DISCUSSION PAPER SERIES

DP13756

ABATEMENT STRATEGIES AND THE COST OF ENVIRONMENTAL REGULATION: EMISSION STANDARDS ON THE EUROPEAN CAR MARKET

Mathias Reynaert

INDUSTRIAL ORGANIZATION AND PUBLIC ECONOMICS

ISSN 0265-8003

ABATEMENT STRATEGIES AND THE COST OF ENVIRONMENTAL REGULATION: EMISSION STANDARDS ON THE EUROPEAN CAR MARKET

Mathias Reynaert

Discussion Paper DP13756
Published 24 May 2019
Submitted 20 May 2019

ABATEMENT STRATEGIES AND THE COST OF ENVIRONMENTAL REGULATION: EMISSION STANDARDS ON THE EUROPEAN CAR MARKET

Abstract

Emission standards are a major policy tool to reduce greenhouse gas emissions from transportation. The welfare effects from this type of regulation depend on how firms choose to abate emissions, i.e., by sales-mixing (changing prices), by downsizing (releasing smaller cars), by technology adoption or by gaming emission tests. Using panel data covering 1998-2011, I find that the introduction of a EU-wide emission standard coincides with a 14% drop in emission ratings. I find that this drop is fully explained by technology adoption and gaming and not by sales mixing or downsizing. I estimate a structural model to find that the regulation missed its emission target and was not welfare improving. Abatement with sales mixing would have reduced emissions, but at high costs. The political environment in the EU shaped the design and weak enforcement of the regulation and explains the choices for abatement by technology adoption and gaming.

JEL Classification: Q5, L5

Keywords: Environmental Regulation, compliance, Carbon Emissions, automobiles, fuel economy

Mathias Reynaert - mathias.reynaert@gmail.com
Toulouse School of Economics and CEPR

Abatement Strategies and the Cost of Environmental Regulation: Emission Standards on the European Car Market.

Mathias Reynaert *
Toulouse School of Economics, University of Toulouse Capitole

May 2019

Abstract

Emission standards are a major policy tool to reduce greenhouse gas emissions from transportation. The welfare effects from this type of regulation depend on how firms choose to abate emissions, i.e., by sales-mixing (changing prices), by downsizing (releasing smaller cars), by technology adoption or by gaming emission tests. Using panel data covering 1998-2011, I find that the introduction of a EU-wide emission standard coincides with a 14% drop in emission ratings. I find that this drop is fully explained by technology adoption and gaming and not by sales mixing or downsizing. I estimate a structural model to find that the regulation missed its emission target and was not welfare improving. Abatement with sales mixing would have reduced emissions, but at high costs. The political environment in the EU shaped the design and weak enforcement of the regulation and explains the choices for abatement by technology adoption and gaming.

Keywords: environmental regulation, compliance, carbon emissions, automobiles, fuel economy
JEL: Q5, L5

*Toulouse School of Economics 21 allée de Brienne, 31015 Toulouse Cedex 6, France, e-mail: mathias.reynaert@tse-fr.eu. I am grateful to Frank Verboven, James Sallee, Bruno de Borger, Johannes Van Biesebroeck and Jan Bouckaert for their support and guidance. The paper benefited from comments and discussions during presentations at Harvard/MIT, Wharton School at University of Pennsylvania, London School of Economics, Yale University, Jornadas de Economia International, University of Chicago, Telecom ParisTech, EARIE 2015, CEPR Applied IO 2015, Mc Master University, University of Leuven, University of Antwerp, UCL-CORE, Toulouse School of Economics, Stockholm School of Economics, Purdue University, Indiana University, Boston College, University of Mannheim and Universitat Pompeu Fabra.

1 Introduction

Transportation accounts for 20% of global greenhouse gas emissions and policy makers are taking up the challenge to reduce the use of polluting petroleum liquids. One of the major policy tools used to achieve this goal are emission standards. This type of regulation sets mandatory limits on average emission rates (or fuel economy) across the fleet. Emission standards are used across the globe in China, the European Union, Japan, the United States of America and other countries. These policies are simple to prescribe and do not explicitly tax consumers or producers. This paper studies the introduction of the emission standard in the EU and tries to estimate the welfare effects of the EU policy.

The EU regulation aimed to reduce CO_2 emissions from passenger cars by 18%. The policy was announced in 2007 and became fully binding from 2015, after a phase-in period that started in 2012. The regulation targets CO_2 emissions which is equivalent to targeting fuel consumption or fuel efficiency.[1] The EU standard is interesting to study for three reasons. First, it is a very demanding standard with a target of 130 g CO_2/km. For comparison, the US standard required only 36 miles per gallon (mpg) in 2016, while the EU standard requires approximately 42 mpg. Second, before the standard, the EU had no regulation on CO_2 emissions. The introduction of the standard thus allows us to study how the market equilibrium changes with the introduction of a stringent emission standard. Third, the EU standard is attribute-based; the target for each firm depends on the average vehicle weight. This means that firms producing heavier vehicles face a less stringent target. Several other countries introduced an attribute-based standard.[2] Understanding the effects of the EU standard is thus helpful to guide the design of this type of regulation in the future and in other markets across the world.

The task of evaluating the welfare impact of emission standards is not an easy one. Firms can choose between different abatement strategies to comply with a standard, and these strategies will have different effects on the market equilibrium. A first strategy is sales-mixing, i.e., shifting the relative prices of vehicles with different CO_2 emissions. A second strategy is downsizing, i.e., releasing smaller but more fuel efficient vehicles. A third strategy is technology adoption, i.e., improving the fuel efficiency of the vehicle fleet. A fourth strategy is gaming, i.e., improving the fuel efficiency as measured in official ratings without improving the actual fuel efficiency on the road. These strategies will change the prices, product attributes, product sets and market outcomes in different ways. Firms will choose the abatement strategy that has the lowest cost, while taking into account the strategy chosen by competing firms. Additionally, the design of the policy will matter for the costs of different abatement strategies. To evaluate the EU standard, this paper

[1] CO_2 cannot be filtered during the combustion process. Fuel consumption translates proportionally into grams of CO_2 per km, with a different CO_2 content per liter for diesel and gasoline. Fuel consumption (liters per kilometer) and CO_2 emissions per kilometer are the inverse of fuel economy (miles per gallon).

[2] The International Council on Clean Transportation (2014) compares the different regulations between countries. The EU has the goal of decreasing emissions to 95 g/km by 2021; the US has communicated a goal of 103 g/km by 2025; Japan has a goal 105 g/km by 2020; and China has a goal of 117 g/km by 2020. The US and Japan have also introduced attribute-basing in their regulations.

presents a model that accounts for firms' abatement strategies and the design of the policy.

In a first step I explain the trend in sales weighted CO_2 emissions between 1998 and 2011 in the EU market. Knittel (2011) shows that vehicles evolve over time in terms of the amount of attributes that are offered relative to fuel economy. He finds that there is a constant improvement in the amount of attributes and the fuel economy offered and interprets this as technological progress. I follow his approach and estimate technological improvements in the trade-off that firms face between emissions and other vehicle characteristics. I find that emissions were reduced by 14% after the regulatory announcement. This reduction is fully explained by increases in the fuel efficiency improvements of vehicles, while the attributes and the product set offered by firms do not change. This shows that the market is not shifting towards smaller vehicles and thus rules out downsizing as an explanation for the drop in emissions. The results show that the technological progress is twice as fast after the regulatory announcement, in line with findings of Klier and Linn (2016). Firms thus respond to emission standards by increasing the speed of technology adoption, at least when we look at the official emission ratings. In Reynaert and Sallee (2019) we studied the fuel efficiency of thousands of vehicles on the road to find that there is significant gaming in the market. Emission ratings improve on paper, but these improvements do not translate to the road. Using these estimates of on-road fuel efficiency, I find that only 30% of the increased technology adoption is measurable on the road, so that 70% can be attributed to gaming.

The observed decrease in the emission ratings is so strong that almost all of the firms reach the emission target before it becomes partly binding in 2012. However, looking at market outcomes before and after the policy is not sufficient to find the welfare effects of the policy. The effect of the regulation cannot be separated from changes in local regulations and taxes. Many EU member states began changing vehicle taxation after 2007, clearly contributing to the downward trend in emissions.[3] There is also an 8 year gap between the policy announcement and implementation so that the demand, costs and market fundamentals potentially change. Therefore, I resort to a structural model to simulate the impact of the emission standard. The model allows me to single out the impact of the policy, and it also allows me to study how the impact changes when we change the design of the policy. In the counterfactual, I assume that the product sets of the firms and the product attributes are fixed; this rules out downsizing, which is in line with the decomposition of emission trends. I then model heterogeneous consumers making discrete choices between vehicles. The firms have the following three strategic choices: price setting (sales mixing), technology adoption and gaming. The model allows me to clearly indicate how the equilibrium changes with each of the strategic choices. To simulate the model I need estimates of the following primitives: preferences and price elasticities, marginal costs, changes in costs from technology adoption and changes in costs from gaming.

For the estimation of these primitives I rely on the rich panel data from before the policy announcement. This has the advantage that firms' decisions are not affected by the policy. I

[3] See for example the French bonus malus discussed in Durrmeyer (2018) and the Swedish rebate discussed in Huse and Lucinda (2014)

estimate consumer tastes and price elasticities following the framework of Berry, Levinsohn, and Pakes (1995). I allow for heterogeneous tastes of consumers for several characteristics, including fuel costs. I extend the model to allow for multiple endogenous characteristics including prices, fuel economy, horsepower and weight. Though I do not model the firms' choices of these variables explicitly in the counterfactual, the demand estimation avoids biases in the taste parameters that would show up when firms choose characteristics with more information on unobserved quality than the econometrician. I instrument these characteristics with a combination of sums of characteristics, globalization and production platform instruments. The globalization instruments exploit changes in the international production network of firms over time. When vehicles become popular in different international markets, this impacts the production location choices as well as the engine design choices. The validity of the instrument relies on the assumption that differences in the global exposure are uncorrelated with unobserved quality in the EU. Given the estimated demand, I invert the first order conditions of the profit function with respect to prices to find the marginal costs of vehicles. Because I use prepolicy data, this is straightforward. I also invert the first order conditions with respect to fuel consumption to find the marginal benefits and marginal costs firms face in deciding how much fuel consumption (and consequently emissions) to offer. This will be a crucial input to model the technology adoption; the regulation will increase the marginal benefit of lowering fuel consumption, as vehicles with lower emissions help the firm attaining the regulatory target. However, in doing so, firms will face increases in the marginal cost of offering lower fuel consumption. To predict how cost changes with decreases in fuel consumption, I rely on two approaches. First, I use the observed relation in the data between fuel consumption and the estimated marginal cost. Second, I use the engineering reports submitted with the policy proposal that give detailed predictions of marginal cost increases. Finally, I need an estimate for the costs of gaming. The costs of gaming consist of the expected litigation and the cost of emission test falsification. These costs are very difficult to estimate and I therefore rely on the level of gaming found in the emission decomposition, rather than modeling a cost function to solve for optimal gaming. The estimated model with multiple endogenous characteristics and multiple strategic decisions on continuous variables is similar to the recent contributions of Fan (2013) and Crawford, Shcherbakov, and Shum (2019).

Given the estimated primitives I introduce the EU policy in the model and simulate how it impacts the market equilibrium. The simulations reveal that technology adoption in combination with gaming is indeed the equilibrium abatement strategy of firms. The increases in costs from technology adoption beyond the willingness to pay for fuel consumption imply that the regulation decreases the consumer surplus and profits. Because of the gaming, the reductions in CO_2 emission are a mere 5% instead of the 18% target. The sum of these emission savings and consumer and profit losses is clearly negative so that the regulation reduces welfare. However, when I consider two additional welfare effects,i.e., reductions in other externalities and corrections in consumer undervaluation, I find the emission standard to have a very small positive impact. The simulation also reveals that when I restrict the firms' abatement strategy to sales mixing, the standard would have had starkly different welfare effects, where the total sales and emissions then strongly decrease at a very high costs for consumers. Why did firms choose for the combination of technology adoption and gaming in response to the EU standard?

I find that the attribute based design of the regulation, which lets the emission target vary with vehicle weight, makes sales mix abatement much more costly for firms. Most firms have less vehicles beneath the attribute based target than below the flat target. Firms will have to distort prices more to lower their sales-weighted emissions due to the attribute basing. If the regulation had a flat design that is not attribute-based, firms would have opted for a mixture of technology adoption and sales mixing. This would have led to CO_2 emission reductions of 11%, which is much closer to the 18% target. Why, then, was the attribute design introduced? The simulations show that the positions of the governments are in line with their firms. The attribute basing redistributes the incidence of the regulation between French, Italian and German producers. This is consistent with newspaper articles at the time citing representatives of different car producing countries. The French and Italian governments were in favor of the flat regulation while Germany lobbied heavily for a steep attribute design. Additionally, the gaming is a product of the political environment. A recent evaluation by the European Parliament (Gieseke and Gerbandy (2017)) has placed responsibility for enforcement failures with the car producing member states. The countries failed to detect and respond to gaming timely, so that the policy missed its target. Together, the political deal regarding attribute basing and the enforcement failures, show the importance of the political environment in which emission standards are introduced.

The paper makes several contributions. First, I show that emission standards can induce technology adoption and gaming by firms. The equilibrium effects of the abatement strategies are not studied in detail in the previous literature. The literature studying the CAFE standard in the US treats changes in the level of technology as a possible longer run effect of emission standards and has not focused on the welfare effects from gaming. Second, by estimating a structural model of demand and supply, I show that the incidence and welfare effects of the regulation vary drastically between different abatement strategies. Third, the model allows me to study how the design of the regulation affects the outcomes. The attribute-basing increased the pressure on firms to adopt technology, and the weak enforcement allowed for gaming so that the policy missed its emission reduction target. Together, these three contributions show that it is crucial to take the supply responses carefully into account when evaluating and designing emission standards. Finally, this is the first paper providing a detailed study of the EU regulation and its impacts on market outcomes.

The framework in this paper builds on the existing work of Knittel (2011), Jacobsen (2013) and Reynaert and Sallee (2019). The emission decomposition follows the estimation in Knittel (2011), but I find that the speed of technology adoption changes when the policy is announced. Jacobsen (2013) incorporates heterogeneous responses from both consumers and producers in a structural model. In an extension Jacobsen (2013) also considers technology adoption using the framework of Austin and Dinan (2005). My analysis contributes by fully considering technology adoption in the economic model. This shows the importance of the starting slope of the technology cost curve and how technology adoption lowers shadow costs endogenously. While Jacobsen (2013) finds that technology adoption limits the welfare losses of standards, I show that technology adoption could lead to welfare losses or gains depending on the cost effectiveness of the technology relative to consumers willingness to pay for fuel consumption.[4] In Reynaert and Sallee (2019) we estimate the amount of gaming in the EU market. I use the on-road fuel efficiency estimates of our work to see how much of the faster technology take up translates to the road. I also use the framework we built to estimate the effects of gaming on consumer surplus. The effects of gaming crucially depend on the consumer awareness and cost avoidance of gaming, and we calibrate several scenarios varying these parameters. Here, the interaction between choice and gaming is embedded in the full equilibrium model interacting with other abatement strategies, and I focus on the overall welfare rather than the narrow consumer welfare.

This paper adds to a body of literature that studies policies targeting vehicle emissions, see Anderson and Sallee (2016) for an overview. Holland, Hughes, and Knittel (2009) show that none of the welfare effects of emission standards are theoretically determined. A regulation might decrease or increase emissions from new vehicles, depending on the distribution of the price elasticities of products below and above the emission target. This shows the importance of estimating a model that allows for rich substitution patterns. However, the welfare effects of standards become even more uncertain when we consider the variety of different abatement strategies. The empirical literature on emission standards has focused on the US CAFE standard. Goldberg (1998) was the first to consider the effect of standards on price setting and the composition of the vehicle fleet. Anderson and Sallee (2011) use a loophole in the regulation to show that the standard is hardly binding in recent years. Jacobsen (2013) finds that the US CAFE standard imposes a large shadow cost on domestic US firms. Both Klier and Linn (2012) and Whitefoot, Fowlie, and Skerlos (2017) extend the analysis by considering the endogenous product characteristics in the model, allowing firms to change the characteristics of vehicles. Both of these papers assume that the level of technology is fixed and consider reoptimization of product offerings on a fixed trade-off function between emissions and attributes. I allow the trade-off relation to change but impose that improvements must reduce fuel consumption, in line with the empirical evidence in Europe. The economic effects of attribute-based regulations are discussed in Ito and Sallee (2018), who focus on distortions in the market for the attribute. Whitefoot and Skerlos (2012) compute the vehicle weight distortions for the CAFE standard with footprint basing. The analysis here is complementary as I study different effects of attribute-basing, i.e., the cost of different abatement strategies and the political economy behind the attribute basing.

This paper focuses on the direct welfare effects of the EU emission standard. Recent work has considered additional margins of the policy. Jacobsen and van Benthem (2015) study the effect of

[4] Other notable differences are the demand model where I estimate price elasticities on the engine level (more than 400 products per market) rather than the broad category level (including 12 products) and I allow for endogenous characteristics. The model of Jacobsen (2013) is richer in other dimensions as it incorporates effects on the second hand market as well as changes in the vehicle miles traveled as modeled in Bento, Goulder, Jacobsen, and von Haefen (2009)

emission standards on vehicle scrap rates. Durrmeyer and Samano (2018) compare a standard with a rebate policy that explicitly subsidizes and taxes emissions above and below the target. Bento, Gillingham, and Roth (2017) focus on the effect of the fuel standard on the dispersion in vehicle weight and its effect on accidents.[5]

The paper is structured as follows. Section 2 describes the policy and the available data. Section 3 decomposes the changes in emissions in the EU automobile market between 2007 and 2011. Section 4 presents the emission standard in a model of supply and demand and discusses the possible effects of the different abatement strategies. Section 5 presents the estimation results. Section 6 presents the results of policy simulations, and Section 7 concludes.

2 The EU emission standard and data

The EU emission standard The European regulation on emission standards for new passenger cars, Regulation (EC) No. 443/2009, sets a mandatory fleet average of κ =130 grams CO_2 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:

$$\frac{\sum_{j\epsilon fleet} q_j(e_j - f(w_j))}{\sum_{j\epsilon fleet} q_j} \quad \leq \quad 130. \tag{1}$$

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 $a = 0.046$.[6] 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.[7] Figure 1 plots the target and the distance from the target for each producer in 2007 and 2011. 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. Deters (2010) gives an overview of the full legislative process and the political background. 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. I will not model the phase-in period. Every firm succeeded in reaching the full

[5]These margins play a role in the current policy debate in the US as discussed in Bento, Gillingham, Jacobsen, Knittel, Leard, Linn, McConnell, Rapson, Sallee, van Benthem, and Whitefoot (2018)

[6]The average SUV in the data weighs 1650 kg, and the average compact car weighs 1250 kg. The SUV's emissions will be scaled down while the compact car's emissions will be scaled up.

[7]Manufacturers can obtain lower average emissions by 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 are ignored in the analysis since they count for a very small share of the total market.

Figure 1: Compliance of Firms in 2007 and 2011

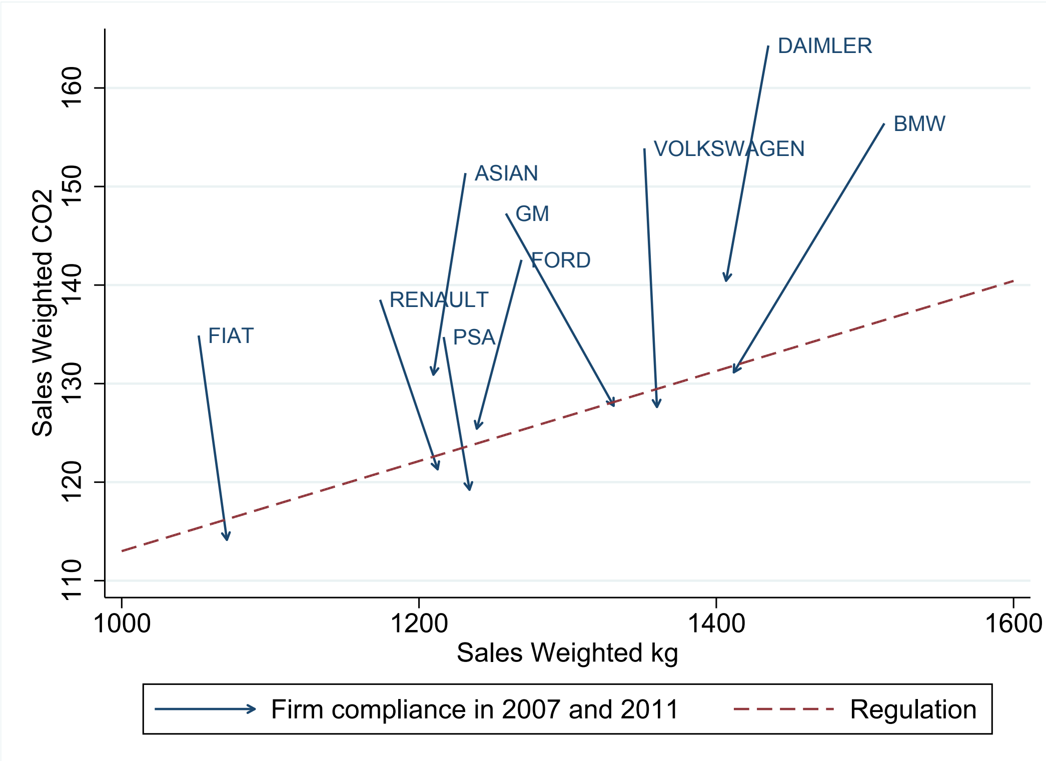

The starting point of each arrow gives the sales weighted CO_2 and mass for each producer in 2007 as observed from the data. The end of each arrow gives the same point in 2011. The dashed diagonal line is the regulation, which is fully binding in 2015.

target by 2014.

The specifics of the regulation were heavily debated during the drafting of the law. Several newspaper reports discuss lobbying efforts by EU member states, firms and environmental groups.[8] France and Italy were strongly in favor of a flat standard, while Germany wanted an upward sloping target function in either weight or footprint (the rectangular area in between the wheels of the vehicle). The German firms BMW, Daimler and Volkswagen, on average, make heavier vehicles than Fiat (Italian), Renault and PSA (French). The production of each of these firms mostly takes place within the boundaries of the home country, and the car sector is an important source for employment.

It is instructive to compare the EU policy with the US CAFE standard since this policy has been the subject of several studies. The CAFE standard came into place in 1978 and after a gradual phase-in has been constant at 27.5 mpg since 1990 (this corresponds to 198 g CO_2/km). From 2009 onward, the CAFE standards are tightened towards 36 mpg in 2016 (this corresponds to 152 g CO_2/km). Contrary to the EU standard, light trucks (SUV's) face a different less demanding

[8]See, for example, "EU unveils tough emissions curbs for cars" - Financial Times, February 7 2007 and "France battles Germany over car emissions" - Financial Times, November 15 2007.

target than passenger cars. Additionally, firms are allowed to trade excess emissions over time and with other firms.[9] From 2012 onward, the CAFE standard also has an attribute-based part, i.e., the target varies with the footprint.

Data The main data set is obtained from a market research firm (JATO dynamics) and contains sales and product characteristics for each passenger car sold during 1998-2011 in the following seven European countries: Belgium, France, Germany, Italy, Great Britain, The Netherlands and Spain.[10] The characteristics and sales are given for several engine variants of a car model at the country level with a yearly frequency. The country will be the geographical market. A model is defined as a brand/model name/body type combination (e.g., Volkswagen Golf Hatchback).[11] The engine variants differ in fuel type (gasoline or diesel) and engine performance. Accounting for fuel type is important in the EU market, as diesel variants have a considerable market share (56% in 2011) and the emissions of diesel variants are lower; a diesel engine emits approximately 20% less $C0_2$ per kilometer.[12]

Sales are defined as new vehicle registrations in each of the countries. The prices are the suggested retail prices (including registration taxes and VAT, as obtained from the European Automobile Association). The product characteristics give information on the vehicle size (footprint defined as length by width, weight and height) and engine performance (horsepower and displacement). The data also contains information on fuel consumption (liters per 100 km and CO_2 emissions per km). These numbers are the official consumption ratings obtained from the New European Driving Cycle (NEDC), a standardized driving cycle to assess the emission levels of car engines. The cycle simulates both urban and extraurban driving patterns and excludes the use of auxiliary features, such as air conditioning. Real world emissions thus differentiate from these test values. In Reynaert and Sallee (2019) we develop a measure of on-road emissions and we show that reductions in the official CO_2 ratings do not fully translate into on-road savings. Car makers are able to calibrate engines with defeat devices and specific software so that they perform much better on the test. We coin the improvements that do not translate to the road gaming. I will use the measure of on-road emissions in the analysis in order to disentangle technology adoption from gaming.[13] This information is available for a much more limited sample of vehicles. Finally, the data on sales and emissions are supplemented with production data from PriceWaterhouseCoopers (PWC) that contain the country and plant of production for each model. I match this with a producer price index and a unit labor cost measure obtained from the OECD. Finally, the data on fuel prices (from

[9] Contrary to the CAFE standards in the US there is also no banking system for excess emissions over time. The penalties in the EU are lower for low excess emissions but increase to higher levels than the penalties for breaking the US CAFE standards.

[10] These markets represent approximately 90% of the total EU market.

[11] The body types are as follows: hatchback, sedan, wagon, coupe, convertible, mini MPV and SUV.

[12] The combustion process and different energy content of the fuel make diesel engines more efficient per kilometer. Diesel cars emit less CO_2 per kilometer, but more other pollutants such as NOX.

[13] In Reynaert and Sallee (2019) we construct on-road emissions from a panel of 12,000 drivers visiting fuel station 22 million times. We observe odometer readings and fuel purchases and use those to construct fuel consumption on the road.

DataStream), GDP/capita and number of households in each country (from Eurostat) are used to construct the fuel costs for consumers, real prices and the number of potential buyers in each year.

To reduce the size of the data and the complexity of the analysis, I leave out firms, brands and models with very low sales. The analysis will focus on the largest producers and their bestselling brands on the EU market. The included firms are BMW, Daimler, Fiat, Ford, General Motors, PSA, Renault and Volkswagen. I treat the largest Asian car makers as one decision maker. This includes the firms Honda, Hyundai, Mitsubishi, Nissan, Suzuki and Toyota.[14] The list of excluded brands and a detailed description of the model selection and data manipulations can be found in the appendix. In total I keep 40,239 market/year/model/engine variants in 98 year/countries, or approximately 400 model engine variants per market. The final data contains 80% of total reported sales in the sample.

Throughout the paper, the full dataset is partitioned over time and markets in several ways. In Section 3, I collapse the data towards a unique model engine variant in each year and leave out the variation over markets. These data are used to make statements on the evolution of the supply of engine characteristics in response to the policy and contains 14,444 unique observations. To estimate the structural model I will rely only on data prior to the policy announcement and use the years 1998-2007. This exploits 30,000 year/market/model-engine observations. The data from the year 2007 will be used as the benchmark for the simulations in Section 6.

Summary Statistics Figure 1 plots each producer's distance from the emission standard in 2007 and 2011. Each firm needs to move below the dotted line, which presents the attribute-based emission standard. In 2007 each firm is far above the target and has the following three options to reach the standard: reduce emissions, increase vehicle weight or combine both. The Asian firms, BMW, Daimler and Ford decrease weight and reduce emissions. Volkswagen reduces emissions while keeping weight constant. Fiat, GM, PSA and Renault all increase the average weight slightly while decreasing emissions strongly. A strong downward trend in emissions towards the standard is observed for all firms. The decrease in emissions is so strong that most of the firms comply with the emission standard four years before it is fully binding.[15] Table 1 shows the change in the sales weighted vehicle characteristics between 2007 and 2011. CO_2 emissions decrease by 14% while there is moderate growth for the other characteristics.

Figure 2 plots the sales weighted characteristics over time from 1998 to 2011 for both the EU (Panel a) and the US (Panel b). Each characteristic is indexed in 1998. The most remarkable trend in the EU is the evolution of sales weighted CO_2 emissions. The level of emissions is constant up until 2002, slightly declines approximately 6% until 2007, and then plunges by 14% in the last

[14]I combine the fleet of the Asian car makers because most of these firms do not have a very broad product set. This makes finding a price equilibrium with sales-mixing impossible in my algorithms. Alternatively, I could choose to only keep Toyota and Nissan, by far the largest Asian firms in the EU, but combining all Asian firms allows me to include more products. Note that the emission standard is sales based, it does not matter where the vehicle is produced. This means that imports sold in the EU are counted and exports are ignored.

[15]This shows that the emission standard is probably not the only mechanism driving down the sales weighted emissions. Below I comment on complementary explanations.

Table 1: Sales weighted vehicle characteristics in 2007 and 2011

Characteristics	2007	2011	% Change
CO_2 (in g/km)	147	126	-14%
Horsepower (in kW)	77	80	3%
Footprint (in m^2)	7.2	7.4	2%
Weight (in kg)	1271	1280	1%
Diesel	56%	56%	0%

The Table presents the sales weighted vehicle characteristics in the EU in 2007 and 2011.

Figure 2: Sales Weighted Characteristics over Time

(a) Europe **(b)** United States

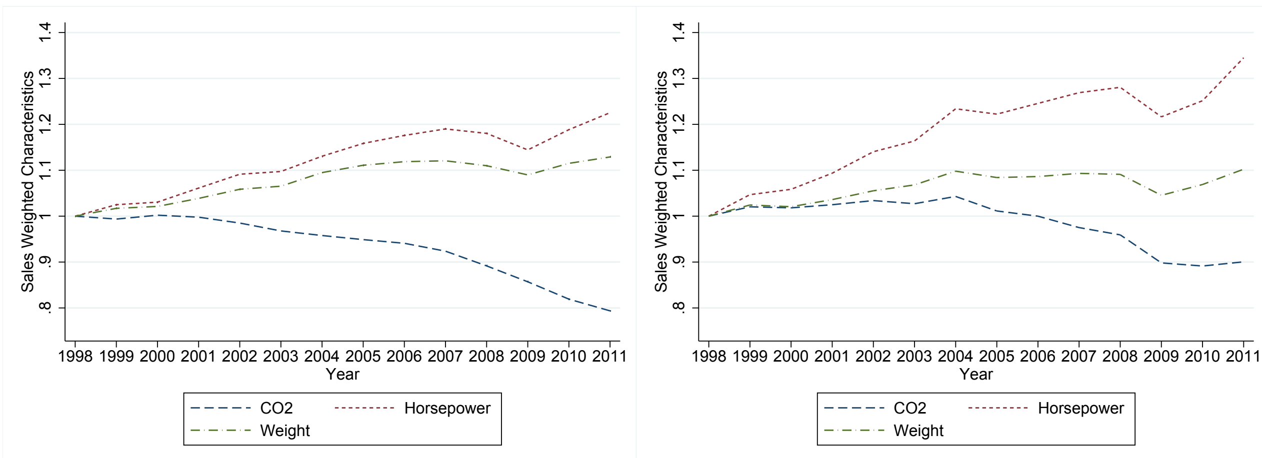

The figure shows the evolution of quantity weighted characteristics from 1998 until 2011, indexed at 1998. The EU trends represent the evolution of sales weighted characteristics as observed in the data. The US trends represent the evolution of the production weighted characteristics as reported by the EPA (http://www.epa.gov/otaq/fetrends.htm).

four years of the sample. This shift coincides exactly with the announcement of the fuel efficiency standard by the European Commission. The CO_2 emissions show a very different pattern in the US than in the EU. Until 2007 there is a very moderate decline in emissions of 3%. Between 2007 and 2009, the emissions of newly produced vehicles decline by 7% but then remain constant at 90% of the 1998 level. In the EU, emissions further decrease in 2010 and 2011 are at 80% of the 1998 level by the end of the sample. In both the EU and the US weight and horsepower grow consistently over time. By 2011, European consumers choose a vehicle that, on average, is 23% more powerful and 13% heavier than in 1998.

3 Market response to the EU emission standard

In this section I decompose the decrease in carbon emissions in the EU vehicle market. The goal of the decomposition is to measure the extent to which different abatement strategies can explain the drop in emissions after the policy announcement. To do this, I estimate a trade-off relation between emissions and other vehicle characteristics. The shifts over time in the relation between characteristics and emissions supplied by car makers are decomposed into changes in technology and changes in the composition of the vehicle fleet. I also investigate to what extent the technological change translates to the road by separating gaming from actual technology.

Estimation of trade-off and technology parameters Following Knittel (2011), I estimate the following regression:

$$\log(e_{jt}) = \zeta_t + \eta \log(x_{jt}) + \epsilon_{jt}, \tag{2}$$

where the technology parameter ζ_t is a time fixed effect, the trade-off parameters η denote how emissions e_{jt} change with a 1% change in characteristic x_{jt} and ϵ_{jt} is an error term.[16] The technology parameter captures shifts over time in the trade-off between emissions and characteristics. When the trade-off parameters η are constant over time, technology ζ_t can be seen as input neutral as it enters multiplicative in levels.[17] I include vehicle model fixed effects such that the identifying variation comes from different engine options within the same model. I assume the remaining unobservable ϵ_{jt} to be i.i.d.. The unobserved error comes from variation in unobserved attributes of engine versions within a model name, such as torque or the valve mechanism. These unobserved engine attributes might be correlated with explanatory variables, so that η cannot be interpreted as a causal relation between the characteristics and emissions. The goal of estimating (2) is to see changes in emissions over time while controlling for the correlation between emissions and other characteristics.

The estimation of (2) is useful because it reveals the abatement strategies that firms resort to without relying on a structural model. In Figure 2, I showed how emissions starkly decreased after the regulatory announcement. If firms choose sales-mixing or downsizing, then the part of the emissions that is explained by characteristics x_{jt} should decrease over time. Both with sales mixing and downsizing firms would sell more small and low performing cars. In contrast, when firms choose to implement technology, we should see shifts in the technology parameters ζ_t over time. As discussed above, there is large concern that firms have gamed the test. This is a concern when estimating (2), as the left hand side variable is the official emission rating obtained from the regulator. To test the extent to which emissions on the road change, I also estimate (2) using a measure of on-road emissions as the left hand side variable. The on-road emissions are taken from our work in Reynaert and Sallee (2019). This allows me to test if the technological improvements revealed in the official data translate to on-road measures. I now present estimates of the trade-off parameters in Table 2 and of the technology parameters in Table 3. Next, I discuss the robustness of estimating (2), which is reported in appendix Table A1. Finally, I show the decomposition of emissions over time in Table 4.

[16] Knittel (2011) assumes that the marginal cost is additively separable in elements related to emissions and other cost elements not related to emissions. Then this estimation can be interpreted as an estimate of the level set or iso-marginal cost relation between emissions and its determinants.

[17] The main specification will have a constant η, but I also show robustness when η changes over time. I estimate 2 using a wide range of robustness checks, which are specified below and are reported in Appendix Table A1.

Table 2 presents the trade-off parameters η from estimating (2). Model 1 is the baseline specification close to that of Knittel (2011), and includes trade-off parameters for horsepower, weight, footprint and height. For Model 1, I find that a 10% increase in horsepower is associated with a 1.8% increase in emissions. Weight is associated with higher emissions, while a 10% increase in the footprint reduces emissions by 2.9%. A diesel engine is approximately 20% more efficient than a gasoline engine, which coincides with engineering numbers and is very robust across specifications. These numbers have the same sign and a similar magnitude as those reported by Knittel (2011) and Klier and Linn (2016) who use similar European data. Model 2 allows for a firm specific trend in technology. The trade-off parameters are robust to introducing firm specific trends. Next, I start introducing the on-road measures of emissions. These measures are available for a much smaller set of vehicles so that the sample shrinks from 14,444 observations to 3,766 observations. Therefore, Model 3 replicates Model 1 with the smaller sample. The trade-off parameters shrink somewhat, but are fairly robust. Model 4 changes the dependent variable so that the relation between the on-road emissions and characteristics is estimated. Model 5 introduces firm specific trends. The trade-off parameters shrink further when looking at on-road emissions. However, it becomes more difficult to interpret these coefficients as a technical trade-off that firms face. On-road emissions vary not only because of physical reasons but also because consumers with different driving patterns select different types of vehicles. For example, the diesel coefficient now deviates from the engineering estimates and is 0.16 instead of 0.2. Typically, long distance drivers choose diesels because of their higher fuel efficiency. If long distance drivers obtain lower fuel economy than short distance drives, this explains the deviation from the engineering estimate. Below, the trade-off parameters will be used to compute the part in emissions that correlates with characteristics. First, I zoom in on the time fixed effects.

The technology parameters ζ_t are derived from the time fixed effects in each regression and are plotted in Table 3 for Models 1 to 5. Note that Models 2 and 5 report the averages of the firm specific trends. The bottom panel of the table presents the difference in the mean of ζ_t for 1999-2007 and for 2008-2011, corresponding to the pre- and postpolicy announcement period. Models 1-3 all explain the official emission ratings, and I find that technology improves significantly faster in the later years of the sample. All estimates show that the technology shifts ζ_t are more than 2% higher after 2007 than in the years before. The difference is statistically significant. After 2007, the estimates thus reveal a significant increase in the pace of technology improvement. The firm specific trends are reported in Appendix Table A2. The change in trends is there for all firms but the magnitude differs. BMW and Renault have the lowest difference in technology growth (0.9% and 0.8%), while Volkswagen has the highest difference (5.4%). In the official emission measures, we find evidence of rapid changes in technology adoption, but these emission reductions might not translate to the road. Indeed, once we explain the shifts in on-road emissions in Models 4 and 5, we find that the technology growth does not increase as rapidly in the postpolicy period. Between 2008-2011, the mean technology change shrinks from 3.9% in Model 2 to a mere 1.5% in Model 5. However, Models 4 and 5 reveal higher average technology growth in later years relative to prior years, but the difference shrinks from 2.4% to 0.7%. Firm specific values, reported in Table Appendix A2, show similar patterns except for BMW and Renault. Overall, this reveals that firms do engage in technology adoption after the policy announcement but that a large part of the technology is aimed at reducing official emission ratings rather than on-road emissions. The estimates show that 70% of the additional reduction in official emissions comes from gaming, while 30% comes from actual technology improvements. This 30/70 ratio will be used in the structural model below.[18]

To show robustness for the estimation of trade-off parameters and trends in technology, I replicate the estimation of 2 in several robustness exercises. Table Appendix A1 shows the estimates and differences in the technology take up for Model 6, which changes the functional form from Cob Douglas to Translog so that higher order terms in x_{jt} are included. Model 7 keeps only the first appearance of each vehicle; Model 8 allows the trade-off parameters to change over time; Model 9 weighs the observations by sales; Model 10 includes the marginal costs as estimated in the structural model as a control; and Model 11 interacts the characteristics with the fuel type. All specifications are ran on the full sample with official ratings as the dependent variable. All specifications estimate that the emissions decrease between 2.1% and 3.3% more rapidly after the policy announcement. The baseline in Model 1 is 2.3%.

Decomposition of the changes in emissions The estimated relation (2) can be used to decompose the decline in emissions over time. The decomposition is done as follows: I predict emissions, $\widehat{e}_{jt}$, as the fitted values of regression (2) with Model 2. I also predict $\overline{e}_{jt}$ using (2) but fixing the technology level at $\zeta_t = \zeta_{2007}$.[19] The evolution in the sales-weighted values of $\overline{e}_{jt}$ only changes when the composition of characteristics x_{jt} changes over time, while $\widehat{e}_{jt}$ changes both because of underlying characteristics and technology.

The results in Table 4 show that between 1998 and 2007, sales-weighted emissions without technology $\overline{e}_{jt}$ declined slightly from 156 to 152. Technology improvements are responsible for the remaining moderate decline in emissions between 1998 and 2007. After 2007, the sales weighted emissions without technology $\overline{e}_{jt}$ remain constant. There is thus no evidence of significant changes that could be attributable to either sales-mixing or downsizing, as these would change the part of the emissions that correlates with the attributes. When I split up the average sales-weighted emissions into vehicle models released after and prior to 2007, the results show that the emissions $\overline{e}_{jt}$ of vehicles released prior to 2007 remain constant.[20] Vehicle models released after the policy

[18] I use the market level estimate of gaming rather than the firm level in the structural model for two reasons. First, the firm level estimates are not statistically precise and introducing firm heterogeneity in gaming will create large differences in welfare outcomes between firms. The preciseness of the estimates in Models 4 and 5 is overstated in the standard errors as the left hand side are estimates of on-road performance and I did not adjust the standard errors for this (bootstrapping is computationally impossible because it would require multiple estimations over 22 million observations. The second reason is that in Reynaert and Sallee (2019) we analyze data up until 2015 rather 2011. Over this longer period, all firms game in large and similar amounts.

[19] I re-scale each of the predicted emissions with the attribute-based target function, such that the numbers can be read as actual distances from the regulation. Not doing this re-scaling has no impact on the results.

[20] An example of a newly released model is the "Citroen DS3 Hatchback", which was released in 2009.

Table 2: Trade-off Estimates between Emissions and Characteristics

	Model 1		Model 2		Model 3		Model 4		Model 5	
	Official Ratings						On-Road Ratings			
	Coef.	St.E.	Coef.	St.E.	Coef.	St.E.	Coef.	St.E.	Coef.	St.E.
ln(Hp)	0.18	0.02	0.18	0.02	0.13	0.02	0.12	0.02	0.12	0.02
ln(Weight)	0.62	0.04	0.64	0.05	0.55	0.06	0.29	0.04	0.29	0.04
ln(Footprint)	-0.29	0.09	-0.33	0.10	-0.34	0.08	-0.08	0.07	-0.13	0.06
ln(Height)	-0.02	0.12	-0.00	0.12	0.00	0.12	0.06	0.08	0.04	0.08
Diesel	-0.20	0.01	-0.20	0.01	-0.21	0.02	-0.16	0.01	-0.16	0.01
Year F.E.?	Yes				Yes		Yes			
YearXFirm?			Yes						Yes	
Car Name F.E.?	Yes		Yes		Yes		Yes		Yes	
Observations	14,444		14,444		3,766		3,766		3,766	
R^2	0,85		0,86		0,78		0,88		0,88	

This table gives the trade-off parameters η between the characteristics and emissions from equation (2). Robust standard errors clustered per firm are reported between brackets. In Models 1, 2 and 3, the official $\ln(CO_2)$ ratings is the dependent variable, in Models 4 and 5, the on-road measures are the dependent. Model 1 includes year fixed effects; Model 2 introduces year by firm fixed effects; Model 3 limits the sample to data for which on-road estimates are available, but the dependent is still the official rating; Model 4 has on-road $\ln(CO_2)$ estimates as the dependent variable and has year fixed effects; Model 5 has year by firm fixed effects.

announcement are, on average, more polluting than the existing vehicle models. The difference between the existing vehicles and the vehicles released after the policy decreases over time however. The observed decline in emissions is thus not attributable to changes in the sales mix or to the release of new downsized fuel efficient vehicles. However, I cannot rule out that the vehicles would have grown without the policy.

The sales weighted emissions with technology $\widehat{e}_{jt}$ are decreasing rapidly after 2007 and this shows that technology adoption is fully responsible for the observed drop in the official emission ratings. Strikingly, the decrease in sales weighted emissions of older vehicles due to technology is as strong as the decrease in newly released vehicles and the engine improvements are installed widely across the fleet.[21]

There are two concerns regarding the results presented so far. First, the Great Recession and the EU debt crisis coincide with the policy announcement and cover the post policy announcement period. In general, an economic downturn would lead consumers to spend less on durables and, thus, would add to the likelihood of finding evidence for shifts in the composition of the fleet towards smaller vehicles. The fact that we do not see this is thus strengthening the argument that firms respond with technology adoption. Technology adoption itself could be affected by the crisis, but the direction of that effect is unclear. A second concern is that the observed response is so strong that most firms already complied with the emission standard in 2011, four years before the regulation is fully binding. Individual member states do increase their emission-based taxation and

[21]When I zoom in on vehicle models I find (not reported) that the likelihood of releasing an engine version with lower than existing emissions is significantly higher in the years after 2007.

Table 3: Technological Progress Estimates

	Model 1		Model 2		Model 3		Model 4		Model 5	
			Official Ratings					On-Road Ratings		
	Coef.	St.E.	Coef.	St.E.	Coef.	St.E.	Coef.	St.E.	Coef.	St.E.
1999	-0.5	0.8	-1.4	0.1	-1.7	1.4	-0.9	0.5	-1.2	0.1
2000	1.9	0.6	2.4	0.1	1.7	0.9	-1.5	0.2	-1.5	0.1
2001	-1.5	0.4	-2.2	0.2	-1.4	0.9	-1.2	0.2	-1.7	0.1
2002	-1.3	0.4	-1.5	0.2	-1.5	0.6	-0.9	0.3	-0.8	0.1
2003	-1.5	0.3	-1.9	0.2	-0.8	0.3	-0.5	0.2	-0.5	0.1
2004	-1.8	0.5	-2.1	0.2	-1.4	0.7	-0.7	0.3	-0.9	0.2
2005	-1.5	0.3	-1.9	0.2	-0.9	0.2	-0.2	0.2	-0.1	0.1
2006	-1.3	0.3	-1.7	0.1	-1.1	0.5	-0.4	0.2	-0.4	0.1
2007	-1.4	0.6	-2.2	0.1	-0.9	0.9	-0.2	0.5	-0.5	0.1
2008	-2.7	0.4	-3.1	0.1	-2.2	0.8	-0.9	0.4	-0.9	0.0
2009	-3.0	0.6	-3.6	0.1	-3.2	1.0	-1.4	0.5	-1.4	0.1
2010	-4.3	0.7	-4.7	0.1	-5.5	1.6	-2.5	0.5	-1.8	0.0
2011	-3.3	0.4	-4.5	0.1	-3.2	1.1	-1.3	0.6	-1.9	0.1
	Difference in Technology Growth 2011-2007 and 2007-1998									
Difference	2.3	0.5	2.42	0.07	2.6	0.9	0.8	0.3	0.65	0.08

The table gives the estimated yearly percentage change in emission rates due to technological improvements. The percentages are obtained from differencing the year fixed effects in (2), and the standard errors are computed with the Delta method. Each of the estimated models correspond to Table 2. The shaded area are years after the policy announcement. The last rows give the difference in average technology growth between 2011-2007 and 2007-1998. For Model 2 and Model 5 the average across firm specific trends are reported. The firm level estimates are reported in Appendix Table A2.

Table 4: Decomposing the Decrease in Emissions

		All Vehicles		Existing Models (2007≤)		New Models (> 2007)	
		No Tech.	Tech.	No Tech.	Tech.	No Tech.	Tech.
	True	$\overline{e}_{jt}$	$\widehat{e}_{jt}$	$\overline{e}_{jt}$	$\widehat{e}_{jt}$	$\overline{e}_{jt}$	$\widehat{e}_{jt}$
1998	169	156	172	156	172		
1999	168	156	170	156	170		
2000	169	155	171	155	171		
2001	167	154	169	154	169		
2002	164	154	166	154	166		
2003	161	153	162	153	162		
2004	158	152	159	152	159		
2005	156	152	157	152	157		
2006	154	152	155	152	155		
2007	151	152	152	152	152		
2008	147	152	147	151	147	157	152
2009	142	152	143	152	142	160	150
2010	135	152	136	152	136	154	137
2011	130	152	130	152	130	154	132

The table reports the observed and predicted levels of the average sales weighted CO_2 emissions. Emissions are corrected with the attribute function $f(w_j)$ and represent the actual target values for the regulation. All predictions use the estimates from Table 2 and Table 3 Model 2. The columns $\overline{e}_{jt}$ contain sales weighted predicted emissions keeping technology constant at $\zeta_t = \zeta_{2007}$. The columns $\widehat{e}_{jt}$ contain sales weighted predicted values for emissions with estimated ζ_t.

regulation in response to the EU wide policy.[22] This combination of new local taxation, together with the standard, can explain why the response is so strong and why compliance is attained early. It would be very interesting to study the interaction between national regulation and the EU wide standard, but this is out of the scope of the current project. In the remaining analysis, I will therefore model firm behavior in response to a standard to single out the effects of the EU emission standard.

In summary, I presented evidence that the observed decline in emissions is attributable to changes in the official CO_2 ratings while the sales mix and characteristics of the fleet are unaffected. The estimates show that 70% of the decline is explained by gaming and 30% is explained by actual technology adaption. These findings form the basis for developing an economic model in which firms can respond to the emission standard with technology adoption or gaming, keeping the other characteristics of their fleet unaffected.

[22] Examples are the bonus/malus system in France and low emission zones in Germany as well as various scrapping schemes.

4 Model

This section introduces an emission standard in a structural model of consumer demand and firm behavior. After describing the model, I discuss how different abatement strategies affect market outcomes. I argue that there are no clear theoretical predictions for the welfare effects of emission standards and discuss what I need to estimate. Finally, I discuss how the design choices of the regulator regarding attribute basing and enforcement affect the firm abatement choice.[23]

Demand There are M geographical markets, indexed by $m = 1, \ldots, M$, each market is observed $t = 1, \ldots, T$ times.[24] I suppress the subscript t. In each market m there are A_m potential consumers. Consumers are assumed to purchase only in the market where they are located. Each consumer i chooses one alternative j, which is either the outside good, $j = 0$, or one of the J differentiated products, $j = 1, \ldots, J$. Consumer i's conditional indirect utility for the outside good is $u_{i0m} = \varepsilon_{i0m}$, and for products $j = 1, \ldots, J$ it is:

$$u_{ijm} = x_{jm}\beta_i^x - \beta_i^e g_{jm}e_{jm} - \alpha_i p_{jm} + \xi_{jm} + \varepsilon_{ijm}, \tag{3}$$

where x_{jm} is a vector of observed product characteristics, $g_{jm}e_{jm}$ are fuel costs (fuel prices g_{jm} times fuel consumption e_{jm}), p_{jm} is the vehicle price and ξ_{jm} is an unobserved characteristic of vehicle j in market m, unobserved by the researcher but observed by consumers and firms. The parameter vector (β_i^e, β_i^x) consists of random coefficients, capturing individual-specific valuations for fuel costs and vehicle characteristics, α_i is the marginal utility of income or price valuation and ε_{ijm} is a remaining individual-specific valuation for product j (assumed to be i.i.d. type I extreme value). The random coefficient for characteristic k is given by $\beta_i^k = \beta^k + \sigma^k \nu_i^k$, where ν_i^k is a random draw from a standard normal distribution, so that β^k represents the mean valuation for characteristic k and σ^k the standard deviation across consumers.

Notice that the coefficient β_i^e measures the consumers' valuation of fuel costs. Consumers use vehicles for several years and thus care about the expected fuel costs over the vehicle lifetime. I follow the literature on fuel cost valuation (see, for example, Allcott and Wozny (2014) and Grigolon, Reynaert, and Verboven (2018)) and make the assumption that consumers expect fuel prices to follow a random walk. This is consistent with the survey evidence, as shown by Anderson, Kellogg, and Sallee (2013). The fuel price at the time of purchase multiplied by the vehicle fuel consumption then captures the expected cost per kilometer of travel in each year of usage. The parameter β_i^e estimates the mean and heterogeneity in the value for fuel costs scaled by mileage

[23]The abatement strategies discussed do not need to occur mutually exclusively. Firms will choose their abatement strategies such that the marginal abatement costs of each strategy are equal. When firms abate by choosing only one strategy the marginal cost of that strategy must be lower than that of the other strategies.

[24]I observe prices and quantities at the country level. Therefore, the geographical market will correspond to a country in the empirical analysis.

and a capitalization factor.[25,26] The valuation for fuel costs will be estimated with data from before the policy announcement in a period where the gap between the official and on-road ratings was constant. I assume that consumers know what the official ratings signal about the actual fuel costs when estimating the model. When I introduce the policy in the simulation, firms will start gaming the emission ratings. In that period it becomes unclear if consumers can determine if rapid decreases in fuel costs are due to gaming or actual technology adoption. I will consider both the case where consumers are sophisticated and where they are fooled.

Each consumer i in market m chooses the alternative $j = 0, \ldots, J$ that maximizes her utility. The predicted market share of vehicle j in market m is the probability that product j yields the highest utility across all available products (including the outside good 0). This is given by the logit choice probabilities, integrated over the individual-specific valuations for the continuous characteristics:

$$s_{jm}(\delta_m, \sigma) = \int \frac{\exp(\delta_{jm} + \mu_{jm}(\sigma, \nu))}{1 + \sum_{l=1}^{J} \exp(\delta_{lm} + \mu_{lm}(\sigma, \nu))} dP_\nu(\nu), \tag{4}$$

where δ_m, the mean utility, which collects all terms in (3) that do not vary across individuals, and μ_{jm} is the term that captures the individual idiosyncratic deviations from the mean utility as follows: $\mu_{jm}=\sum_k \sigma^k \nu_i^k x_{jm}^k$. To complete the demand side, I set the observed market share $s_{jm} = q_{jm}/A_m$ equal to the predicted market share (4). In vector notation, the demand side in market m can then be described by the market share system as follows: $\mathbf{s}_m = \mathbf{s}_m(\delta_m, \sigma)$.

Firm Behavior To study firms' responses to the emission standard, I model a game in which firms have the following three choices: price setting, technology adoption and gaming. I assume that the vehicle fleet of firms is given exogenously. Each firm sells hundreds of differentiated products and I do not model the decision regarding which vehicles firms choose to offer. I also assume that at the time of decision, firms have perfect information on all observable and unobservable characteristics. Given their vehicle fleet, firms have the option to change the prices and fuel consumption of each of the products. The fuel consumption can be changed by adding technology or by gaming the test that determines the level of fuel consumption. When firms face a regulation, the choices will be altered to comply. The regulation is binding across the total sales from all EU markets. The goal of this section is to understand how firms will change decisions on pricing and fuel consumption levels when facing a binding regulation.

[25] The total expected fuel costs can be written as $E[\sum_{s=1}^{S}(1+r)^{-s} m_i g_s e_j]$ where s=1 is the time of purchase, S is the time of scrappage, r is the interest rate and m is mileage. Using the random walk assumption for g, we can write expected fuel costs as $\rho m_i g_1 e_j$, with ρ being the capitalization coefficient. See Grigolon, Reynaert, and Verboven (2018) for a detailed discussion. In the utility specification in (3) β_i^e absorbs ρm_i.

[26] The specification does not allow consumers to separately care about emissions (a 'green glow' effect). These preferences might be captured by the standard deviation in tastes for fuel costs as green consumers will value fuel costs more than others.

The total profit per year t is the sum of the profits from each country m, as follows:

$$\max_{p,\tau,g}(\sum_m[\pi_{fm}(p,\tau,g)] - C(\tau,g))\ s.t.\ \frac{\sum_m\sum_{j\epsilon F_f} q_{jm}((1-\tau_j-g_j)e_j - f(w_j))}{\sum_m\sum_{j\epsilon F_f} q_{jm}} \leq \kappa, \tag{5}$$

where π are variable profits, C are sunk costs of changing τ and g, κ is the level of the standard and $f(w_{jm})$ is the attribute-basing on weight w_j. For a flat standard $f(w_j) = 0$, when $f(w_j) \neq 0$ vehicles with different weights will obtain reductions or penalties on their emissions.[27] I model technology adoption and gaming as percentage reductions in fuel consumption so that $0 \geq \tau \leq 1$, $0 \geq g \leq 1$ and $0 \geq \tau + g \leq 1$. I follow Goldberg (1998) and Jacobsen (2013) and write the Lagrangian of the problem. The variable profits of firm f in year t are then given by the following:

$$\pi_f = \sum_m \sum_{j\in F_f}[(p_{jm} - c_{jm} - \lambda_f L_j)s_{jm}A_m], \tag{6}$$

$$L_j = (1-\tau_j-g_j)e_j - f(w_j) - \kappa \tag{7}$$

where λ_f is the shadow cost of the regulation and L_{jm} is the distance of each product from the target. When $L_{jm} < 0$ (> 0), an additional sale of vehicle j will bring the average sales-weighted emissions closer to (further away) from the target. The per vehicle shadow cost λ_f gives the cost of deviating one unit from the standard. If the standard is non-binding, $\lambda_f = 0$ and (6) reduces to a standard multiproduct profit function. The shadow cost λ_f is firm specific because trading of excess emissions between firms is not allowed. The shadow cost takes the same value for each vehicle in the fleet F_f because, in equilibrium, firms will equalize the shadow costs over their fleet to be cost efficient.

The optimal solution of the profit maximization problem is described by three first order conditions. I introduce the following matrices to write the first order conditions concisely as follows: Φ denotes a JxJ ownership matrix with each element (i,j) equal to one if i and j are owned by the same firm and zero otherwise; $\triangle_k$ denotes the JxJ matrix of first order derivatives of market shares with respect to $k = p, e$ or g. The J first order conditions with respect to prices, technology and gaming can be written as follows:

$$\frac{\partial\mathcal{L}}{\partial p} = 0 = \mathbf{q} + \mathbf{\Phi} \circ \triangle_{\mathbf{p}}(\mathbf{p} - \mathbf{c} - \lambda\mathbf{L}) \tag{8}$$

$$\frac{\partial\mathcal{L}}{\partial \tau} = 0 = (-\mathbf{c}'_\tau + \lambda\mathbf{e})\mathbf{q} + \mathbf{\Phi} \circ \triangle_\tau(\mathbf{p} - \mathbf{c} - \lambda\mathbf{L}) - C'_\tau \tag{9}$$

$$\frac{\partial\mathcal{L}}{\partial g} = 0 = \lambda\mathbf{eq} + \mathbf{\Phi} \circ \triangle_{\mathbf{g}}(\mathbf{p} - \mathbf{c} - \lambda\mathbf{L}) - C'_g \tag{10}$$

The first order conditions with respect to prices show the standard trade off between increases in mark up (first term) and losses from reduced sales (second term) when increasing the price. The first order conditions with respect to technology adoption shows that firms will trade-off marginal

[27]Here I specify the attribute-basing as a simple additive penalty or reduction but one could design a regulation where the target is any function $g(e_{jm}, w_{jm})$ of emissions and the attribute.

cost increases $\mathbf{c}'_\tau$ and the fixed costs of adoption with slacking the regulatory constraint and the benefits of increased market share (we expect $\triangle_\tau > 0$). The first order conditions with respect to gaming is very similar but gaming will induce sunk cost changes and no marginal cost changes. Note that I model the first order conditions as joint contemporaneous choices and ignore dynamics.[28]

Externalities The regulation limits the sales-weighted emissions and not the total externalities from the new vehicle fleet. The amount of CO_2 emitted depends on both the size and composition of the fleet. Given the fuel consumption of each vehicle, we can compute the total lifetime damages from CO_2 emissions. This requires information on mileage and an estimate of the value of a ton of CO_2. When consumers purchase vehicles with lower fuel consumption, emissions will decrease. However, it might be that the standard results in more fuel efficient vehicles but also in more purchases. It is thus not certain that total emissions will decline. In the welfare calculation I will also consider other externalities, such as traffic, accidents and local pollutants.

Next, I discuss how equilibrium outcomes change when we move from a market without a standard (or a nonbinding standard), $\lambda_f = 0$ to a market with a binding standard, $\lambda_f > 0$. The changes in the market will depend on the abatement strategy of firms, i.e., sales mixing, technology adoption or gaming. I will discuss the effect of each strategy in turn, keeping the other strategies fixed.

Abatement by sales-mixing A first mechanism to abate emissions is to change the relative prices of high and low emission vehicles. Firms can decrease the prices of vehicles with emissions below the target ($L_j < 0$) while increasing prices of vehicles with emissions above the target ($L_j > 0$). The first order conditions with respect to prices (8) shows that the shadow cost of the regulation, λ_f, determines to what extent prices will be distorted from a no-policy equilibrium, as follows:

$$\mathbf{p} = \mathbf{c} + \lambda\mathbf{L} - (\mathbf{\Phi} \circ \triangle_{\mathbf{p}})^{-1}\mathbf{q} \tag{11}$$

The regulation implicitly increases costs and taxes vehicles with $L_j > 0$, while it is an implicit subsidy for vehicles with $L_j < 0$. This change in the relative prices of products will shift sales towards vehicles with lower fuel consumption, thus resulting in a different sales-mix, explaining the name of this strategy.

The incidence and effectiveness of this abatement strategy largely depends on the responsiveness of consumers to these price changes as captured in $\triangle_p$. Holland, Hughes, and Knittel (2009) show that when the price elasticities of subsidized products differ from those of taxed products, total sales and emissions might increase or decrease. Without knowledge of own and cross price elasticities we cannot make statements about the effect of the regulation on sales, emissions or consumer surplus. The effects on profits will depend on the consumers' responses to price changes but will also depend on the position of the fleet relative to the target. Firms with a fleet that

[28] In reality, firms choose to incrementally adopt technology and gaming over the several years between the policy announcement and enforcement. In ignoring this process, I miss how firms could make strategic abatement choices that take into account the path dependency of their choices.

is better adapted to the standard might increase profits, as their prices will need less distortion than the prices of competitors. The empirical model will allow me to identify the own and cross price elasticities for all products and to simulate the shifts in sales when we introduce a binding regulation.

Abatement by technology adoption Firms can reduce the emissions of existing vehicles by adapting the engines, the combustion process or features that only affect the fuel consumption.[29] From (9) we see that the regulation gives incentives to the firms to adopt more technology. The first term, $(-c'_\tau + \lambda e)q$, specifies that increases in the fuel efficiency change the marginal cost of each unit sold, but the technology also makes the regulation slacker, resulting in a marginal benefit of λ. Without a regulation, $\lambda = 0$ and the marginal benefits of technology adoption would thus be lower. The second term captures the change in sales from technology, $\Phi \circ \triangle_\tau$, multiplied by profits per unit, $(p - c - \lambda L)$. Finally, implementing technology can increase the sunk costs captured by C'_τ.

It is important to stress that each unit of technology adoption lowers L. The fleet of the firm shifts closer to the regulatory target and, as such, the regulatory constraint becomes slacker. This reduces the shadow cost of the regulation λ_f as firms need less distortions from the preferred price schedule to comply. When determining to what extent firms adopt technology or sales mixing in equilibrium, the shadow cost λ and the amount of technology will be determined endogenously.

The welfare effects of this strategy are again undetermined theoretically. This time, there are two offsetting effects for consumers. There is upward pressure on prices as marginal costs increase with c'_τ. This reduces consumer surplus. However, offsetting this, firms offer vehicles with lower fuel consumption, decreasing the cost of operating a vehicle. The sum of the purchase price and operating costs might thus increase or decrease. Even though the resulting vehicle fleet will have lower emissions, it is not clear that emissions will decrease, i.e., if consumers benefit from technology, they might purchase more vehicles. The changes in marginal costs and the degree of pass through will determine the overall effect. Given that vehicles have different fuel consumption, the equilibrium pricing will also adjust, creating another welfare change for consumers. For firms, technology adoption will be a preferred abatement strategy if the shadow cost of changing prices is high and the marginal cost and sunk cost changes of technology adoption are low.

Abatement by gaming I define gaming as the efforts of the firm that decrease the official emission ratings but not the actual on-road emissions. Gaming helps with compliance in a very similar way as technology adoption. Each unit of gaming will reduce L and thereby the shadow cost of the regulation λ_f until firms can price without regulatory constraints. However, the benefits

[29] Knittel (2011) gives several examples of specific technologies that are implemented. The International Energy Agency reported a possible 40% improvement in fuel efficiency from available technologies in 2005. These include low rolling resistance of tires, reduced driveline friction, combustion improvements, thermal management, variable valve actuation and lift, auxiliary systems improvement, thermodynamic cycle improvements and dual clutch transmission. See http://www.iea.org/publications/freepublications/publication/technology-roadmap-fuel-economy-of-road-vehicles.html.

and costs of gaming are very different. In Reynaert and Sallee (2019) we empirically study the extent of gaming and discuss the effects of gaming on buyer welfare in detail; here, I summarize what is needed to understand firms choices and welfare implications.

The costs of gaming are very different than those of technology adoption. As summarized in media reports following the Volkswagen crisis, gaming is essentially a sunk cost that does not affect the marginal costs of production. The New European Driving Cycle (NEDC) is the procedure to determine CO_2 emissions and fuel consumption.[30] The procedure takes a single vehicle (the golden vehicle) and optimizes that vehicle for the test. Tape, non-resistant tires and software, referred to as defeat devices, make the engine run in a program specialized for the test. Note that none of these features impact how the consumers experience vehicles on the road, nor do they impact the marginal cost of production. The sunk costs of gaming are the costs of installing the defeat devices and more importantly, the potential legal costs and consumer blow-back when gaming is uncovered. When deciding to game, firms will trade off these costs against the benefits of cheap compliance. Since firms have opted to game the regulation, it reveals that firms did not expect these reputation costs to be larger than the costs of other compliance mechanisms. In the simulation I discuss computation of bounds on these sunk costs.

Note that the second term in first order conditions (10) accounts for market share changes of gaming. Gaming will not affect the on-road experience of consumers and will not reduce the actual fuel consumption. However, when purchasing a vehicle, consumers might be fooled by the advertised fuel consumption and believe it to be the truth. The official ratings were designed to be used in advertising and windshield stickers to inform consumers about fuel costs and pollution. As such, gaming will have demand effects that the distort consumer choice. In the results, I will discuss what happens when consumers are fooled or when they are fully aware of the gaming. The effects for consumer welfare will critically depend on this consumer sophistication. If consumers do not see through gaming, the information in the official ratings becomes a noisy signal of fuel costs and might therefore distort consumer choices. This choice distortion reduces consumer surplus, and might also lead to higher prices. Firms with market power will increase prices if the perceived fuel costs reduce, but gaming might also have positive effects for consumers, as it allows firms to avoid costly compliance and high pricing. If technology adoption or sales-mixing is very costly, then consumers might actually prefer firms to game, benefiting from avoiding the cost increases or price distortions. The effects on emissions from gaming can be pervasive. First, each reduction in fuel consumption that is obtained with gaming does not reduce actual emissions. Second, consumers will respond to the decreases in perceived fuel costs by substituting to higher performance cars. Because cars look cheaper on paper, the overall sales might increase. Both of these effects potentially increase emissions relative to honest compliance.

[30]The NEDC is also used to determine compliance with the Euronorms that regulate local pollutants such as NO_x, CO and PM. The Volkswagen scandal was about NO_x emissions in diesel vehicles. The fallout of the scandal made clear that the majority of EU car makers are not in compliance with Euronorms, see the ICCT reports.

Abatement by downsizing I assume the fleet of each firm to be given. Previous work, see for example Klier and Linn (2012) and Whitefoot, Fowlie, and Skerlos (2017), has tried to relax this assumption by modeling not only improvements in fuel consumption but by allowing firms to optimize a larger set of attributes. This strategy is coined downsizing as firms reduce the size and power of vehicles to lower fuel consumption. Klier and Linn (2012) find that compliance costs decrease by approximately 40% per year when firms down size instead of sales mix. The consumer loss is similar. Whitefoot, Fowlie, and Skerlos (2017) use an interesting engineering model. In this paper I will not discuss downsizing for two reasons. First, as I showed in Section 3, I find no evidence of downsizing in the EU. Second, a model that would allow strategic choices over several characteristics for the rich engine version level data that I use is out of the scope of this paper.

Design of emission standards One of the key contributions of this paper is to show how political choices about the design features of emission standards interact with the firms' abatement choices. In the counterfactual, I will consider two design choices of the regulation. First, I study the EU's choice to have a policy that is attribute-based and therefore upsloping, as depicted in Figure 1. Second, I study the implications of weak enforcement of the regulation.

It is instructive to compare the attribute-based regulation with a flat tax. For a flat standard, $L'_{jm} = [e_{jm} - \kappa']$ and $f(w_{jm}) = 0$. The target function is a horizontal line at κ' and all firms need to reach exactly the same level of sales-weighted emissions. Each firm will have a different set of vehicles with $L_{jm} < 0$ and $L'_{jm} < 0$. Because of the upsloping ABR in the EU, the sales of many small lightweight vehicles do not help with compliance, while they would have helped with compliance under a flat standard. As such, the attribute basing reduces the number of products that have $L_{jm} < 0$. It follows that sales mixing becomes much costlier because firms have less products to which they can shift sales. Reducing fuel consumption by technology adoption or gaming helps to increase the number of products that have emissions underneath the target. In the empirical section I will show that the ABR increases the stringency for French and Italian firms but reduces the stringency for German firms. This matches the lobbying efforts by the governments described at the time of the negotiations. I will also show that sales mixing became so expensive for firms such that they had to resort to lowering fuel consumption.

The attribute-based regulation might have other economic consequences. Ito and Sallee (2018) point out that attribute-based standards create a distortion in the demand and supply of the attribute itself. If heavier cars help with attaining the target, weight is indirectly subsidized and producers will choose to add more weight to their vehicles.[31] In this exercise, I keep the weight fixed and assume that there are no distortions in the attribute itself. The reason is that I do not observe any changes in weight consistent with this distortion and that endogenizing additional characteristics is computationally costly.

The EU's political structure has also led to weak enforcement of the policy. In the empirical section I will discuss how failures in delegation of enforcement have led to the gaming crisis in

[31]This creates distortions, which might be significant if weight is associated with other external costs. See, for example, the analysis by Anderson and Auffhammer (2014) relating weight to accident risk.

the EU market. Theoretically, enforcement is captured by the sunk cost of gaming $C(g)$. When enforcement would be strict and legal punishment for gaming high, the sunk costs of gaming would increase. This would mean that gaming as an abatement strategy becomes less attractive so that firms would resort to sales mixing or technology adoption. In the empirical section I will discuss what would have happened if firms had not gamed the EU regulation.

5 Estimation

In this section, I describe the estimation. I use a panel of 7 countries over 10 years to estimate the taste and cost parameters. The sample is restricted to 1998-2007, the years before the policy announcement. This allows me to estimate a model in which firms maximize unconstrained profits as given in (6) with $\lambda = 0$. It also allows me to estimate the consumer valuation of fuel costs before firms start to game emission ratings. I first discuss the demand estimation that allows for endogeneity of both prices and engine characteristics. Next, I discuss the supply side and estimation of costs. I also discuss the functional form assumptions needed to make a projection of marginal costs out of sample to model technological improvements. With the estimates of consumer preferences, marginal costs and the functional form assumption we can simulate the impact of the emission standard in the next section.

Demand Estimation The vector of parameters θ to be estimated consists of the taste parameters β_i^e, β_i^x and α_i. I estimate both a mean and a standard deviation of the taste for fuel consumption, horsepower, and a dummy for foreign perceived cars (e.g. a BMW in France). I specify α_i to be proportional to income y_{mt} in market mt, so $\alpha_i = \alpha/y_{mt}$.[32] A set of controls is added for which I only estimate the mean taste. These include weight, footprint, height, a dummy for 3 or 5 doors, market fixed effects, diesel by market interactions, months on market dummies (for vehicles introduced within a calendar year), and a market specific time trend.[33] Finally, I add fixed effects on the vehicle model level so that all identifying variation for the taste parameters comes from different engine versions within the same vehicle model. The remaining unexplained variation in market shares is ξ_{jmt}. The parameters are obtained by minimizing the GMM criterion as follows:

$$\min_{\theta} \xi\left(\theta\right)' g(z)' A \xi\left(\theta\right)' g(z)' \tag{12}$$

where ξ is a vector of demand unobservables stacked over all markets, $g(z)$ is the matrix of instruments and A is a weighting matrix. I use the two step efficient GMM estimator so that the second

[32] Assume that the utility is Cob Douglas in income and characteristics of the good, as in Berry, Levinsohn, and Pakes (1995). The logarithm of indirect utility can be written as $u_{ij} = \alpha log(y_i - p_j) + u(x, g, \xi, \epsilon, \theta)$, where the first part is utility from income and the second part utility from consumption of the good. Assuming that $p_j << y_i$ we can write $log(y_i - p_j) \approx log(y_i) - p_j/y_i$ and substitute $\alpha_i = \alpha/y_{mt}$ in (3).

[33] I introduce random coefficients on 3 variables that capture important margins on which I expect consumer heterogeneity to matter. Height, weight and footprint do not have a random coefficient because the identification of multiple random coefficients proved impossible with multiple endogenous characteristics. The remaining variables try to control for market or time level shifts and seem less fit to be candidates to model individual heterogeneity.

step A is the estimate of the optimal weighting matrix. I follow the estimation algorithm described in Berry, Levinsohn, and Pakes (1995) and Nevo (2001). I take into account recent cautionary warnings and improvements and carefully check the properties of the obtained minimum.[34]

A common exclusion restriction in the literature is $E[\xi|x] = 0$, so that any function of observed characteristics is a valid candidate to form the unconditional moments. This means that we allow for correlation between prices and ξ but assume that all the vehicle characteristic choices of firms are independent of ξ. The counterfactual considers strategic choices in response to a regulation. However, even before the regulation firms will design products to maximize profits. The concern when estimating demand is that firms know more about consumer tastes than the econometrician when the product is designed. This potentially introduces omitted variable bias through correlation of the demand unobservable ξ and product characteristics, despite the rich set of fixed effects. To account for this, I introduce additional instrumental variables. The identification assumption is $E[\xi|z^k] = 0$ so that the demand unobservables are mean independent from instruments z. There are three sets of parameters for which we need to construct moments, which are as follows: the taste parameter for price, the taste parameters for endogenous characteristics and the nonlinear parameters. First, I will consider instruments z^1 for the model with price endogeneity. Second, I will consider instruments z^2 for the more general case with endogenous characteristics.

Instruments The instrument set z^1 contains all characteristics and demand shifters, except price, as included instruments. The excluded instruments are the sums of other product characteristics (both the sum across all competing firms' products and the sum across products of the same firm). These are valid instruments for prices only when the characteristics are mean independent of the demand unobservable. Using the location where each vehicle is produced, I include the logarithm of local labor costs in the country of production as a cost shifter.

For instrument set z^2, I follow the recent work by Whitefoot, Fowlie, and Skerlos (2017) who use engineering procedures to distinguish between fixed vehicle characteristics and more mutable characteristics. In the first engineering step, the dimensions of the vehicle (footprint and height) are set. In following steps, the engines are fitted in the design. I therefore assume that the footprint and height are exogenous and remain in the set of included instruments. Fuel costs, horsepower and weight are allowed to be endogenous so that we need excluded instruments for the following four variables: the three mutable characteristics and prices.

Because the mutable characteristics are potentially correlated with ξ, their sums (across own or other firm products) are not valid instruments anymore. I therefore form additional instruments that aim to capture variations in the engine design decisions of firms. Using global production data,

[34] More specifically I do the following: (i) I use a nested-fixed point (NFP) algorithm, BLP's contraction mapping with a tight convergence criterion (1e-12) to solve for ξ_{jmt}. I use a NFP because Mathematical Programming under Equilibrium Constraints proved to be slower in this application once I paralyzed the computation of the contraction mapping, (ii) I re-estimate the model with 25 different starting values for the nonlinear parameters, (iii) I check first and second order conditions at the obtained minimum, (iv) I use the Interior/Direct algorithm of Knitro, (v) I compute the integral over individual market shares using sparse grids, see Heiss and Winschel (2008), (vi) I estimate the variances of the random coefficients rather than the standard deviations, see Ketz (Forthcoming).

I exploit changes in the exposure of each vehicle model to non-EU markets over time. I compute the share of each vehicle model that is produced in Africa, Asia, Eastern Europe, North America and South America. These production shares evolve within a model over time as some models become popular in the US (e.g. the BMW X5) or in China (e.g. the AUDI A8). A vehicle designed for the EU and US market will be different than a model designed for the EU and China. Because of the vehicle fixed effects, the identifying variation of the IV's comes from the trend in globalization within vehicle models. I assume that the differences in exposure to other markets impact the design choices but are orthogonal to the EU unobserved demand ξ_{jmt}. In the production data I also observe a size variable, scaled from 1 to 10, for each vehicle produced. Using these observations, I compute a weighted sum of size for each model and brand. The weights are the regional production shares and the size is the production weighted average size of all vehicles produced in a region. Similar to the sums of characteristics instruments, this captures the degree of competition in the product space for each vehicle from other vehicles across all the regions where the vehicle is produced. Next, I follow Klier and Linn (2012) and include the average footprint and height of vehicles in different classes that are produced on the same production platform. For example, the AUDI A5, which is in the luxury class, is produced on the same platform, named MLB, as the AUDI Q5, an SUV. The average attributes of the Q5 are used as an instrument for the A5, and vice versa. The idea behind the instrument is that vehicles produced on the same platform will share both fixed and mutable characteristics.[35] Finally, fuel costs are the interaction of fuel prices times fuel consumption. Fuel prices are exogenous, so I interact them with the projection of fuel consumption on all instruments.

To improve the efficiency of the estimates of the variances of the random coefficients, I compute approximate optimal instruments for the nonlinear parameters following the approach described in Berry, Levinsohn, and Pakes (1999) and Reynaert and Verboven (2014). To approximate the infeasible optimal IV's, I perform a two stage approach, first estimating the nonlinear model with a guess for the nonlinear parameters in the first step of the GMM and then updating the approximation at the first stage estimates.[36] This procedure generates a number of additional instruments that is equal to the number of standard deviations of random coefficients to estimate but these instruments are nonlinear functions of the previously described included and excluded instruments. Appendix 7 gives a detailed overview of all the instrumental variables.

Costs The marginal costs c_{jm} are not directly observed but are obtained from the first order conditions of the firms' profit maximization. Using equation (8), the demand estimates and the fact that $\lambda = 0$ in the estimation sample, I compute c_{jm}. Similarly, we can use equation (9) to back out the product level estimates of c'. These estimates give prepolicy estimates of the slope of the marginal cost with respect to reductions in fuel consumption. Information on c' is useful because

[35] Notice that Klier and Linn (2012) include all the characteristics of different class vehicles on the same platform. This means that the exclusion restriction related to characteristics is specific to vehicle classes. I allow the exclusion restriction to be more general and include only the footprint and height in the computation of this instrument.

[36] The guess for the nonlinear parameters is the Logit taste parameter divided by 10. I compute the approximated optimal IV at $\xi = 0$ and use projections of the endogenous variables on the included and excluded instruments in the computation, see Reynaert and Verboven (2014)

the slope of c' in the final year reveals the cost changes that firms are willing to incur to provide the current levels of fuel consumption of their products. Further improvements would be more costly, but those would not be made because the benefit from the improvements is smaller than the costs. If the supply model is correctly specified, any cost function for further reductions in emissions must start at $c'_{\tau=0}$, otherwise it is at odds with the revealed choices of firms and consumers. I will therefore use the estimated c' as the intercept for the cost slope in the simulations.[37]

We also need an estimate for values of $c'_{\tau>0}$ when the regulation requires fuel consumption reductions above what we observe. I will use two approaches for this problem. The first approach is to rely on cost estimates by engineers. The EU Comission has relied on several studies to support the design of the EU emission standard, see TNO (2011). These documents describe several existing technologies to reduce emissions, at zero development cost. The policy report only includes improvements that should be readily available for the car makers at no fixed costs. The studies specify convex cost functions for percentage reductions in emissions. I will use these functions in the main results and I will refer to them as the engineering cost curve. However, I shift the intercept in the engineering cost function from zero to the point where the slope of the engineering function equals $c'_{\tau=0}$.[38]

The second approach is to rely on historical marginal costs in the data to obtain information on how costs change with fuel consumption. Under the assumption that marginal costs are log-linear we have the following:

$$\log(c_{jm}) = \gamma^e e_{jm} + d_{jm}\gamma^d + \omega_{jm}, \tag{13}$$

in which d_{jm} is a $1 \times L$ vector of observed product characteristics, market controls and cost shifters and ω_{jm} is unobserved. Fuel consumption enters the marginal cost, as all else equal, it is likely to be more expensive to produce engines with lower fuel consumption. The estimated parameter on emissions informs us how marginal costs change with changes in e_{jm}. The functional form allows us to make predictions on the costs of further reductions in fuel consumption that are not observed in the sample. In the counterfactuals, I will show how the results change when we rely on the estimated cost function rather than the engineering cost function.

Estimation Results Table 5 reports the estimated parameters and standard errors for the demand model. The table presents four specifications. The first two specifications assume exogenous characteristics and use instrument set z^1 in both a logit and a RC logit. The last two specifications allow for endogenous characteristics and use instrument set z^2. The first stage results are reported in Appendix Table A3. When prices are the sole endogenous variable, the excluded IV's are strong and have an F statistic of 67. The labor cost instrument has the expected positive sign. With mul-

[37] The counterfactual results will also consider a case where we start from another intercept below the estimated one. As I will discuss below there might be reasons, such as market failures in technology adoption, to think that the actual slope is smaller than the estimated $c'_{\tau=0}$.

[38] The engineering estimates specify a polynomial function for the cost of emission reduction. I use that function and compute its derivative at each point. See Table 79 in the report TNO (2011). I use separate cost functions for diesel and gasoline and the coefficients for medium sized vehicles (the reported cost functions only differ between size categories for very large reductions in emissions of more than 40%).

tiple endogenous variables, the instruments are weaker with F stats that account for the multiple endogenous variables being between 5.7 and 21.2.

The demand parameters for the model with endogenous prices show that consumers dislike higher prices, higher fuel costs and cars perceived as foreign (BMW in France). Consumers prefer vehicles that are more powerful, heavier and larger. In the RC logit, the standard deviations on horsepower and foreign show that there is considerable variation in the taste for horsepower and weight. The estimates change considerably when I allow for multiple endogenous characteristics. The price coefficient decreases from -6.7 to -8.1, while the fuel cost parameter increases from -2.6 to -1.2. This will matter for the welfare simulations, i.e., when consumers care less about fuel costs relative to price, they will derive lower utility from technologies that drive up costs to decrease fuel costs. The remaining coefficients have larger standard errors than in the model with less endogeneity, and several mean tastes switch sign in the RC Logit model. The standard deviations show considerable heterogeneity in the taste for fuel costs, horsepower and foreign vehicles in this case.

In Appenix Table A4 I discuss the fit of the demand model. By definition the model fits perfect within sample. However, I show how well different parts of the demand explain sales-weighted attributes. First, I predict sales by only relying on the taste parameters for attributes while setting the vehicle model fixed effects and demand unobservable equal to zero. This shows that, based on taste parameters alone, consumers would buy vehicles with lower attributes. The model name fixed effects, on average, explain purchases of vehicles with high attributes and prices. Out of the estimation sample, in calendar year 2011, we cannot rely on fixed effects and demand unobservables. Prediction based on taste parameters alone shows the demand model explaings attributes purchased relatively well, and not worse than the prediction error from leaving out the unobservable and fixed effects within sample.

The second panel of Table 5 gives the results from a regression of marginal costs on product characteristics. I show the results under perfect competition (a regression of prices on characteristics) and imperfect competition for both the RC Logit models. These results show that cost shifters have the expected sign, i.e., increases in labor cost increase marginal costs and larger more powerful cars are also costlier. All marginal cost regressions show that increasing the fuel efficiency of the vehicle is costly. A one unit decrease in the liters per 100km increases the cost by 2.8% to 4.7% in the two specifications.

Finally, I use the model to compute the cost slope that rationalizes the fuel consumption choices of firms before the policy. Additionally, there is an important difference here when I allow for more endogenous variables. When consumers care less about fuel costs it becomes less interesting for firms to provide fuel consumption reductions. As the marginal benefit of fuel consumption reductions decreases, so does the implied marginal cost slope in (9). Indeed, the model with endogenous characteristics implies a prepolicy intercept of the cost slope that is five times less steep than what the price endogeneity model implies. The level of cost increases caused by the policy will be lower when the initial intercept is lower, so that it is more likely that the emission standard will increase

Table 5: Estimation Results

	Demand Estimates							
	Price Endog. (z^1)				Charact. Endog. (z^2)			
	Logit		RC Logit		Logit		RC Logit	
	Est.	St. Err.	Est.	St. Err.	Est.	St. Err.	Est.	St. Err.
Price/Inc.	-6,709	0,365	-6,476	0,325	-8,074	0,569	-8,452	0,694
Fuel Cost	-2,649	0,120	-2,842	0,127	-1,291	0,192	-1,036	0,269
Horsepower	2,889	0,227	1,232	0,152	3,868	0,673	-1,546	0,760
Weight	0,666	0,181	1,248	0,156	-6,130	2,143	1,530	1,300
Base	0,580	0,056	0,526	0,051	1,499	0,229	0,878	0,179
Height	0,183	0,041	0,215	0,037	0,477	0,091	0,291	0,060
Foreign	-0,848	0,023	-1,218	0,043	-0,698	0,042	-0,973	0,060
			Standard Deviations					
Fuel Cost			0,000	0,033			1,196	0,484
Horsepower			1,604	0,272			2,878	2,445
Foreign			1,276	0,192			0,897	0,185
			Marginal Cost Estimates					
	Perf.Comp.		Imp.Comp.		Imp.Comp.			
Fuel Cost	-0,028	0,001	-0,047	0,001	-0,041	0,001		
Horsepower	0,629	0,005	0,692	0,006	0,627	0,004		
Weight	0,664	0,009	0,670	0,011	0,600	0,008		
Base	0,092	0,002	0,026	0,003	0,021	0,002		
Height	-0,005	0,001	0,051	0,001	0,035	0,001		
Log Labor Cost	0,229	0,010	0,081	0,012	0,111	0,009		

The Table reports the estimated parameters for the demand and marginal cost equations. The first two columns in the Demand Estimates are the Logit and RC Logit for the model with endogenous prices and instrument set z^1. The last two columns are the Logit and RC Logit for the model with endogenous prices and characteristics and instrument set z^2. The additional controls in all the demand specifications are: market fixed effects, a dummy for 3 doors, months on market dummies (for vehicles introduced within a calendar year), fuel type by market dummies, a time trend and 331 vehicle model fixed effects. Marginal Costs slope estimates are reported for perfect competition in the first column (regression of price on characteristics). Columns two and three give marginal cost estimates under the assumption of a Nash Bertrand game in prices (imperfect competition) for both the RC Logit Models. Note that standard errors are not corrected for the uncertainty in the marginal cost levels.

welfare.

I conclude this section by emphasizing that emissions enter the model through two channels. First, all else equal, consumers dislike vehicles that have higher emissions because they are more costly. There is considerable and significant variation in the degree to which consumers dislike fuel costs. Second, producing vehicles with lower emissions and fuel costs is costly for manufacturers. The first channel matters for all compliance strategies, while the second channel matters to evaluate the cost changes from technology adoption.

6 Welfare effects

In this section I use the estimated model to compare the welfare effects of emission standards with various compliance strategies. I start by describing the solution methods and computation of welfare. After that, I describe the welfare impact of the EU emission standard, the role of the attribute basing, the role of enforcement and the importance of the cost assumptions. These results show the main contribution of the paper, i.e., the design of emission standards has an impact on the abatement choices of firms and the abatement choices matters for the welfare impact of the standard.

Simulation setup The goal of the welfare simulation is to find the welfare effects of the introduction of the EU emission standard, to test the robustness to the economic assumptions and to compare its effects with alternative policy designs. The counterfactual is computed by solving the system of first order conditions (8), (9) and (10) for a binding regulation with a target to reduce sales-weighted emissions to 130 g/km. The goal of the counterfactual is to simulate the optimal changes in prices, technology and gaming in response to the emission standard. To do so I take the vehicle fleet of each firm in 2007 as given and improve the emissions by 6%. This adjusts for the fact that there would have been reductions in emissions even without the regulation.[39] All vehicles and characteristics of this fleet are fixed in the simulation except price, technology τ, and gaming g. Furthermore, in line with the engineering documents, I assume that there are no fixed costs to develop or adopt the technology, so that $C'_\tau = 0$. All cars would have gone through a redesign cycle in the eight years between policy announcement and implementation. I will comment on fixed costs of gaming below. Finally, I solve for the equilibrium so that each firm exactly complies. In reality, this does not need to be the case, as firms may not comply with the standard and pay fines.[40] Note that no firm has incentives to over-comply, as this would involve distorting the optimal choices further from the prepolicy equilibrium, as such the counterfactual singles out

[39] The different models estimated in Section 3 show that there is a 0.7% to 1% annual decrease in emissions in the prepolicy period. To adjust for the fact that the vehicle fleet would have improved without the policy, I endow the 2007 fleet with a 6% reduction in emissions. This 6% equals 8 years (between 2007-2015) of 0.7% technology growth. The hypothetical fleet is the baseline for the welfare computations.

[40] In equilibrium, this will end up to not be a constraint on the solution as all firms prefer compliance over paying fines in the main scenario. Except in the case where I restrict firms to sales mixing, the fines will matter and change the outcome. I discuss this in Appendix 7.

the effect of the policy.

Solving for the system of equations (8), (9) and (10) and the policy target is computationally infeasible. Solving the full system would require a solution for j unknown prices, τ's and g's in each country, as well as a level of λ for each firm. I proceed by making a couple of additional assumptions. First, the first order conditions for gaming (10) are satisfied only when the sunk cost of gaming is positive. This cost is very difficult to determine as it depends on future legal fees, brand reputation and hidden efforts. Therefore, I will use the reduced form results from estimating the trade off relationship to fix the level of gaming. Given the estimates derived in Section 3, I assume that each percentage improvement in fuel consumption includes 70% gaming and 30% actual technology. This means that I solve (8) and (9) with respect to prices and technology while imposing that each 0.3 units of τ imply 0.7 units of g. When discussing the role of gaming and enforcement, I vary this ratio so that either each unit of technology has no gaming with it, or technology is full gaming. I further reduce the number of variables to solve for by assuming that technology is implemented fleet wide for each firm. This reduces the number of equations in (9) from j to the number of firms.

Given these simplifications, I run the following algorithm. Step 1 is to choose a value for the firm level technology τ and shadow cost λ. Given these values, I compute the changes in the marginal costs from technology and the product price distortion implied by the shadow cost λL for each product. Step 2 solves for the Nash equilibrium in prices given the values chosen in Step 1. This step thus solves for (8). Step 3 evaluates how far we are from satisfying (9) and the regulatory target given the guess from Step 1 and the prices from Step 2. I then update the guess and repeat Steps 1-3 until we are sufficiently close to (9) and the policy target. This final iteration gives the solutions for τ, λ, prices and the implied amount of gaming. To update the guess between Step 3 and Step 1, I use the least square nonlinear equation solver provided by Knitro, with bounds on the parameters. The parameter space is bounded because the shadow costs must be positive and the technology improvements must be between 0 and 1.

Given the solution vectors of τ, λ and prices, I compute the changes in outcomes between the initial equilibrium and the new equilibrium. All welfare changes give the total vehicle lifetime changes for one year of new vehicle sales. The direct effects of the regulation will be changes in the consumer surplus, profits and gains from correcting externalities. The consumer surplus is calculated using the log sum formula of Small and Rosen (1981). There is an additional step to account for gaming. When consumers are fooled by gaming, it creates a difference between the decision and experience utility. To compute consumer surplus for non-sophisticated consumers I compute the size of the choice distortion created by the differences in decision and experience utility. The consumer surplus partly comes from reduced fuel expenses. Approximately 60% of these expenses are fuel taxes paid to the government. Depending on whether these taxes are efficient, this part of the consumer gains could be seen as a transfer from the government to consumers and not as a pure welfare change. Changes in profits are obtained from prices, marginal costs and quantities. Note that emission standards do not result in monetary transfers from firms to the government when every firm complies. Appendix 7 show when firms would opt not to comply and pay fines. To compute the changes in CO_2 emissions, I assume a vehicle lifetime of 15 years, a yearly mileage of 14,000km and a discount rate of 6% to capitalize the yearly gains/losses in externalities.[41] The mileage is assumed to be constant, ignoring the possible rebound effects on the intensive margin. To compute the value of the CO_2 reductions, I assume that each ton of CO_2 has an external cost of €28.[42]

Finally, I compute two additional welfare effects of the policy that were not explicit targets of the policy maker. The regulation will change the number of vehicles sold, so that the size of the market changes. Parry, Walls, and Harrington (2007) give an estimate for the total external cost from driving for the US market. These externalities include local pollution, accident risks and congestion, and together, these are estimated to be more important than the CO_2 externality. They report an externality of €12 cents per kilometer. This number is probably not directly applicable to the EU market but at least gives a sense of the relative importance of emissions and other externalities.[43] I will report the gains from shrinking the size of the vehicle market using this €12 cents per kilometer number. Note that this is very optimistic, as the emission standard only targets the sales of new vehicles, not when they drive or how much they drive. When the new vehicle fleet shrinks, it is very uncertain congestion will decrease, as existing vehicles might fill the gap. A second additional welfare effect is related to the behavioral biases of consumers. It has been argued that emission standards are a more effective tool to reduce pollution if consumers undervalue future fuel savings. Using the same data and a similar methodology, in Grigolon, Reynaert, and Verboven (2018) we find that the consumer undervaluation of fuel costs in the EU is at most modest.[44] For the model with endogenous characteristics I find that consumers value fuel costs less and purchase prices more. At these estimates using the annual mileage and interest rate of 6%, I find that consumers only value a 1 euro reduction in future fuel costs at 0.42 cents. The emission standard will result in consumers purchasing vehicles with lower fuel costs and I include the future gains consumers obtain from this. To do this, I change the consumers experience utility from vehicles so that their valuation of net present fuel costs is equal to their valuation of price. The change in the consumer surplus is paternalistic in the sense that I take a stance on what consumer valuation should be, see Allcott (2016).

I will now discuss the welfare effects of the EU emission standard, the role of attribute basing and enforcement, and the importance of the assumptions about costs.

Welfare Effects Table 6 Column I shows the central welfare estimates for the EU emission standard.[45] The simulation assumes that 70% of emission reductions are due to gaming, the

[41] Yearly mileage and vehicle lifetime are chosen to match statistics reported by Eurostat.

[42] This number comes from the Interagency Working Group on the Social Cost of Carbon.

[43] This number is probably an upper bound for the EU since taxes on fuel and driving are, on average, higher than in the US.

[44] Additionally, recent work in the US has found limited to no undervaluation for fuel costs. See Busse, Knittel, and Zettelmeyer (2013), Allcott and Wozny (2014) and Sallee, West, and Fan (2016)

[45] The table presents the simulation outcomes at the estimated parameters and an 80% confidence interval. The C.I. is computed by taking 50 draws from the estimated parameter variance covariance matrix (assuming a joint

standard is the actual attribute-based regulation, consumers are sophisticated so that choices do not respond to gaming and technology increases the marginal costs through the engineering cost function that starts at the estimated level of $\hat{c}'$. Column I solves for the optimal firm strategy so that firms can employ both sales mixing and technology (and the gaming that comes mechanically with it). Column II restricts the firm abatement strategy to technology adoption and Column III restricts it to sales-mixing.

Columns I and II are almost equivalent, showing that firms choose to abate almost fully by technology adoption and gaming.[46] In the baseline estimate of Column I, we see that the emission standard reduces sales by 1% and emissions by 4.9%. The standard did not change the share of small vehicles (defined as the compact and subcompact vehicle classes). If firms would have responded by sales mixing (Column III) we would have seen stark decreases in sales and emissions and a stark increase in the market share for small vehicles. This illustrates already that the abatement strategy chosen by firms is crucial for the market outcomes of the standard.

The second panel of Table 6 shows the welfare effects of the regulation. I find that the EU standard decreased both the consumer surplus and profits by 2.6 billion and 0.6 billion euro, respectively. These losses are incurred because the firms are forced to take up technology beyond the marginal benefit. In turn, this increases the prices for consumers above their willingness to pay for fuel consumption. The CO_2 reduction of 5% is much lower than the policy target because of the gaming. The total emission savings are worth only 0.3 billion, which is much less than the private losses. Dividing private losses by tons of CO_2 I find an implied value of the government for a ton of CO_2 of €2633. This is much higher than the current estimated levels of the social cost of carbon. There are other policies available, such as a gasoline tax, or more general an economy wide carbon tax, that would reduce carbon at a much lower cost. Note that the effects on consumer surplus and profits would have been much more negative with sales mixing, so that the regulation would have been worse for welfare.

However, the policy has two indirect potential benefits that I report in final panel of Table 6. First, a decrease in the fleet of vehicles on the road, reduces other externalities from traffic. Given the assumption of 12 cents per km, the 1% reduction in traffic increases welfare by 2.2 billion. This is a large number and, as explained above, is most likely an absolute upper bound.[47] Second, correcting undervaluation reduces the consumer surplus losses significantly by 1.5 billion. Because consumers undervalue fuel consumption in the RC Logit II model, the reduced future fuel expenses of the new choices will benefit consumers in the future more than they value today. Both these

normal distribution). For each draw of parameters, the supply side is re-estimated and new simulation outcomes are computed. A bootstrap interval is computed using the differences between these 50 outcomes and the outcome at the mean parameters.

[46] The percentage improvements in technology in both scenarios are almost equal with differences of 0.1% points while the shadow costs are very close to zero. On the margin, a very small amount of sales mixing will always be efficient.

[47] Less sales of new cars increases the lifetime of existing old and dirtier cars (see Jacobsen and van Benthem (2015), efficient vehicles might be driven more, congestion does not necessarily reduce with less vehicles and the effects of local pollution depend heavily on where the vehicles are driven.

indirect welfare effects need to be added in full to find positive welfare effects. [48]

The regulation was introduced with the goal of reducing CO_2 emissions. The regulation failed in meeting its target because most reductions in CO_2 happened only on paper and not on the road. However, the regulation reduced private welfare by a substantial amount relative to the emission gains. Other savings, either in other externalities or from correcting undervaluation are both needed to find positive welfare effects from the regulation. Why did firms choose abatement with technology and gaming? I continue to discuss the role of attribute basing and the enforcement.

[48] Appendix Table A5 shows the simulation results with the estimates of RC Logit I. Consumer valuation for fuel costs is estimated to be higher, so that we start the simulation from a steeper point on the convex cost curve. Decreases in sales, emissions, consumer surplus, profits and other externalities are much larger because the compliance costs are higher. There are no savings from correcting undervaluation (in fact there are losses because of estimated overvaluation).

Table 6: Simulation Outcomes

	I	II	III	IV	V	VI	VII
	Opt.	Tech	Sales Mix	Flat	Fooled	Enforce	Engin.
Solve For:	λ, τ	τ	λ	λ, τ	λ, τ	λ, τ	λ, τ
Gaming:	70%	70%	-	70%	70%	0%	0%
Consumer Soph.:	Yes	Yes	Yes	Yes	No	-	-
	Market Size						
Total Sales (%)	-1.08	-1.30	-7.34	-0.98	-0.12	-1.73	0.20
	[-3.11,0.19]	[-3.41,0.02]	[-8.74,-3.67]	[-3.46,0.30]	[-2.95,1.30]	[-5.00,-0.04]	[-33.69,0.74]
Emissions (%)	-4.85	-4.31	-19.16	-11.71	-3.81	-12.60	-5.47
	[-6.53,-3.63]	[-5.98,-3.33]	[-21.66,-12.92]	[-14.70,-7.19]	[-5.98,-2.55]	[-14.77,-10.59]	[-24.48,-4.43]
Share small (%-p.)	0.24	-0.01	12.61	9.01	0.06	0.85	3.33
	[0.09,0.41]	[-0.06,0.06]	[7.41,16.29]	[6.38,14.10]	[-0.04,0.37]	[0.29,1.31]	[-0.05,5.66]
	Direct Welfare Effects (Δ in billion €'s)						
Consumer Surplus	-2.57	-2.51	-19.10	-7.62	-2.62	-4.45	-2.68
	[-6.17,-0.22]	[-6.12,-0.16]	[-20.74,-15.44]	[-13.82,-3.26]	[-6.22,-0.29]	[-8.90,-1.63]	[-104.5,-2.07]
Profits	-0.60	-0.65	-2.58	-0.96	-0.13	-0.95	-0.19
	[-1.71,0.05]	[-1.77,0.03]	[-3.60,-0.12]	[-2.45,-0.23]	[-1.56,0.59]	[-2.54,-0.11]	[-0.42,183.38]
CO2 Value	0.34	0.30	1.32	0.81	0.26	0.87	0.38
	[0.21,0.44]	[0.21,0.41]	[0.62,1.60]	[0.34,1.05]	[0.20,0.40]	[0.57,1.06]	[0.25,1.85]
Total	-2.83	-2.86	-20.36	-7.76	-2.49	-4.53	-2.49
	[-7.45,0.12]	[-7.49,0.12]	[-22.83,-15.01]	[-15.77,-2.77]	[-7.08,0.49]	[-10.29,-0.87]	[-4.48,78.92]
Implied Value CO_2	2633	2962	4564	2952	2911	1730	2113
	[1038,7081]	[1138,8187]	[3328,5447]	[570,5685]	[1020,8844]	[657,3607]	[1944,24149]
	Indirect Welfare Effects (Δ in billion €'s)						
Other Ext.	2.19	2.65	14.94	2.00	0.24	3.52	-0.41
	[-0.20,6.42]	[0.16,7.02]	[5.19,18.13]	[-0.97,7.39]	[-2.52,6.60]	[0.36,10.54]	[-1.60,76.38]
Undervaluation	1.52	1.28	6.75	3.62	1.19	3.89	1.92
	[1.21,2.02]	[1.07,1.79]	[3.89,8.03]	[2.43,4.16]	[0.84,1.82]	[3.14,4.58]	[1.44,6.87]
Total:	0.88	1.07	1.33	-2.14	-1.07	2.87	-0.98
	[0.57,1.53]	[0.90,1.84]	[-5.78,4.71]	[-6.62,1.41]	[-2.04,1.34]	[2.13,4.65]	[-2.99,163.19]

This table gives the welfare effects of policy simulations with 80% C.I. in brackets. Column I solves for the optimal abatement strategy given the following baseline assumptions: engineering cost function starts at $\hat{c}'$, each emission reduction is 30% technology and 70% gaming and consumers are sophisticated (gaming does not affect choice). Column II only allows for technology and gaming. Column III only allows for sales mixing. Column IV introduces a flat standard with a target of 130g of CO_2/km. Column V allows gaming to affect consumer choice. Column VI sets gaming to zero. Column VII implements the engineering cost function without adjusting the intercept.

Attribute Basing In the baseline result, we find that firms choose to abate almost exclusively by lowering the emission rates. The reason for this choice is the attribute basing of the regulation. Table 6 Column IV gives the welfare effects of compliance to a flat regulation. This scenario would have led to more sales mixing and more emission savings. The share of small cars increases by 9% points. Why does a flat standard allow for more sales mixing? The slope in the target has the effect that many low weight products are above the target rather than below. This is illustrated for Fiat in Figure 3. The figure plots all the products in the fleet, scaled by sales, in the emission-vehicle weight space. This shows that many products in the lower left of the figure are below the red flat target but are above the green sloped target. For all these products the policy is an implicit subsidy under a flat regulation but is an implicit tax for the attribute based regulation. This makes sales mixing much more costly because the firm has much fewer products to which it can shift sales.

Why was the EU emission standard attribute-based? Deters (2010) describes the legislation process in detail. He gives the following quote from French president Nicolas Sarkozy clearly favoring a flat regulation:*"There is no legitimate reason to give the buyer of a heavy vehicle a right to more pollution than any other buyer."* While Romano Prodi, Italian Prime Minister, stated the following: *"A steeper value curve would lead to a significant distortion of competition and an illegitimate hardship for the producers of small cars"*. Both these statements contrast nicely with Angela Merkel, Chancellor of Germany, who stated the following: *"The proposed value curve is already a reduction duty far above average for larger cars"*. A steeper target function (the Germans proposed a slope $a = 0.06$ instead of 0.04) would have resulted in lower effort needed from the German firms. The attribute basing was therefore the result of a political agreement between the car producing countries.

In Table 7, I compare the effects of attribute-based and flat regulation. I compare the profits, technologies and shadow costs for different firms averaged per production region, i.e., Asia, France, Germany, Italy and the US. The left panel gives changes of compliance to the ABR when firms choose optimally or when they are restricted to sales mixing. The right panel does the same for the flat target. Restricting firms to sales mixing shows that the shadow costs of sales mixing increases with the slope of the regulation. The mean shadow costs of sales mixing triple from 0.29 to 0.94 when introducing the attribute basing. To see this, compare λ and λ' in II and IV. Fiat, the only Italian firm, would have automatically complied with a flat target, while they have the highest shadow cost under the attribute-based standard. Because sales mixing to the ABR is so costly for Fiat (and also the Asian firms), it will abate emissions mostly by technology and gaming. This has large implications for market equilibrium, because when some firms start improving emission ratings (and consumers benefit from lower fuel costs), other firms face competitive pressure to follow. This equilibrium behavior explains why we see a shift from some technology and some sales mixing in III to almost exclusively technology and gaming in I.

The changes in profits in Table 7 show that French and Italian firms benefit from a flat standard, while all the compliance costs fall on the German firms. This is in line with the strong positions the countries took when bargaining over the regulation. The policy debate in 2007 focused mainly on

Figure 3: Position of the Fiat fleet relative to flat and attribute based target

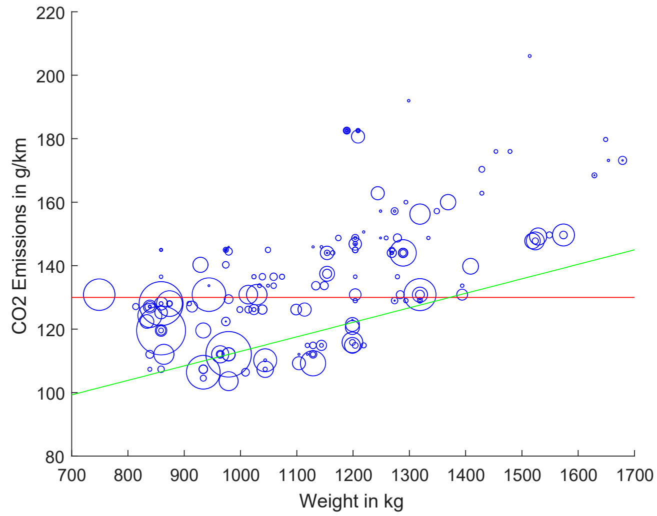

The figure plots every vehicle sold by Fiat in 2007 as it appears in the base data for the counterfactual. Each blue circle is a product, and the circles scale with the quantity sold. The red line is the flat target at 130 grams of CO_2 per km. The green line is the attribute-based target. More sales of products below the target help with compliance.

these distributional issues and not on the effect of the slope on the likelihood of different abatement strategies. This clearly shows the importance of the political economy of the regulation. The ABR was agreed such that all firms would have similar distances from the target and thus would face similar compliance efforts. However, this design made sales mixing more expensive so that the industry needed to reduce the official emission numbers to be able to comply with the regulation.

Enforcement Interestingly, it is again national interests that led to the weak enforcement of emission testing and ultimately led to gaming. In a recent report the European Parliament (Gieseke and Gerbandy (2017)) blamed both the European Commission and the member states for allowing firms to game emission tests. The report states that member states contravened their legal obligation to monitor and enforce defeat devices, while France, Germany and Italy had evidence that emission control systems were not focused on the use of vehicles in on-road conditions. The report states that these countries did not take steps to understand the performance gap between official and on-road emission, thus indicating maladministration. The report also blames member states for under-funding testing facilities (in practice, car makers funded the testing facilities themselves). Even after the Volkswagen scandal in the US most member states did not start immediate and consistent investigations, nor did they adopt an effective and dissuasive penalty system. Additionally, the European Commission failed to oversee the enforcement of member states. In summary, the main car producing countries were aware of the gaming but failed to enforce the regulation. The

Table 7: Profits and Stringency per Producing State

	ABR					Flat				
	I			II		III			IV	
Solve for:	τ, λ			λ		τ, λ'			λ'	
	Δ Prof.	τ	λ	Δ Prof.	λ	Δ Prof.	τ	λ	Δ Prof.	λ
Asia	-212	14	0.09	-1660	1.33	-380	2	0.33	-163	0.41
France	23	4	0.04	1527	0.45	382	0	0.04	1296	0.05
Germany	-635	13	0.08	-1855	0.93	-1287	8	0.33	-3196	0.57
Italy	-84	8	0.07	-682	1.12	28	0	0.00	281	0.00
US	-44	8	0.06	86	0.85	299	3	0.21	578	0.43

The table gives the average shadow costs and changes in profits in millions of euros relative to no policy for German, Italian, French, Asian and US firms from abatement in response to both the Attribute Based Regulation and a Flat Regulation. In all solutions there is no gaming, only actual technology adoption. Solution I solves for equilibrium abatement to the ABR, II allows only for sales mixing. Solution III solves for equilibrium abatement to the flat standard, IV allows only for sales mixing.

overseeing European Commission in its turn failed to follow up on the signals that downstream enforcement was failing.

Columns I, V and VI of Table 6 shed light on the economic consequences of the weak enforcement. The difference between Columns I and V is consumer's awareness of the gaming. In I, the gaming does not affect consumer choice, while it does in V. When consumers are fooled by gaming they perceive cars to have lower fuel consumption, and discover this to be wrong while driving. This causes a choice distortion to consumers and thus, a further reduction in consumer surplus. Additionally, the firms increase prices because products are perceived to be of higher quality. The situation is worse for the environment as there are almost no reductions in sales and emissions now. Overall, this causes the regulation to have a clear negative welfare effect (even when taking into account other externalities and undervaluation). Column VI shows what the welfare effects would have been had the standard been fully enforced. This would have increased private losses in consumer surplus and profits, but would have led to much higher CO_2 and other externality savings. Column VI is the simulation that shows the highest welfare numbers, but the EU failed to attain this. Finally, if abatement would be 100% gaming there would be no changes in the static equilibrium with sophisticated consumers. Car makers would report different emissions to the regulator but prices and fuel costs would remain the same

The evidence presented here shows that the design of the regulation is responsible for the observed outcome. First, the political bargaining between France, Germany and Italy led to an attribute-based standard with a steep slope. This increased the cost of sales mixing and thus, increased the likelihood of compliance by gaming and technology adoption. Next, enforcement failures on the level of the member states then enabled firms to resort to gaming.

Cost function and investment inefficiencies In the main simulation we start from the economic point where the marginal benefits of emission reductions equal marginal costs. The engineering reports however present a cost function that starts with a lower slope than the estimated point concurring with the economic model. If the engineers are correct, this means that firms leave money on the table, i.e., they do not choose to adopt available cost effective technology for which there is willingness to pay.[49] Additionally, economists have argued that there might be market failures in the supply and adoption of technology. Jaffe, Newell, and Stavins (2005) point to spillovers in technology, spillovers in adoption and incomplete information about future returns of the investment as possible market failures. The result of these market failures could be a socially sub-optimal equilibrium with no or too little investment and technology adoption. The regulation gives clear and binding efficiency targets for the whole industry and thus might have succeeded in moving the industry out of a sub-optimal equilibrium by inducing technology adoption. There has been very little work on the empirical validation of these supply side market failures but the framework here allows the testing of the welfare effects of the emission standard when the supply side undervalues technology.[50]

By comparing Column VI and VII in Table 6, we see that supply side failures in technology adoption are not necessarily bad for welfare when externalities are at play. Because of the cheap technology, the regulation makes firms reduce fuel consumption for less than the consumers' willingness to pay. This means that vehicles become cheaper, as consumers receive better characteristics for a price below their willingness to pay. As such, the market size will increase rather than decrease. The regulation now has a rebound effect on the extensive margin.[51] The cheap technology reduces consumer and profit losses but also wipes out the savings in other externalities so that, surprisingly, more technology adoption is not necessarily better for welfare. Note that the regulation pushes firms beyond what could be explained by market failures, a policy that would only require cost efficient technology would be less stringent. In general, I believe this scenario not to be very credible. In line with Anderson and Sallee (2011) we expect car makers to abate emissions with the least costly compliance strategies. If technology is so cheap and would have costed only €190 million in variable profits, then why would firms have resorted to gaming?

In Appendix Table A5 Column II I also show the results with the estimated cost function specified in (13). The results are very similar to the main scenario but the estimated cost function is less convex than the engineering cost function. This reduces the consumer and profit losses. The question is, to what extent historical cost increases identify future cost changes.[52]

[49] Notice that the main results already account for the welfare effects of potential undervaluation by consumers; here, I consider the suboptimal technology adoption of firms.

[50] Recent work, such as Hashmi and Van Biesebroeck (2016) and Aghion, Dechezlepretre, Hemous, Martin, and Reenen (2016), has looked at R&D patterns in the automobile industry through patents.

[51] See Gillingham, Kotchen, Rapson and Wagner (2013) for an overview of the possible sources of rebound effects. A second rebound effect that might be expected is an increase in vehicle usage, a rebound effect on the intensive margin. A further rebound effect could come from the use of savings on vehicle expenses on other energy intensive activities. This is known as the indirect rebound effect. Lastly, a decrease in the demand for fuels might lower the price of oil, thus causing further shocks in the economy,as macro-economic rebound effect. Here, I only focus on the rebound effect on the extensive margin, the reported emission savings are thus an upper bound on the total savings.

[52] Additionally, the slope of the estimated cost function with respect to fuel costs is highly dependent on the functional form. In many specifications, the regression of the predicted marginal costs on fuel costs results in a positive sign, such that the fuel economy lowers the marginal costs. These issues could be due to multicollinearity

Sunk Costs The profit changes reported in Table 6 are changes in variable profits. Both technology adoption and gaming potentially have sunk costs. I explain the underlying sources of sunk costs and then I compute an upper bound on these costs by computing deviations in variable profits from optimal strategies.

The fixed costs of technology adoption contain the development of technology, implementation of technology and the redesign of the vehicle. Engineers state that the technology to attain the emission reduction was available at the time of the policy. The policy allows for a period of 8 years between announcement and enforcement. Vehicles typically go through faster redesign cycles, see Blonigen, Knittel, and Soderbery (2017). This means that the additional sunk costs of redesigning are also likely to be small. Because of the availability of technology and no additional redesign cycles we expect little additional sunk costs of technology except for the adoption itself.

Gaming has expected sunk costs. The defeat devices used to game the emission tests have to be purchased or designed. Then there is also the risk of facing liability for noncompliance and class actions from consumers and shareholders. A complicating factor is that similar defeat devices are used to avoid pollution standards so that the legal cases are both about misrepresenting pollutants and fuel economy. It is thus unclear if we should attribute all of these costs to the emission standard.

Given the estimated static model, I can compute an implied upper bound of sunk costs from the changes in variable profits when firms deviate from the optimal strategy.[53] What is a firm willing to pay in sunk costs to not comply with sales mixing? The upper bound of these sunk costs can be computed by comparing variable profits in the optimal equilibrium with variable profits in an equilibrium where the firm is restricted to sales mixing, while all other firms continue playing optimal strategies. The variable profits of the deviating firm will be lower and this provides information on the costs this firm would be willing to sink to not play the sales mixing strategy. The difference in variable profits gives the cost of deviating for a single year of sales, but the regulation is binding for several years. We thus need to scale and discount the difference in variable profits by the expected investment horizon of the firm. Because the 2015 emission standard binds until 2021 (when it will be replaced by an even more stringent regulation), I foresee a horizon of 6 years.[54] Computing this upper bound requires solving an additional equilibrium for each firm and is thus costly. I find that the sum of the upper bounds for all firms equals €70 billion. The lowest upper bound equals €500 million for Fiat and the highest is €20 billion for Volkswagen. The results for all firms and for smaller deviations are shown in Appendix Table A6. This shows that firms are willing to sink high amounts to be able to comply with gaming and technology adoption. This also means that the modest positive welfare numbers presented in Table 6 are not sufficient to claim that the regulation had any positive effect. There are currently multiple legal cases against car makers, the

as well as changes in mark ups that are correlated with fuel costs but not captured in the first order condition, see Langer and Miller (2013).

[53] In this setting, I can compute deviations from the single equilibrium with the optimal compliance strategies. From observing multiple compliance choices in multiple markets, one could use a moment inequality estimator to estimate the sunk costs, see Pakes, Porter, Ho, and Ishii (2015). Here, I can only compute a single deviation profit for each firm implied by the simultaneous game.

[54] I use a 6% discount rate to compute the net present value of the stream of variable profit losses.

efforts and costs related to these cases are an additional welfare loss so that the regulation has at best zero effect on welfare.[55]

7 Conclusion

This paper has evaluated the response to a recent EU wide emission standard. I find that between 2007 and 2011, sales-weighted emissions from new vehicle sales have decreased by 14%. Decomposing this decrease, I find that two-thirds of the decrease in emission is attributable to firms gaming the emission test procedure. One-third of the decrease stems from actual technology adoption. A structural model of demand and supply, allowing for endogenous abatement strategy choices, revealed that the overall effect of the regulation has been negative for consumers and producers and did not save large amounts of carbon. For the regulation to be welfare-improving, we need to optimistically scale the reduction in the size of the new vehicle fleet with a cost of saving other externalities, such as accidents and local pollution, while this was not the goal of the regulation. The projected carbon savings did not materialize because of noncompliance and gaming on emission tests. The reasons that firms chose this abatement strategy are the attribute basing and the lack of enforcement. Both of the reasons are a product of the political environment in the EU.

This exercise provides important lessons for emission standards as a policy tool. First, the compliance strategy that firms use matters for the welfare and environmental outcomes of the policy. The design of the policy has potentially large impacts on which strategies are employed. I showed that the attribute basing made sales mixing so costly that some firms had almost no choice but to decrease the emission ratings of all their vehicles. Second, emission standards are designed, implemented and enforced in a world where politics matter. I find that even in a developed world setting, such as the EU, this political reality matters for the outcome, as both the attribute basing and lack of enforcement are the result of the political economy. Third, technology adoption is a crucial mechanism of compliance and needs to be considered when evaluating emission standards. This is very difficult because the evaluation depends on assumptions about the cost curve of technology adoption. Both consumers and firms might undervalue fuel consumption reductions. It is crucial to understand to what extent this is happening, especially on the supply side, as we lack evidence of market failures so that our economic models are at odds with engineering estimates.

Overall, I want to stress that emission standards are a risky and unpredictable policy instrument when the goal is to reduce carbon emissions. Standards rely on emission tests that could be potentially gamed when compliance costs become very high. It is uncertain that emissions actually reduce because of the policy. If the engineers are correct and there is cost effective technology available, then standards might increase the size of the market rather than decrease it, causing an increase in emissions. A correct fuel tax in combination with subsidies to adopt technology would fare better, as the first instrument corrects for the externality and the second for potential supply

[55]In the fallout of the VW diesel scandal the EU parliament listed dozens of ongoing lawsuits (see: European Parliament Briefing PE 583.793). In recent years, several other firms have become defendants in claims as it became clear VW was not the only firm to have gamed the EU emission tests.

side failures. Another option would be to bring the transport sector into the emission trading system so that the vehicle market contributes to the prices of carbon permits and abatement costs equalize across sectors. Since all EU countries have very high fuel taxes that cover more than the carbon externality, it is difficult to understand why the EU chose this policy as a mechanism to reduce carbon emissions. One of the reasons is that the EU has limited fiscal authority to further raise taxes. The EU parliament recently approved ever more stringent emission standards for 2021 to 2030. These standards are so stringent that classic combustion engines will not be capable of reaching the target, which implies that a shift to alternative fuels is imposed on the next decade.

The numbers derived in this paper are obtained under various assumptions and one should keep in mind the limitations of the model and the data. First, I focus only on the sales of new vehicles and assume there will be no effects on prices and vehicle lifetimes in the used car market. It would be interesting to study if there are different effects from gaming, technology adoption and sales mixing on the used car market. Second, all welfare numbers are obtained ignoring possible rebound effects on driving behavior. Third, I did not include dynamics in the analysis. Important fixed costs on the firm side might have effects on the market structure and rapid technology adoption might result in consumers strategically timing purchases. Even without these complications, the counterfactual outcomes show us that the welfare effects from emission standards are far from obvious and that the design of the regulation matters for which abatement strategies will be chosen.

References

Aghion, Philippe, Antoine Dechezlepretre, David Hemous, Ralf Martin, and John Van Reenen. 2016. "Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry." Journal of Political Economy 124 (1):1–51. URL https://doi.org/10.1086/684581.

Allcott, Hunt. 2016. "Paternalism and Energy Efficiency: An Overview." Annual Review of Economics 8 (1):145–176. URL https://doi.org/10.1146/annurev-economics-080315-015255.

Allcott, Hunt and Nathan Wozny. 2014. "Gasoline Prices, Fuel Economy, and the Energy Paradox." Review of Economics and Statistics 96 (10):779 – 795.

Anderson, Michael and Maximilian Auffhammer. 2014. "Pounds that Kill: The External Costs of Vehicle Weight." Review of Economic Studies 81 (2):535–571.

Anderson, Soren T., Ryan Kellogg, and James M. Sallee. 2013. "What Do Consumers Believe About Future Gasoline Prices?" Journal of Environmental Economics and Management 66 (3):383–403.

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. URL http://www.aeaweb.org/articles.php?doi=10.1257/aer.101.4.1375.

———. 2016. "Designing Policies to Make Cars Greener: A Review of the Literature." Annual Review of Resource Economics 8:157–180. URL http://www.annualreviews.org/doi/abs/10.1146/annurev-resource-100815-095220.

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. URL http://www.sciencedirect.com/science/article/pii/S0095069605000550.

Bento, Antonio, Kenneth Gillingham, and Kevin Roth. 2017. "The Effect of Fuel Economy Standards on Vehicle Weight Dispersion and Accident Fatalities." Working Paper 23340, National Bureau of Economic Research.

Bento, Antonio M., Kenneth Gillingham, Mark R. Jacobsen, Christopher R. Knittel, Benjamin Leard, Joshua Linn, Virginia McConnell, David Rapson, James M. Sallee, Arthur A. van Benthem, and Kate S. Whitefoot. 2018. "Flawed analyses of U.S. auto fuel economy standards." Science 362 (6419):1119–1121. URL http://science.sciencemag.org/content/362/6419/1119.

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):667–99. URL http://ideas.repec.org/a/aea/aecrev/v99y2009i3p667-99.html.

Berry, Steven T., James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." Econometrica 63 (4):841–890.

———. 1999. "Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy." American Economic Review 89 (3):400–430.

Blonigen, Bruce A., Christopher R. Knittel, and Anson Soderbery. 2017. "Keeping it fresh: Strategic product redesigns and welfare." International Journal of Industrial Organization 53:170 – 214. URL http://www.sciencedirect.com/science/article/pii/S0167718716301813.

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):220–56. URL http://ideas.repec.org/a/aea/aecrev/v103y2013i1p220-56.html.

Crawford, Gregory S., Oleksandr Shcherbakov, and Matthew Shum. 2019. "Quality Overprovision in Cable Television Markets." American Economic Review 109 (3):956–995. URL https://ideas.repec.org/a/aea/aecrev/v109y2019i3p956-95.html.

Deters, Henning. 2010. "Legislating on Car Emissions. What Drives Standards in EU Environmental Policy." Tech. rep.

Durrmeyer, Isis. 2018. "Winners and Losers: The Distributional Effects of the French Feebate on the Automobile Market." TSE Working Papers 18-950, Toulouse School of Economics (TSE). URL https://ideas.repec.org/p/tse/wpaper/32928.html.

Durrmeyer, Isis and Mario Samano. 2018. "To Rebate or Not to Rebate: Fuel Economy Standards Versus Feebates." The Economic Journal 128 (616):3076–3116.

Fan, Ying. 2013. "Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market." American Economic Review 103 (5):1598–1628. URL http://www.aeaweb.org/articles.php?doi=10.1257/aer.103.5.1598.

Gieseke, Jens and Gerben-Jan Gerbandy. 2017. "Report on the Inquiry into emission measurements in the automotive sector." Tech. Rep. 2016/2215(INI), European Parliament. URL http://www.europarl.europa.eu/sides/getDoc.do?pubRef=-//EP//TEXT+REPORT+A8-2017-0049+0+DOC+XML+V0//EN.

Gillingham, Kenneth, Matthew Kotchen, David Rapson, and Gernot Wagner. 2013. "The Rebound Effect is Over-played." Nature 493:475–476.

Goldberg, Pinelopi Koujianou. 1998. "The Effects of the Corporate Average Fuel Efficiency Standards in the US." Journal of Industrial Economics 46 (1):1–33.

Grigolon, Laura, Mathias Reynaert, and Frank Verboven. 2018. "Consumer Valuation of Fuel Costs and Tax Policy: Evidence from the European Car Market." American Economic Journal: Economic Policy 10 (3):193–225. URL http://www.aeaweb.org/articles?id=10.1257/pol.20160078.

Hashmi, Aamir and Johannes Van Biesebroeck. 2016. "The Relationship between Market Structure and Innovation in Industry Equilibrium: A Case Study of the Global Automobile Industry." The Review of Economics and Statistics 98 (1):192–208. URL http://EconPapers.repec.org/RePEc:tpr:restat:v:98:y:2016:i:1:p:192-208.

Heiss, Florian and Viktor Winschel. 2008. "Likelihood approximation by numerical integration on sparse grids." Journal of Econometrics 144 (1):62 – 80. URL http://www.sciencedirect.com/science/article/pii/S0304407607002552.

Holland, Stephen P., Jonathan E. Hughes, and Christopher R. Knittel. 2009. "Greenhouse Gas Reductions under Low Carbon Fuel Standards?" American Economic Journal: Economic Policy 1 (1):106–46. URL http://www.aeaweb.org/articles.php?doi=10.1257/pol.1.1.106.

Huse, Cristian and Claudio Lucinda. 2014. "The market impact and the cost of environmental policy: evidence from the Swedish green car rebate." The Economic Journal 124 (578):F393–F419.

International Council On Clean Transportation,. 2014. "Global passenger vehicle standards." http://www.theicct.org/info-tools/global-passenger-vehicle-standards .

Ito, Koichiro and James M. Sallee. 2018. "The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel Economy Standards." The Review of Economics and Statistics 100 (2):319–336. URL https://doi.org/10.1162/REST_a_00704.

Jacobsen, Mark R. 2013. "Evaluating US Fuel Economy Standards in a Model with Producer and Household Heterogeneity." American Economic Journal: Economic Policy 5 (2):148–87. URL http://www.aeaweb.org/articles.php?doi=10.1257/pol.5.2.148.

Jacobsen, Mark R. and Arthur van Benthem. 2015. "Vehicle Scrappage and Gasoline Policy." American Economic Review 105:1312–1338.

Jaffe, Adam B., Richard G. Newell, and Robert N. Stavins. 2005. "A tale of two market failures: Technology and environmental policy." Ecological Economics 54 (2-3):164–174.

Ketz, Philipp. Forthcoming. "On Asymptotic Size Distortions in the Random Coefficients Logit Model." Journal of Econometrics .

Klier, Thomas and Joshua Linn. 2012. "New-vehicle characteristics and the cost of the Corporate Average Fuel Economy standard." The RAND Journal of Economics 43 (1):186–213. URL http://dx.doi.org/10.1111/j.1756-2171.2012.00162.x.

———. 2016. "The effect of vehicle fuel economy standards on technology adoption." Journal of Public Economics 133:41 – 63. URL http://www.sciencedirect.com/science/article/pii/S0047272715002030.

Knittel, Christopher R. 2011. "Automobiles on Steroids: Product Attribute Trade-Offs and Technological Progress in the Automobile Sector." American Economic Review 101 (7):3368–99. URL http://ideas.repec.org/a/aea/aecrev/v101y2011i7p3368-99.html.

Langer, Ashley and Nathan Miller. 2013. "Automakers' Short-Run Responses to Changing Gasoline Prices." Review of Economics and Statistics 95 (4):1198–1211.

Nevo, Aviv. 2001. "Measuring Market Power in the Ready-to-Eat Cereal Industry." Econometrica 69 (2):307–42.

Pakes, A., J. Porter, Kate Ho, and Joy Ishii. 2015. "Moment Inequalities and Their Application." Econometrica 83 (1):315–334.

Parry, Ian W. H., Margaret Walls, and Winston Harrington. 2007. "Automobile Externalities and Policies." Journal of Economic Literature 45 (2):373–399. URL http://www.aeaweb.org/articles.php?doi=10.1257/jel.45.2.373.

Regulation (EC) No. 443/2009, . 2009. "Setting emission performance standards for new passenger cars as part of the Community's integrated approach to reduce CO2 emissions from light-duty vehicles." Official Journal of the European Union L 140/1.

Reynaert, Mathias and James M. Sallee. 2019. ""Who Benefits When Firms Game Corrective Policies?" Working Paper 22911, Toulouse School of Economics. URL https://docs.google.com/viewer?a=v&pid=sites&srcid=ZGVmYXVsdGRvbWFpbnxtYXRoaWFzcmV5bmFlcnR8Z3g6NGUyMTZkNDBiN2NkNzdlOQ.

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 http://www.sciencedirect.com/science/article/pii/S0304407613002649.

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. URL http://www.sciencedirect.com/science/article/pii/S0047272716000049.

Small, Kenneth A. and Harvey S. Rosen. 1981. "Applied Welfare Economics with Discrete Choice Models." Econometrica 49 (1):105–130. URL http://www.jstor.org/stable/1911129.

TNO. 2011. "Support for the revision of Regulation (EC) No 443/2009 on CO_2 emissions from cars." URL 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.

Whitefoot, K. S. and S. J. Skerlos. 2012. "Design Incentives to Increase Vehicle Size Created from the U.S. Footprint-based Fuel Economy Standards." Energy Policy 41 (1):402–411.

Whitefoot, Kate S., Meredith L. Fowlie, and Steven J. Skerlos. 2017. "Compliance by Design: Influence of Acceleration Trade-offs on CO2 Emissions and Costs of Fuel Economy and Greenhouse Gas Regulations." Environmental Science & Technology 51 (18):10307–10315. URL https://doi.org/10.1021/acs.est.7b03743. PMID: 28825797.

Appendix For Online Publication Only

Details on Data Selection

I focus the analysis on the largest EU firms that sell more than 50 000 vehicles in each year of the sample. These are as follows: BMW, Daimler, Fiat, Ford, GM, PSA, Renault and Volkswagen. I consider the largest Asian manufacturers as being one firm in the model. This firm includes the following: Honda, Hyundai, Mazda, Mitsubishi, Nissan, Suzuki and Toyota. The following firms are not considered in the analysis: Alpina, Aston Martin, Brilliance Auto, Chana, DR Motor, Geely Group, Great Wall, Isuzu, Jensen, Jiangling, Lada, Mahindra & Mahindra, MG Rover, Morgan, Perodua, Porsche, Proton, SAIC, Santana, Spyker, Ssangyin, Subaru, Tata, TVR, Venturi and Wiesmann. Daimler and Chrysler merged during the sample period, and I treat them as one firm in the whole sample.

For the included firms I focus on the most popular brands. I drop the following brands which mostly include luxurious sports cars and temporary owned brands: Abarth, Bentley, Buick, Cadillic, Corvette, Daimler, Dodge, Ferrari, Galloper, Hummer, Infiniti, Innocenti, Iveco, Jaguar, Lamborghini, Land Rover, Lincoln, Maserati, Maybach, Pontiac, Rolls-Royce and Tata. In total, the firms and brands that are not included account for 3.5% of the sales.

Additionally, to reduce the number of observations I select only the top 50% highest selling models which are a combination of a Brand/Model/Body indicator, e.g. "Volkswagen Golf Hatchback". Of the top 50% most popular models, I select the engine variants that are sold at least 20 times. Because of this selection, which is necessary to make the number of market share equations tractable, I lose another 14% of sales such that the final data set includes 81.5% of the total reported sales. I lose another 3% of total reported sales due to missing values and unrealistic outliers in the characteristics.

The definition of the vehicle weight changes throughout the sample from the curb weight before 2010 to the gross vehicle weight in the years 2010 and 2011. I transform the gross vehicle weight to the curb weight by matching vehicles that are identical in all characteristics between 2009 and 2010. I regress curb weight on gross vehicle weight, doors and displacement and use the predicted value of that regression to obtain the curb weight in 2010 and 2011. The R^2 of that regression is 0.95. The curb weight is approximately 72% lower than gross vehicle weight. The observed and imputed curb weight are then used to compute each vehicle's compliance with the regulation.

Details on Excluded Instruments

For the specification with endogenous prices the following 13 excluded instruments are used in the first stage:

- Sum of the characteristics of fuel consumption, horsepower, weight, footprint, height of all other products sold by the same firm in the market (5 instruments);
- Sum of the characteristics of fuel consumption, horsepower, weight, footprint, height of all other products in the market (5 instruments);
- Number of products sold by the same firm in the market (1 instrument);
- Number of products in the market (1 instrument); and
- Log of the labor cost in the country of production of the vehicle (1 instrument).

For the specification with endogenous prices, fuel costs, horsepower and weight, the following 17 excluded instruments are used in the first stage:

- Sum of the characteristics of the footprint, height of all other products sold by the same firm in the market (2 instruments);
- Sum of the characteristics of the footprint, height of all other products in the market (2 instruments);
- Number of products sold by the same firm in the market (1 instrument);
- Number of products in the market (1 instrument);
- Log of the labor cost in the country of production of the vehicle (1 instrument);
- Production share of the vehicle model in Africa, Asia, East Europe, North America, and South America (5 instruments);
- Weighted sum of the average size of all vehicles produced in each region, weights equal the production share of the vehicle model (1 instrument);
- Weighted sum of the average size of all vehicles produced in each region, weights equal the production share of the vehicle brand (1 instrument);
- Sum of the characteristics of the footprint, height of vehicles of different vehicle segment produced on the same platform (2 instruments); and
- Fuel consumption projected on all included and excluded instruments interacted with fuel prices (1 instrument).

The role of Fines

As explained in Section 2, the fines are given by €5 per unit sold for the first excess g/km and increase to €95 per unit above 134 g/km. These are the pure monetary fines; it could be that noncompliance with the regulation brings other reputation costs. Though, when firms choose to game these reputation costs might not matter anymore.

The fines are an increasing schedule, so that minor deviations are punished lightly at €5 per vehicle sold and a deviation of more than 4 grams is punished at €95 per vehicle. In between fines are €15 and €25 for gram 2 and 3 for noncompliance. I consider the smaller fines as minor punishments for unexpected changes in the fleet averages and the fine of €95 per vehicle as the punishment for actual noncompliance. This fine will matter for the abatement strategies.

In principle, the fines give an upper bound for the Lagrangian multipliers λ. If the per unit shadow cost of the regulation becomes higher than the fine, then a firm would prefer to pay the fine above further price distortions. In practice, I find that this matters only in the scenario Flat reported in Column IV of Table 6. In all other solutions the equilibrium value of λ is far below this upper-bound. The framework can accommodate fines flexibly however. As I solve the model with bounds on λ and τ, I can replace the upper-bound of λ from infinity to the level of the fine. Column III of Table A5 presents the results where the firms pay fines. Only Volkswagen ends up in an equilibrium where they deviate from the emission standard and they pay 400 million euro. This is a profit loss, but the profit loss flows to the state, so it is not necessarily a welfare loss. Otherwise, the equilibrium outcomes are very comparable.

The solution presented in Table A5 also bounds technology. When firms hit the upper-bound on λ, the algorithm will resort to increasing τ to attain compliance. Therefore, I use the solution of τ when there is no upper-bound on λ as an upper-bound in this algorithm. This is slightly over-restrictive as somewhat more technology will have a lower cost than the fine. The τ also has an implicit per vehicle cost that should be lower than the fine. This cost depends on the equilibrium conditions and is thus harder to solve for (I could, in principle, restate the variable to solve for the per vehicle technology costs rather than percentage reductions).

Additional Figures and Tables

Table A1: Robustness for Table 2 and Table 3

	Model 6	Model 7	Model 8	Model 9	Model 10	Model 11
ln(Hp)	0.43	0.21	0.08	0.13	0.22	0.26
	(0.58)	(0.02)	(0.03)	(0.02)	(0.02)	(0.05)
ln(Weight)	-1.34	0.68	0.75	0.65	0.67	0.55
	(2.05)	(0.06)	(0.05)	(0.03)	(0.05)	(0.03)
ln(Footprint)	0.35	-0.15	-0.28	-0.40	-0.26	-0.33
	(3.02)	(0.10)	(0.09)	(0.10)	(0.08)	(0.10)
ln(Height)	-17.28	-0.07	0.01	0.04	-0.04	-0.06
	(5.90)	(0.11)	(0.10)	(0.15)	(0.13)	(0.12)
Diesel	-0.20	-0.20	-0.20	-0.21	-0.20	-0.65
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.15)
Year F.E.?	Yes	Yes	Yes	Yes	Yes	Yes
YearXFirm?						
Car Name F.E.?	Yes	Yes	Yes	Yes	Yes	Yes
Observations	14,444	3,441	14,444	14,444	14,444	14,444
R^2	0.87	0.81	0.86	0.86	0.86	0.86
	Difference in Technology Growth 2011-2007 and 2007-1998					
Difference	2.6	2.1	3.3	2.2	2.3	2.5
	0.5	0.5	0.7	0.7	0.5	0.5

The table presents the robustness of the findings in Table 2 and Table 3. Each Model estimates equation (2) and the Table presents the first order terms of trade off parameters and the difference in technological change between 2011-2007 and 2007-1998. Standard errors are robust and clustered per firm. Standard error for the difference is computed using the Delta method. Model 6 changes the functional form from Cob Douglas to Translog so that higher order terms in attributes are included. Model 7 keeps only the first appearance of each vehicle. Model 8 allows the trade off parameters to change over time. Model 9 weighs observations by sales. Model 10 includes the marginal costs as estimated in the structural model as a control. Model 11 interacts characteristics with fuel type.

Table A2: Technological Progress Estimates per Firm for Model 3 and Model 5

	BMW		Daimler		Fiat		Ford		GM		PSA		Renault		VW		Asian	
	Off.	Road	Off.	Road	Off.	Road	Off.	Road	Off.	Road	Off.	Road	Off.	Road	Off.	Road	Off.	Road
1999	0,0	-2,0	-3,9	-1,2	-4,2	1,1	-7,6	-2,3	0,0	-1,1	-1,3	-1,2	-1,2	-2,2	1,3	0,0	3,7	-1,1
2000	4,1	-1,0	3,9	-1,2	-2,7	0,0	7,6	-2,2	3,9	-2,2	1,3	-2,3	0,0	-2,1	2,6	-1,1	1,3	-1,1
2001	-3,1	-1,9	-2,6	-1,1	-2,6	-3,2	-3,8	-1,1	-1,3	-2,1	-5,2	-2,2	-2,4	-1,1	0,0	-1,1	-1,3	-1,1
2002	0,0	0,0	-1,3	-1,1	0,0	-1,1	-1,3	0,0	0,0	0,0	-2,5	-1,1	-2,4	-1,0	-1,3	-1,1	-3,7	-2,2
2003	-1,0	-0,9	-1,3	-1,1	-2,6	0,0	-2,5	0,0	-2,6	0,0	-1,3	0,0	-2,3	-1,0	-2,6	0,0	-1,2	-1,1
2004	1,0	0,0	-2,5	0,0	-1,3	-1,0	-1,2	-1,1	-3,8	0,0	-6,1	-2,2	-3,4	-2,0	0,0	-1,1	-1,2	-1,1
2005	-1,0	0,0	-3,7	-1,1	-1,3	0,0	-3,6	-1,1	-2,4	-1,0	-2,3	0,0	0,0	0,0	-1,3	0,0	-1,2	2,1
2006	-2,0	0,0	-1,2	-1,1	-3,7	0,0	-1,2	1,1	-1,2	0,0	-2,3	-1,1	-1,1	0,0	-1,3	-1,1	-2,3	-1,1
2007	-9,3	-6,4	0,0	0,0	-2,4	0,0	1,2	0,0	-1,2	1,0	-2,3	2,1	-1,1	0,0	-1,2	0,0	-1,2	-1,1
2008	-5,2	-2,6	-3,5	-1,1	-3,5	1,0	-2,3	0,0	-2,3	-2,1	-2,2	-1,1	0,0	0,0	-6,0	-2,1	-3,4	0,0
2009	-0,8	-0,8	-6,8	-2,1	-3,4	-1,0	-1,1	1,1	-3,4	0,0	-2,2	-2,1	-4,3	-1,0	-5,7	-2,1	-4,4	-4,1
2010	-1,6	-0,8	-2,2	1,1	-6,6	-3,1	-6,6	-2,2	-7,6	-3,0	-3,2	-4,1	-2,1	0,0	-7,5	-3,1	-4,2	-1,0
2011	-0,8	0,0	-7,3	-7,3	-6,2	-1,0	-7,3	-7,3	-5,1	0,0	-4,1	-1,0	-3,0	0,0	-4,1	-1,0	-3,0	0,0
Difference in Technology Growth 2011-2007 and 2007-1998																		
Difference	0,9	-0,3	3,5	1,5	2,6	0,5	3,0	1,3	3,7	0,7	0,5	1,2	0,8	-0,8	5,4	1,5	3,0	0,4

The table gives the estimated firm specific yearly change of technology in the CO_2 production function as derived from the year fixed effects in (2) for Model 3 and Model 5 in Table 2. The first column for every firm reports percentage technological change in official emission ratings, and the second column reports the change in the on-road emission ratings. The shaded areas are the years after the policy announcement.

Table A3: First Stage Estimates

	(1) Price/Income	(2) Price/Income	(3) Euro per Km	(4) Horsepower	(5) Weight
Log Labor Costs	0.191***	0.159***	-0.394**	-0.0734*	-0.0236*
	(0.0252)	(0.0269)	(0.120)	(0.0316)	(0.0105)
Sum of own Fuel Consumption	-2.440***				
	(0.487)				
Sum of own Horsepower	-1.476				
	(2.408)				
Sum of own Weight	1.024**				
	(0.384)				
Sum of own Footprint	3.215***	3.602***	-4.856	-0.967	-1.303***
	(0.750)	(0.596)	(2.779)	(0.756)	(0.234)
Sum of own Height	1.944***	1.853***	3.418	0.646	0.561***
	(0.497)	(0.412)	(1.924)	(0.525)	(0.162)
Sum of own Products	-4.659***	-5.330***	-1.473	-0.156	0.144
	(0.876)	(0.863)	(4.036)	(1.101)	(0.339)
Sum of other Horsepower	-33.81***				
	(4.241)				
Sum of other Weight	-0.701				
	(0.572)				
Sum of other Footprint	9.614***	-2.211*	2.147	0.280	-1.054**
	(1.522)	(0.920)	(4.244)	(1.145)	(0.361)
Sum of other Height	3.818***	2.917***	3.444	0.602	-0.0873
	(0.854)	(0.831)	(3.834)	(1.035)	(0.326)
Sum of other Products	-8.803***	-2.649*	-6.613	-0.973	0.944
	(1.480)	(1.273)	(5.910)	(1.603)	(0.500)
Gasoline Price by proj. Li		-0.00975***	1.013***	0.00254	-0.00209*
		(0.00209)	(0.00978)	(0.00266)	(0.000822)
Production Share Africa		0.137*	0.247	-0.0252	0.0895***
		(0.0598)	(0.273)	(0.0730)	(0.0234)
Production Share Asia		-0.0423**	-0.111	-0.00323	0.0104
		(0.0160)	(0.0650)	(0.0162)	(0.00622)
Production Share East Europe		0.0453	-0.700*	-0.289***	-0.0154
		(0.0624)	(0.285)	(0.0762)	(0.0245)
Production Share North America		-0.270**	1.051*	-0.0481	-0.0877*
		(0.0958)	(0.438)	(0.118)	(0.0375)
Production Share South America		0.149***	-0.303	-0.112*	-0.0302
		(0.0442)	(0.203)	(0.0547)	(0.0173)
Brand Prod. shares by Size		-0.0345**	0.162***	0.0495***	0.0245***
		(0.0118)	(0.0474)	(0.0117)	(0.00458)
Model Prod. shares by Size		0.120***	-0.189	0.0332	0.0327*
		(0.0326)	(0.150)	(0.0406)	(0.0128)
Plant Height Other Segment		1.069	-15.06	-6.392**	-1.424
		(1.953)	(9.059)	(2.452)	(0.766)
Plant Footprint Other Segment		-3.810	26.97	10.98*	2.666
		(3.720)	(17.27)	(4.677)	(1.460)
SW F Stat	67.35	5.82	21.20	5.75	11.20
# End. Vars	1	4	4	4	4
# Excl. Instr.	13	17	17	17	17
Observations	28775	28775	28775	28775	28775

The table gives the first stage estimates for the specification with endogenous prices (1) and the specification with endogenous prices, fuel costs, horsepower and weight (2-4). The coefficients and robust standard errors for all excluded instruments are reported (the included instrument coefficients are not reported). The Sanderson-Windmeijer multivariate F test of excluded instruments is reported for every endogenous variable, this statistic equals the standard F-test of excluded variables with a single endogenous variable in (1).

Table A4: Model Fit

	Within Sample				Out of Sample	
	True	Pred. I	Pred. II	Pred. III	True	Pred. I
Emission	147	142	155	139	126	125
Weight	1.27	1.19	1.33	1.17	1.28	1.24
Horsepower	0.78	0.66	0.84	0.66	0.80	0.71
Footprint	7.23	7.01	7.43	6.95	7.39	7.31
Price/Income	0.71	0.56	0.76	0.56	0.69	0.58
Diesel	0.56	0.50	0.53	0.49	0.56	0.49

The table presents sales-weighted averages of characteristics within the estimation sample (for year 2007) and out of sample (for year 2011). The prediction columns present sales-weighted measures based on predicted sales rather than observed sales. Prediction I predicts sales using the estimated utility parameters without fixed effects and demand unobservable. Prediction II predicts sales using the estimated utility parameters with fixed effects but without demand unobservable. Prediction III predicts sales using the estimated utility parameters with demand unobservable but without fixed effects.

Table A5: Simulation Outcomes

	I	II	III
	RC Logit I	Estimated Tech	Flat with Fines
Solve for:	λ, τ	λ, τ	λ, τ
Gaming:	70%	70%	70%
Consumer Soph.:	1	1	1
	Market Size		
Total Sales	-14,78	-0,14	-0,11
Emissions	-19,45	-5,69	-10,66
Share small	2.27	3.33	9.26
	Direct Welfare Effects (Δ in billion €'s)		
Consumer Surplus	-29,67	-3,04	-6,19
Profits	-9,50	-0,25	-0,44
CO2 Value	1,35	0,39	0,74
Total	-37,82	-2,89	-5,89
Implied Value for CO2	8109,22	2330,24	2338,86
	Indirect Welfare Effects (Δ in billion €'s)		
Other Externalities	29,66	0,28	0,22
Paternalism	-7,21	1,95	3,31
Fines:			0,44
Total:	-15,37	-0,67	-1,92

The table gives the aggregated effects over all countries and firms for each policy simulation. Column I solves for the optimal abatement strategy given baseline assumptions but at parameter estimates of the RC Logit I model without endogenous characteristics. Column II is the same as Column I Table 6 but using the estimated cost function. Column III is the same as Column IV of Table 6 but introducing fines as an upperbound for the Lagrangian multiplier. See the text for the assumptions behind the welfare calculations.

Table A6: Sunk Cost Upper Bound Estimates

Deviation:	$0.9 * (\tau + g)$	$0 * (\tau + g)$
BMW	0.07	4.92
Daimler	0.49	14.46
Fiat	0.08	5.08
Ford	0.04	2.24
GM	0.14	7.86
PSA	0.01	0.51
Renault	0.05	2.71
VW	0.29	20.86
Asian	0.21	11.38
Total for Industry	1.38	70.00

The table gives the estimated upper bounds on sunk costs from 18 different simulations. Column I presents the difference in the variable profits obtained from optimal compliance and from deviating from the optimal strategy. In the deviation $(\tau + g)$ is restricted to 90% of the optimal $(\tau + g)$. The second column is the loss in variable profits for each firm when it is restricted to fully comply with sales mixing while all other firms are responding optimally. Simulation I from Table 6 is used as the base to compute the deviations in the variable profits. The differences in variable profits are counted for 6 years and discounted with a rate of 6%.