings.<sup>47</sup> That literature suggests that buyers pay attention to government test ratings, but it does not necessarily imply that buyers are aware of gaming or are able to correctly deflate government ratings in the presence of rapid declines in test accuracy.

**R3.7** Section 3's main focus is on the impact on consumer surplus with a little extension in Section 3.6 to discuss total welfare. I found this sub-section somewhat informal and incomplete. I see a good opportunity to show formally under which circumstance gaming may be socially beneficial or not.

Thank you for this comment. When we began this project, we anticipated writing an analysis that was more focused on overall welfare, but it became apparent to us that our unique contribution was really in regards to points about buyer surplus. In terms of total welfare, the price effects that are our focus are just transfers that wash out. Moreover, it is clear that a policy can be beneficial or not, depending on the compliance costs and the size of the externality. That is, gaming is either good or bad for welfare, depending on whether it is relaxing an overly stringent constraint, or relaxing one that is already too loose. We don't need a formal model to show this and do not wish to take a stand on the optimality of the stringency of the EU standard. What our model does is articulate the unique channels through which gaming impacts buyer welfare, and we believe that our core contribution is the identification of these channels, the description of how they hinge on the existence of policy that is our value added, and the illustration that they map directly into empirical objects estimable in discrete choice settings.

**R3.8** First, the simulation model considers a different mechanism leading to gaming that the one that has been considered in the theory section and demonstrated in the empirical section. To be

---

<sup>47</sup>A recent literature (see [Busse, Knittel, and Zettelmeyer 2013](#), [Allcott and Wozny 2014](#), [Sallee, West, and Fan 2016](#), [Grigolon, Reynaert, and Verboven Forthcoming](#)) has shown that buyers do take fuel consumption ratings into account when purchasing a vehicle. In a field experiment [Allcott and Knittel \(2017\)](#) show that treating buyers with detailed information about fuel consumption has no impact on their choice.

consistent, it would be better if the simulation model were to endogenize fuel economy and not only pricing. A key parameter on the supply-side is how the marginal cost of each product varies with fuel economy and how this additional cost is passed to consumers. In light of the recent discussion of Leard, Linn, and Springel (2019), it may not be trivial to exclude the impact of the regulation on product attributes, and to solely focus on strategic pricing. My prior is that solely focusing on sales-mixing, may not provide an accurate quantification of the effect of gaming in this market. In a multi-product oligopoly market, short-term pricing decisions versus medium-term pricing decisions where technology (i.e., fuel economy) and prices are strategically set should lead to different equilibrium outcomes. Crucially, markups could be very different in both scenarios.

Thank you for this comment. In the revised version, we now also simulate what happens to buyer surplus when firms adjust fuel consumption, as we discussed in our response to your comment above regarding undervaluation. In these simulations, improving fuel consumption is costly to the automaker, but in our base scenario the engineering cost estimates are below private benefits. In this scenario, the regulation is correcting an externality. Gaming subsequently harms consumers because it prevents this externality from being corrected. We show those results in column VI of Table 2. We discuss the setup and assumptions in section 4.1 and the results in section 4.2.

Note that we also pursue two other scenarios where we shift up the engineering cost curve so as to equate marginal costs and buyer benefits prior to the advent of policy. In this case, compliance via technology is privately costly (in effect we have eliminated the externality), and the results are qualitatively the same as our results based on mix shifting as the compliance strategy: price effects from gaming benefit consumers and outweigh negative choice distortions.

**R3.10** A suggestion: The cost of changing fuel economy could be estimated, or calibrated. An amount of gaming could then be taken as given, say 30%, and fuel economy and prices could then be endogenized.

Thank you for this suggestion. Given that our main simulation is a calibration exercise to uncover the importance of the several parameters that determine buyer surplus from gaming we have chosen to stick with a simulated cost curve in this project. See Reynaert (2019) where one of the current co-authors estimates a model that considers strategic price setting, adoption of technology and gaming and compares estimated technology costs with engineering data on costs.

**R3.11** The demand system is somewhat of a black-box. This makes it hard to assess how robust are the results. I like the sensitivity test where the size of the outside option is changed. This makes me want to see more. Equally important to quantify the price and choice distortion effects is the degree of substitution between products. Observed and unobserved preferences heterogeneity are thus important, but it is currently difficult to see how these features play out. The fact that there



are little choice distortions when the outside option is set at 10% makes me wonder how rich are the substitutions patterns induced by the demand model. This is a little surprising and worrisome that most of the choice distortions occur via a market expansion effect. Is it realistic to have so much consumers that decide to own car because of lower expected fuel costs, but to have little (in relative terms) consumers switching between different types of car? My prior is that the latter effect should be more prevalent.

In response to comments of the editor and other referees, we have kept discussion of the demand estimation in the paper as concise as possible. To give more detail about the estimation results and resulting substitution patterns we have added Tables A.8 and A.9 to the Appendix. These tables summarize the own price elasticities by fuel consumption quartile. To summarize cross-price elasticities we have added the diversion ratios between the fuel consumption quartiles and the outside good. We refer to these tables in section 4.1.

In response to your concern about substitution to the outside good, Column II of Table A.8 and A.9 give the diversion ratios to the outside good by fuel consumption quartile. This shows that substitution to the outside good is at most capturing 40% of the diversion. In the case where the outside good is small we find diversion to the outside good smaller than 3% for all fuel consumption quartiles. Switching between vehicles is thus more important quantitatively than market expansion in all of the results that we show.

The reason that the value of the choice distortion shrinks between Column II and Column IV of Table 2 (we guess this raises your concerns about substitution patterns) is related to changes in the value of consumer welfare rather than changes in substitution patterns. As market shares change (when changing the value of the outside good), the value of the consumer utility also changes as the model is a direct mapping from shares to indirect utility. This explains why the numerical value of the choice and price distortions changes.

**R3.12** Related to the above point, what is the role of income in the demand model? I expect that the elasticity with respect to the outside should vary with income and this could have important implications. For instance, lower income households should be more likely to purchase smaller, cheaper, more fuel-efficient vehicles, and more elastic with respect to the outside option. On one hand, the market expansion effect should be pronounced for these segments. On the other hand, those are also the segments that are the less likely to be affected by gaming in the first place. (A 30% gaming on fuel efficient vehicles might translate in small changes in fuel costs.)

Thank you for this interesting comment. We have clarified that we have aggregated sales per country in Footnote 28 and that our demand model does not incorporate data on household level incomes. We fully agree that studying differences in incidence of gaming across income groups is very interesting. However, given the lack of more granular data in the current study we leave this for future research.

Notice that the market expansion effect (see Table A.8) is indeed most pronounced in the lowest fuel consumption quartile, in line with your intuition.

**R3.13** Does parametric uncertainty in the demand model matter? Table 2 should have standard errors.

Thank you for point this out. We have added a 95% confidence intervals for all the simulated numbers. To do this, we draw from the estimated variance-co-variance matrix of the parameters assuming joint normality. For each of the N draws we compute the difference between the counterfactual outcome and the outcome at the point estimate of the parameter. In this way we obtain an empirical bootstrap interval for the counterfactual outcomes.

**R3.14** I believe that this paper is not the first to provide evidence of gaming in the context of fuel economy ratings. Isn't Saltee and Slemrod (2012) that provide the first empirical evidence that automakers game fuel economy estimates.

Thank you for pointing out the relationship between that paper and our current manuscript. We now discuss this in footnote 19 (p. 17): "Similarly, the first paper to show manipulation of fuel economy ratings in response to incentives uses tax notches to identify manipulation (Saltee and Slemrod 2012)."

**R3.15** The economic literature seems to refer to the Goodhart's Law when "self-fulfilling expectations" make a policy ineffective. It is not quite the case here, and I find the references, especially to Lucas, not quite on point.

Thank you for causing us to think about this again. We have eliminated the reference to the Lucas critique, which is more narrowly about self-fulfilling expectations and is arguably not that closely related to our paper. Goodhart's Law, however, is the notion that when you create incentives around a particular measure, that measure becomes a joint measure of whatever it was measuring before and strategic manipulation of that measure. We think this is exactly the phenomenon that we describe around fuel consumption ratings, so we have elected to keep that reference. We clarify this when summarizing our reduced form results (p. 17): "We see the decrease in the correlation between official and on-road ratings at a time when stakes on official ratings increase as a manifestation of Goodhart's law—once policy incentives were loaded onto the fuel economy test results, those test results became a less reliable measure of actual fuel economy performance."