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 2020

Abstract

This paper studies the introduction of an EU-wide emission standard on the automobile market. Using panel data from 1998-2011, I find that firms decreased emission ratings by 14%. Firms use technology adoption and gaming of emission tests to decrease emissions, rather than shifting the sales mix or downsizing. I find that the standard missed its emission target, and from estimating a structural model, I find that the standard was not welfare improving. The political environment in the EU shaped the design and weak enforcement and resulted in firms' 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 policymakers are taking up the challenge to reduce the use of polluting petroleum liquids. Emission standards are one of the primary policy tools used to achieve this goal. This type of regulation sets mandatory limits on average emission rates (or fuel economy) across the fleet. 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 152 g CO_2/km in 2016. Second, before the standard, the EU had no regulation on CO_2 emissions. The introduction of the standard thus allows me to study how the market equilibrium changes with the introduction of an emission standard. Third, the EU standard is an attribute-based regulation (ABR); the policy target not only depends on CO_2 emission but also on the attribute vehicle weight. The attribute basing makes the policy target less stringent for firms producing heavier vehicles.[2] Understanding the effects of the EU standard is helpful to guide the design of this type of regulation in the future and in other markets across the world.

Evaluating the welfare impact of emission standards is not an easy task. Firms can choose between different abatement strategies to comply with a standard, and these strategies have different effects on market equilibrium. The first strategy is to change pricing to shift the sales mix to vehicles with CO_2 emissions below the target. I define this strategy as mix-shifting. 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 emissions as measured in official ratings without improving the actual emission on the road. Throughout the paper, I refer to official CO_2 emissions as the emission ratings that firms report from lab tests and actual CO_2 emissions as the emissions that vehicles emit on the road. These strategies change the prices, product attributes, product sets, and market outcomes in different ways. Firms 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 matters for the costs of different abatement strategies. This paper presents a model that accounts for firms' abatement strategies and the design of the policy to evaluate the EU standard.

[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.

In the first step, I explain the trend in sales-weighted official CO_2 emissions between 1998 and 2011 in the EU market. Knittel (2011) shows that vehicles sold on the US market change between 1980 and 2006 in terms of the amount of attributes that automakers offer relative to fuel economy 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 official emissions reduce by 14% after the policy announcement. This reduction is explained by decreases in official emissions of vehicles, while the attributes and the product set offered by firms do not change, showing no downsizing. The results show that technological progress is twice as fast after the regulatory announcement, in line with the findings of Klier and Linn (2016). Firms respond to emission standards by increasing the speed of technology adoption, at least when we look at the official emission ratings. Reynaert and Sallee (2019) study the fuel-efficiency of thousands of vehicles on the road to find 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% is due to gaming.

The observed decrease in the official emission ratings is so strong that almost all of the firms reached the emission target before it became partly binding in 2012. However, looking at market outcomes before and after the policy is insufficient to find the welfare effects of the policy as we cannot separate the policy from the effects of 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] Demand, costs, and market fundamentals potentially change in the eight-year gap between the policy announcement and implementation. 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 to study how that impact changes when I change the policy design. In the counterfactual, I assume that the product sets of the firms and the product attributes are fixed. I then model heterogeneous consumers making discrete choices between vehicles. The firms have three strategic choices: mix-shifting, technology adoption, and gaming. I study how equilibrium changes with each strategic choice. 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.

To estimate these primitives, I rely on the rich panel data from before the policy announcement when firms' decisions are not affected. I estimate demand following the framework of Berry, Levinsohn, and Pakes (1995) and extend the model to allow for multiple endogenous characteristics: fuel economy, horsepower, and weight in addition to prices. I instrument the endogenous characteristics with variables that use exogenous variation in the competitive environment, the global sales of vehicles, and the manufacturing structure. Given the estimated demand, the first-order conditions of the profit function provide estimates of the marginal costs of vehicles as well

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

as the marginal benefits and costs of supplying fuel consumption (and emissions). The regulation increases the marginal benefit of lowering fuel consumption, as vehicles with lower emissions help the firm attain the standard. To predict how marginal cost increases with technology adoption, I rely on two approaches. First, I use the observed correlation in the data between fuel consumption and marginal cost. Second, I use engineering predictions submitted with the policy proposal, see TNO (2011). In line with recent evidence presented by Whitefoot, Fowlie, and Skerlos (2017), these engineering predictions show that fuel consumption is relatively adjustable in the short term with available technology. Given this evidence, I choose to model improvements in fuel consumption as a continuous choice that is made simultaneously with pricing. 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 challenging to estimate, and I rely on the level of gaming from the emission decomposition. The estimated model with multiple strategic decisions on continuous variables is similar to the recent contributions of Fan (2013) and Crawford, Shcherbakov, and Shum (2019).[4]

Given the estimated primitives, I introduce the EU policy in the model and simulate how it impacts the market equilibrium. Increases in costs from technology adoption beyond the willingness to pay for fuel consumption imply that the regulation decreases consumer surplus and profits. Because of the gaming, the reductions in actual CO_2 emissions are a mere 5% instead of the 18% target. The sum of the value of emission savings and consumer and profit losses is negative so that the regulation reduces welfare. However, when I consider two additional non-targeted welfare effects, I find the emission standard to have a small positive impact. The two additional welfare effects are reductions in other externalities (local pollution, congestion, and accident risk) and correction of consumer undervaluation of fuel economy. I estimate that consumers do not fully take into account future fuel expenses, so when the regulation provides consumers with more fuel-efficient vehicles, this increases their experience utility. The simulation also reveals that when I restrict firms to mix-shifting, the standard has different welfare effects. The total sales and actual emissions sharply decrease at a very high cost for consumers and firms.

Next, I study how the design of the regulation induced firms to choose the combination of technology adoption and gaming. I first study the attribute basing which makes the emission target dependent on vehicle weight. Firms selling more lightweight vehicles face a more stringent attribute-based target. I find that attribute basing makes abatement with mix-shifting much more costly because firms have to distort prices more to reach the target. If the regulation has a flat target without attribute basing, firms would opt for mix-shifting and technology adoption. The flat target reaches actual CO_2 emission reductions of 11%, much closer to the 18% target. Why, then, was the attribute basing introduced? The simulations show that the positions of the governments are in line with the interests of their firms. The attribute basing redistributes the incidence of the regulation between French, Italian, and German producers. Newspaper articles show that French and Italian

[4] Another option would have been to model technology adoption as a discrete choice for firms in a two-stage model, as in Eizenberg (2014) or Wollmann (2018). However, this approach is computationally less suited for a market with hundreds of products and a continuous characteristic, such as improvements in fuel consumption.

governments were in favor of regulation without attribute basing, while Germany lobbied for a steep attribute design.[5] Additionally, 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. The attribute basing and enforcement failures show the importance of the political environment for the effectiveness of emission standards.

The paper makes several contributions. First, I show that emission standards can induce technology adoption and gaming by firms. Previous literature does not study the equilibrium effects of the abatement strategies in detail. 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 of 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 abatement strategies. Third, the model allows me to study how the design of the regulation affects the outcomes. Attribute basing increased the pressure on firms to adopt technology, and weak enforcement allowed for gaming so that the policy missed its emission reduction target. These 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.

Related Literature 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 to evaluate the US CAFE standard. In an extension Jacobsen (2013) also considers technology adoption using the framework of Austin and Dinan (2005) in which firms respond to an emission standard by changing both prices and technology. 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.[6] My analysis contributes by considering an equilibrium model where firms choose between the compliance options of mix-shifting, technology adoption, and gaming. Reynaert and Sallee (2019) estimate the amount of gaming in the EU market. I use these on-road fuel efficiency estimates to disentangle official from actual emissions. I also use the framework to estimate the effects of gaming on consumer surplus which depend on consumer awareness of gaming. Here, I embed the interaction between choice and gaming in the full equilibrium model with other abatement strategies, and I focus on overall

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

[6] 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)

welfare rather than consumer welfare.

This paper adds to a body of literature that studies policies targeting vehicle emissions, reviewed in Anderson and Sallee (2016). Holland, Hughes, and Knittel (2009) show that none of the welfare effects of emission standards are theoretically determined. A standard might decrease or increase emissions from new vehicles, depending on the distribution of the price elasticities of products below and above the emission target. 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 substantial shadow cost on domestic firms. Both Klier and Linn (2012) and Whitefoot, Fowlie, and Skerlos (2017) evaluate the welfare effects of emission standards while introducing endogenous product characteristics in the model and allowing firms to change vehicle characteristics. Both of these papers assume that the level of technology is fixed and consider the optimization 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. Ito and Sallee (2018) and Whitefoot and Skerlos (2012) discuss the economic effects of attribute-based regulations. 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. Recent work has considered additional margins of the policy. Jacobsen and van Benthem (2015) study the effect of 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.[7]

The paper is structured as follows. Section 2 describes the policy and 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

[7] 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).

target for a firm is as follows:

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

The target binds over each producer's fleet of new vehicles sold in a calendar year. The trading of excess emissions between producers is not allowed.[8] 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 $a = 0.046$.[9] 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 law in 2009. Deters (2010) gives an overview of the full legislative process and political background. In 2012, 65% of manufacturer's sales had to comply with the emission standard. This share rose to 75% in 2013 and 80% in 2014, and the standard was fully binding from 2015 onward. I do not model the phase-in period. Every firm succeeded in reaching the full target by 2014.

Figure 1: Compliance of Firms in 2007 and 2011

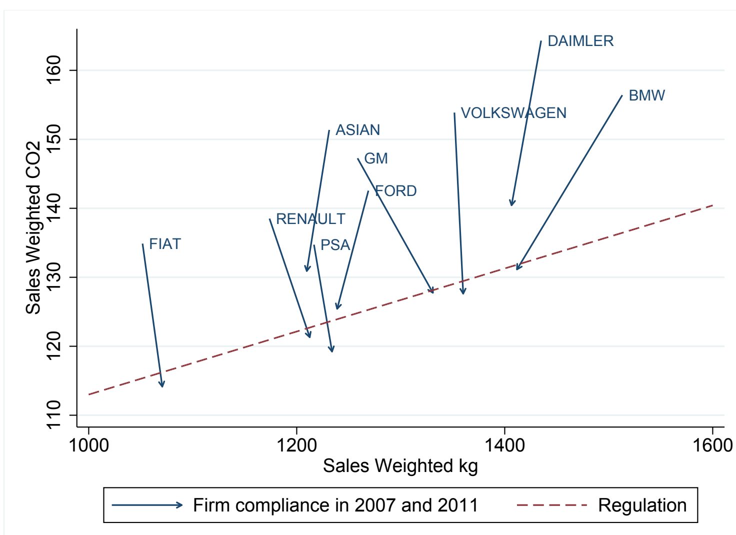

The starting point of each arrow gives the sales-weighted CO_2 and mass for each producer in 2007. The end of each arrow gives the same point in 2011. The dashed diagonal line is the regulation, which binds in 2015.

[8]Manufacturers can obtain lower average emissions by gathering super credits. These credits are given for vehicles that emit less than 50 g/km. Small manufacturers making less than 30,000 vehicles per year face separate standards. I ignore both of the exceptions as they count for a tiny share of the total market.

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

EU member states heavily debated the specifics of the regulation during the drafting of the law. Lobbying efforts by EU member states, firms, and environmental groups were covered by news outlets; see Footnote 5. France and Italy were strongly in favor of a flat standard. At the same time, Germany wanted an ABR with 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 political agreement was to have an ABR with a slope of $a = 0.046$ that aimed at equalizing compliance efforts between firms. The production of each firm mostly takes place within the home country, and the automobile sector is an important source for employment.[10]

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 (SUVs) face a different, less demanding target than passenger cars. Additionally, firms are allowed to trade excess emissions over time and with other firms.[11] From 2012 onward, the CAFE standard also has an attribute-based part, i.e., the target varies with the footprint.

The EU Commission relied on several studies to support the design of the EU emission standard, see TNO (2011). The policy report only includes possible technology adoptions that should be readily available for the car makers at no fixed or development costs. In designing the regulation, the policymaker clearly had the channel of technology adoption in mind. The document lists which technologies are available and how adoption affects marginal cost.[12]

The New European Driving Cycle (NEDC) is the procedure to determine official CO_2 emissions and fuel consumption.[13] 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. None of these features impact how the consumers experience vehicles on the road, nor do they impact the marginal cost of production. I define gaming as the firm efforts that decrease the official emission ratings but not the actual

[10]The car industry represents 6.1% of EU employment and 11.4% of manufacturing employment, see https://www.acea.be/statistics/tag/category/employment-trends, read on March 4, 2020.

[11]Contrary to the CAFE standards in the US, the EU standard has no banking system for excess emissions. 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.

[12]The technologies are reduced friction, direct injection, variable valve timing and actuation, more efficient transmissions, engine start-stop, improved aerodynamics, low rolling resistance tires, and improvements in ancillary systems and auxiliaries can be used to reduce emissions. The document also considers engine downsizing, weight savings, and hybridization, which I do not consider to be technology adoption as they involve downsizing the vehicle. Knittel (2011) also gives several examples of specific technologies implemented in the US. The International Energy Agency reported a possible 40% improvement in fuel efficiency from available technologies in 2005. See http://www.iea.org/publications/freepublications/publication/technology-roadmap-fuel-economy-of-road-vehicles.html.

[13]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 carmakers are not in compliance with Euronorms.

on-road emissions. The costs of gaming are 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 higher than the costs of other compliance mechanisms. 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.

Data The main data set is obtained from a market research firm (JATO dynamics). It contains sales and product characteristics for each passenger car sold during 1998-2011 in seven European countries: Belgium, France, Germany, Italy, Great Britain, The Netherlands, and Spain.[14] The characteristics and sales are given for several engine variants of a car model at the country level with a yearly frequency. The country is the geographical market. A model is defined as a brand/model name/body type combination (e.g., Volkswagen Golf Hatchback).[15] 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 $C0_2$.[16]

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 highway driving patterns and excludes the use of auxiliary features, such as air conditioning. Real-world emissions thus differentiate from these test values. Reynaert and Sallee (2019) develop a measure of on-road emissions and show that reductions in the official CO_2 ratings do not fully translate into actual savings. Carmakers can calibrate engines with defeat devices and specific software so that they perform much better on the test. Gaming is defined as improvements that do not translate to the road. I use the measure of on-road emissions in the analysis to disentangle technology adoption from gaming.[17] This information is available for a limited sample of vehicles. Next, I use 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, I use the data on

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

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

[16] 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 increases emissions of other pollutants such as N_{OX}.

[17] Reynaert and Sallee (2019) construct on-road emissions from a panel of 12,000 drivers visiting a fuel station 22 million times. Odometer readings and fuel purchases are used to construct fuel consumption on the road.

fuel prices (from DataStream), GDP/capita, and the number of households in each country (from Eurostat), 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, 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, Honda, Hyundai, Mitsubishi, Nissan, Suzuki, and Toyota, as one decision-maker.[18] The list of excluded brands and a detailed description of the data manipulations are in Appendix A1. 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 the total reported sales in the sample.

Throughout the paper, I split the dataset 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. I use these data 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 rely only on data before the policy announcement and use the years 1998-2007. This exploits 30,000 year/market/model-engine observations. I use the data from the year 2007 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 was far above the target and had 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 sharply. I observe a sharp downward trend in emissions towards the standard 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.[19] Table 1 shows the change in the sales-weighted vehicle characteristics between 2007 and 2011. CO_2 emissions decrease by 14% while other characteristics grow moderately. In Appendix Figure A1, I plot the sales-weighted characteristics over time from 1998 to 2011 for the EU and the US. The most remarkable trend in the EU is the evolution of CO_2 emissions. Sales-weighted emissions are constant until 2002, decline approximately 6% until 2007, and then plunge by 14% in the last four years of the sample. This shift coincides with the announcement of the emission standard by the European Commission. In

[18]19 I combine the fleet of the Asian carmakers because most of these firms do not have a broad product set. This makes finding a price equilibrium with solely price changes 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.

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

the US, emissions decline 3% between 2007 and 2009, but then, in contrast to the EU, emissions remain constant at 90% of the 1998 level.

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.

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. I decompose the shifts over time in the correlation between characteristics and official emissions into changes in technology and changes in the composition of the vehicle fleet. I also investigate to what extent technological change translates to the road by comparing actual with official emissions.

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

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

where e_{jy} are emissions, the technology parameter ζ_y is a year fixed effect, x_{jy} are characteristics, η trade off parameters denoting how emissions change with characteristics and ϵ_{jy} is an error term.[20] 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_y is input neutral as it enters multiplicative in levels.[21] 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_{jy} 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 correlate with explanatory variables, so that η cannot be interpreted as a causal relationship between the characteristics and emissions. The goal of estimating (2) is to see

[20] 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.

[21] The main specification has 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 reported in Appendix Table A3.

changes in emissions over time while controlling for the correlation between emissions and other characteristics.

Table 2: Technological Progress Estimates

	Model 1		Model 2		Model 3		Model 4		Model 5	
			Official Ratings					actual 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 yearly percentage change in emissions due to technological improvements obtained from differencing the year fixed effects in (2). The standard errors are computed with the Delta method. Years after the policy announcement are shaded. Model 2 and Model 5 report the average of firm-specific trends, Appendix Table A4 reports firm-level estimates.

The estimation of (2) is useful because it reveals the abatement strategies that firms resort to without relying on a structural model. In Figure 1, I show how emissions starkly decreased after the regulatory announcement. If firms choose to alter the sales mix or to downsize, then the part of the emissions that is explained by characteristics x_{jy} should decrease over time. With mix-shifting 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_y over time. As discussed above, firms gaming the test is a concern. 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 actual emissions as the left-hand-side variable. I use the actual emissions constructed in Reynaert and Sallee (2019). Looking at actual emissions allows me to test if the technological improvements revealed in the official data translate to on-road measures. I now present estimates of the technology parameters in Table 2 and the decomposition of emissions over time in Table 3. I discuss the trade-off parameters and robustness in Appendix A2, Table A2 and Table A3.

The technology parameters ζ_y are derived from the time fixed effects in each regression and are plotted in Table 2 for Models 1 to 5. Model 1 is the baseline specification close to Knittel

(2011), and includes trade-off parameters for horsepower, weight, footprint, and height. Model 2 allows for a firm-specific trend in technology. 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. Model 3 replicates Model 1 with the smaller sample. Model 4 changes the dependent variable so that the relation between the actual emissions and characteristics is estimated. Model 5 introduces firm-specific trends. Models 2 and 5 report averages of firm-specific trends. The bottom panel of the table presents the difference in the mean of ζ_y for 1999-2007 and 2008-2011, corresponding to the pre- and post-policy 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_y are more than 2% higher after 2007 than in the years before. The difference is statistically significant. After 2007, the estimates reveal a significant increase in the pace of technology improvement.[22] In the official emission measures, I find evidence of rapid changes in technology adoption, but the emission reductions might not translate to the road. Once I look at changes in actual emissions in Models 4 and 5, I find that technology growth does not increase as rapidly in the post-policy 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%. [23]

Overall, this reveals that firms do engage in technology adoption after the policy announcement but that a large part of the technology aims at reducing official emission ratings rather than actual emissions. The estimates show that 70% of the additional reduction in official emissions comes from gaming, while 30% comes from actual technology improvements. I will use this 30/70 ratio in the structural model below.[24]

Decomposition of the changes in emissions I use the estimated relation (2) to decompose the decline in emissions over time. The decomposition is as follows: I predict emissions, $\widehat{e}_{jy}$, as the fitted values of regression (2) with Model 2. I also predict $\overline{e}_{jy}$ using (2) but fixing the technology level at $\zeta_y = \zeta_{2007}$.[25] The evolution in the sales-weighted values of $\overline{e}_{jy}$ only changes when the composition of characteristics x_{jy} changes over time, while $\widehat{e}_{jy}$ changes both because of underlying characteristics and technology.

The results in Table 3 show that between 1998 and 2007, sales-weighted emissions without technology $\overline{e}_{jy}$ declined slightly from 156 to 152. Technology improvements are responsible for the

[22]Table A4 reports firm-specific trends with similar patterns except for BMW and Renault. All firms change the trend of technology adoption, but the magnitude of the change differs. BMW and Renault have the lowest difference in technology growth (0.9% and 0.8%), while Volkswagen has the highest difference (5.4%).

[23]Table A3 presents further robustness checks with official emission decreases estimated to be between 2.1% and 3.3% more rapid after the policy announcement.

[24]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 creates substantial differences in welfare outcomes between firms. The second reason is that Reynaert and Sallee (2019), based on a extended period of data until 2015, find that all firms game in large and similar amounts.

[25]I re-scale each of the predicted emissions with the attribute-based target function, such that the numbers are actual distances from the regulation. Not re-scaling gives similar results.

Table 3: Decomposing the Decrease in Emissions

		All Vehicles		Existing Models (2007≤)		New Models (> 2007)	
		No Tech.	Tech.	No Tech.	Tech.	No Tech.	Tech.
	True	$\overline{e}_{jy}$	$\widehat{e}_{jy}$	$\overline{e}_{jy}$	$\widehat{e}_{jy}$	$\overline{e}_{jy}$	$\widehat{e}_{jy}$
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 corrected with the attribute function $f(w_j)$ to represent the target values of the regulation. All predictions use the estimates from Table A2 and Table 2 Model 2. The columns $\overline{e}_{jy}$ contain sales-weighted predicted emissions keeping technology constant at $\zeta_y = \zeta_{2007}$. The columns $\widehat{e}_{jy}$ contain sales-weighted predicted values for emissions with estimated ζ_y.

remaining moderate decline in official emissions between 1998 and 2007. After 2007, the sales-weighted emissions without technology $\overline{e}_{jy}$ remain constant. Because $\overline{e}_{jy}$ is constant, mix-shifting or downsizing cannot explain the decline in official emissions. When I split up the average sales-weighted emissions into vehicle models released after and before 2007, the results show that the emissions $\overline{e}_{jy}$ of vehicles released before 2007 remain constant.[26] Vehicle models released after the policy 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. The observed decline in emissions is not attributable to changes in the sales mix or the release of new downsized fuel-efficient vehicles. However, I cannot rule out that the vehicles would have grown in terms of weight, size, and horsepower without the policy.

The sales-weighted emissions with technology $\widehat{e}_{jy}$ 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 engine improvements are installed widely across the fleet.

These results present two concerns so far. First, the Great Recession and the EU debt crisis took place just after the policy announcement. In general, an economic downturn would lead consumers to spend less on durables and 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 strengthens 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 regulation in response to the EU-wide policy.[27] This combination of new local taxation and 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 model firm behavior in response to a standard to single out the effects of the EU emission standard.

In summary, I present 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 gaming explains 70% of the decline in official emissions, technology adaption explains the remaining 30%. 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.

[26] An example of a newly released model is the ‘Citroen DS3 Hatchback’, released in 2009.

[27] 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. The supply-side model describes how firms can use different abatement strategies in response to a binding emission standard. This model will be the basis to simulate the welfare effects of the emission standard in Section 6. In the estimation, I use pre-policy data to estimate a standard supply model in which firms do not face the regulatory constraint. Finally, I discuss how the design choices of the regulator regarding attribute basing and enforcement affect abatement choices.[28]

Demand A market is defined as a country, indexed by m, observed in a calendar year y.[29] In each market my there are A_{my} 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 $u_{i0my} = \varepsilon_{i0my}$, and for products $j = 1, \ldots, J$ it is:

$$u_{ijmy} = x_{jmy}\beta_i^x - \beta_i^e d_{jmy} e_{jmy} k_j - \alpha p_{jmy} + \xi_{jmy} + \varepsilon_{ijmy}, \tag{3}$$

where x_{jmy} is a vector of observed product characteristics, $d_{jmy}e_{jmy}k_j$ are fuel costs (fuel prices d_{jmy} times emissions e_{jmy} times a fuel type specific factor k_j that translates emissions into fuel consumption as explained in footnote 1), p_{jmy} is the vehicle price and ξ_{jmy} is an unobserved characteristic of vehicle j in market my, 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 is the marginal utility of income or price valuation and ε_{ijmy} 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.

The coefficient β_i^e measures the consumers' valuation of fuel costs. Consumers use vehicles for several years and 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. The random walk assumption 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 and

[28] The abatement strategies discussed do not need to occur mutually exclusively. Firms 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.

[29] The geographical market corresponds to a country in the empirical analysis because I observe prices and quantities at the country level.

a capitalization factor.[30,31] I estimate the valuation for fuel costs with data from before the policy announcement where firms had no incentive to game.[32] I assume that consumers know what the official rating signals about the actual fuel costs when estimating the model. When I introduce the policy in the simulation, firms can game the emission ratings and it becomes unclear if consumers can determine if decreases in fuel costs are due to gaming or technology adoption. I consider both the case where consumers are sophisticated and where they are fooled.

Each consumer i in market my chooses the alternative $j = 0, \ldots, J$ that maximizes her utility. The predicted market share of vehicle j in market my 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:

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

where δ_{my}, the mean utility, which collects all terms in (3) that do not vary across individuals, and μ_{jmy} is the term that captures the individual idiosyncratic deviations from the mean utility as follows: $\mu_{jmy}=\sum_k \sigma^k \nu_i^k x_{jmy}^k$. To complete the demand side, I set the observed market share $s_{jmy} = q_{jmy}/A_{my}$ 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: $\mathbf{s}_{my} = \mathbf{s}_{my}(\delta_{my}, \sigma)$.

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. The timing of the game is assumed to be as follows. First, I assume that the vehicle fleet of firms is exogenous. Each firm sells hundreds of differentiated products and I do not model the decision regarding which vehicles firms choose to offer. Second, at the time of their strategic decision, firms have perfect information on all observable and unobservable characteristics. Third, given their vehicle fleet, firms have the option to change the prices and fuel consumption of each of the products simultaneously. The fuel consumption can be changed by adding technology or by gaming the test that determines the level of fuel consumption. Finally, when firms face regulation, the choices are altered to comply with the regulation that binds across total sales from all EU markets. The goal of this section is to understand how firms change decisions on pricing and fuel consumption levels when facing a binding regulation.

[30] The total expected fuel costs can be written as $E[\sum_{s=1}^{S}(1+r)^{-s} m_i d_s e_j k_j]$ 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 $\rho m_i d_1 e_j k_j$, 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.

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

[32] In Reynaert and Sallee (2019), we show that the gap between the official and on-road ratings was constant up until 2007. We expect a gap between lab and on-road ratings because driving circumstances differ.

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

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

where π_{fm} are variable profits for firm f in market m, F_f is the set of products firm f sells, κ is the level of the standard and $f(w_j)$ 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 obtain reductions or penalties on their official emissions.[33] I model technology adoption and gaming as percentage reductions in fuel consumption so that $0 \leq t \leq 1$, $0 \leq g \leq 1$, and $0 \leq t + g \leq 1$. Firms adjusting t or g induce sunk costs $C(t,g)$, which I discuss in Section 6. I model simultaneous joint choices of p, t and g and ignore dynamics.[34] I follow Goldberg (1998) and Jacobsen (2013) and write the Lagrangian of firm f net of sunk costs in market m as:

$$\mathcal{L} = \sum_m \sum_{j \in F_f} [(p_{jm} - c_{jm}(t_j) - \lambda L_j(t_j, g_j)) s_{jm}(p,t,g) A_m], \tag{6}$$

where c_{jm} denotes marginal costs, λ is the shadow cost of the regulation per unit of sales. The shadow cost is multiplied by $L_j = (1 - t_j - g_j)e_j - f(w_j) - \kappa$, the distance of each product from the target. When $L_j < 0$ (> 0), an additional sale of vehicle j brings the average sales-weighted emissions closer to (further away from) the target. The shadow cost λ gives the cost per unit sold of deviating one unit from the standard. If the standard is non-binding, $\lambda = 0$ and (6) reduces to a standard multi-product profit function. The shadow cost λ takes the same value for each vehicle in the fleet F_f because, in equilibrium, firms equalize the shadow costs over their fleet to be cost-efficient. The shadow cost differs between firms because trading of excess emissions is not allowed.

Three first-order derivatives of $\mathcal{L}$ to p, t, g, together with sunk cost, determine the optimal solution to the profit maximization problem. I introduce the following matrices to write the derivatives concisely: Φ denotes a $J \times J$ 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 $J \times J$ matrix of first-order derivatives of market shares with respect to $k = p, t$ or g. Define $\circ$ as the Hadamard product. Bold vectors ($\mathbf{q}, \mathbf{p}, \mathbf{c}, \lambda \circ \mathbf{L}$ and $\mathbf{e}$) have dimension $J \times 1$, with λ now a vector with the firm-specific shadow cost when $j \epsilon F_f$. Vector $\mathbf{c}'_{\mathbf{t}}$ is $J \times 1$ with scalar derivatives $\partial mc_j / \partial t_j$ as elements. The J first-order

[33] 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.

[34] In reality, firms choose to incrementally adopt technology and gaming over 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.

derivatives with respect to prices, technology, and gaming are:

$$\frac{\partial \mathcal{L}}{\partial p} = \mathbf{q} + \mathbf{\Phi} \circ \triangle_{\mathbf{p}}(\mathbf{p} - \mathbf{c} - \lambda \circ \mathbf{L}) \tag{7}$$

$$\frac{\partial \mathcal{L}}{\partial t} = (-\mathbf{c}'_{\mathbf{t}} + \lambda \circ \mathbf{e}) \circ \mathbf{q} + \mathbf{\Phi} \circ \triangle_{\mathbf{t}}(\mathbf{p} - \mathbf{c} - \lambda \circ \mathbf{L}) \tag{8}$$

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

The first-order derivatives with respect to prices show the standard trade off between increases in mark up (first term on the right-hand side) and losses from reduced sales (second term on the right-hand side) when increasing the price. The first-order derivatives with respect to technology adoption show that firms trade-off marginal cost increases $\mathbf{c}'_{\mathbf{t}}$ with relaxing the regulatory constraint and the benefits of increased market share (we expect $\triangle_t > 0$). The first-order derivatives with respect to gaming are very similar but gaming induces no marginal cost changes. The system is defined at the product level for t and g. In the counterfactual, I impose g because I lack cost data and I solve for firm level choices of t to shrink the dimensionality of (8) to the number of firms.

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. The computation 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 decrease. However, it might be that the standard results in more fuel-efficient vehicles but also in more purchases. It is not certain that total emissions will decline. In the welfare calculation, I 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), $\lambda = 0$ to a market with a binding standard, $\lambda > 0$. The changes in the market depend on the abatement strategy of firms, i.e., mix-shifting, technology adoption, or gaming. I discuss the effect of each strategy in turn, keeping the other strategies fixed.

Abatement by mix-shifting 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 the prices of vehicles with emissions above the target ($L_j > 0$). The first-order conditions with respect to prices (7) show that the shadow cost of the regulation, λ, determines to what extent prices will be distorted from a no-policy equilibrium, as follows:

$$\mathbf{p} = \mathbf{c} + \lambda \circ \mathbf{L} - (\mathbf{\Phi} \circ \triangle_{\mathbf{p}})^{-1}\mathbf{q} \tag{10}$$

The regulation implicitly taxes vehicles with $L_j > 0$, while it implicitly subsidizes vehicles with $L_j < 0$. This change in the relative prices of products shifts the sales mix towards vehicles with lower fuel consumption. It is important to stress that the abatement strategy of changing prices to shift the sales mix is driven by the term $\lambda \circ \mathbf{L}$. This term captures the first-order effect of the regulation on prices and shows how the emission target constrains the pricing equilibrium. Abatement by technology adoption and gaming, as discussed immediately below, results in changes in attributes of different products. These changes also affect pricing, as firms face different marginal costs and competition. But, I do not interpret these price changes as mix-shifting.

The incidence and effectiveness of this abatement strategy largely depend 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 the 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 depend on consumers response to price changes but also depend on the position of the fleet relative to the target. Firms with a fleet that is better adapted to the standard might increase profits, as their prices need less distortion than their competitors. The empirical mode allows me to identify the own and cross-price elasticities for all products and to simulate the shifts in sales when I 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 fuel consumption. From (8), we see that the regulation gives incentives to firms to adopt more technology. The first term, $(-c_t' + \lambda \circ e) \circ q$, specifies that increases in fuel efficiency change the marginal cost of each unit sold. But, the technology also relaxes the regulatory constraint, resulting in a marginal benefit of λ. Without a regulation $\lambda = 0$ and the marginal benefits of technology adoption is lower. The second term captures the change in sales from technology, $\Phi \circ \triangle_t$, multiplied by profits per unit, $(p - c - \lambda \circ L)$. Finally, implementing technology can involve sunk costs $C(t, g)$ related to making changes in the production.

Each unit of technology adoption lowers L. The fleet of the firm shifts closer to the regulatory target and the regulatory constraint is relaxed. Technology adoption reduces the shadow cost of the regulation λ as firms need less distortion away from the preferred price schedule to comply. When determining to what extent firms adopt technology or mix-shift in equilibrium, the shadow cost λ and the amount of technology is determined endogenously.

The welfare effects of this strategy are again undetermined theoretically. Consumer face two offsetting effects. The standard puts upward pressure on prices as marginal costs increase with c_t'. Increased prices reduce 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 increase or decrease. Even though the resulting vehicle fleet has lower emissions, it is not clear that emissions decrease, i.e., if consumers benefit from technology, they might purchase more vehicles. The changes in marginal costs and the degree of pass-through determine the overall effect. Given that different vehicles have different fuel consumptions, the equilibrium prices also adjust, creating another welfare change for consumers. For firms, technology adoption is a preferred abatement strategy if the shadow cost of mix-shifting is high, and the marginal cost and sunk cost changes of technology adoption are low.

Abatement by gaming Gaming helps with compliance in a very similar way as technology adoption. The first part of (9) shows each unit of gaming reduces L and the shadow cost of the regulation λ until firms can price without regulatory constraints. The second term in (9) accounts for market share changes due to gaming. Gaming does not affect the on-road experience of consumers and does not reduce the actual fuel consumption. However, when purchasing a vehicle, consumers might be fooled by the advertised fuel consumption. The EU introduced the official ratings for advertising and windshield stickers to inform consumers about fuel costs and pollution. As such, gaming has demand effects that distort consumer choice. In the results, I discuss what happens when consumers are fooled or when they are aware of gaming. The effects on consumer welfare 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 distort consumer choices. This choice distortion reduces consumer surplus and might also lead to higher prices. Gaming might also have positive effects for consumers, as it allows firms to avoid costly compliance and high pricing. The effects on emissions from gaming can be pervasive. Each reduction in fuel consumption obtained with gaming does not reduce actual emissions. Consumers respond to the decreases in perceived fuel costs by substituting toward 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.

The costs of gaming are very different than those of technology adoption. As summarized in Section 2, gaming is essentially a sunk cost that mainly consists of expected legal fees and consumer blow-back. This cost $C(t, g)$ needs to be positive to prevent firms from choosing for gaming as the only abatement strategy. Both terms in (9) are benefits to the firm, and these benefits will be weighted against the expected costs. Without empirical information on $C(t, g)$, I choose to pin down gaming using the evidence from Section 3.

Abatement by downsizing I assume the fleet of each firm as given. In previous work, Klier and Linn (2012) and Whitefoot, Fowlie, and Skerlos (2017) tried to relax this assumption by modeling not only improvements in fuel consumption but by allowing firms to optimize a broader set of attributes. This strategy is described as downsizing: 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 downsize instead of changing prices. Consumer loss is similar. Whitefoot, Fowlie, and Skerlos (2017) use an interesting engineering model. I do not model downsizing for two reasons. In Section 3, I find no evidence of downsizing in the EU. A model that allows strategic choices over several characteristics is computationally challenging for the detailed engine version level data used in this analysis, and it would introduce difficulties with multiple equilibria.

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 consider two regulation design choices. First, I study the EU's choice to have a policy that is attribute-based and upward sloping, 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 standard. For a flat standard, $L'_j = [e_j - \kappa']$, and $f(w_j) = 0$. The target function is a horizontal line at κ' and all firms need to reach the same level of sales-weighted emissions. Each firm has a different set of vehicles with $L_j < 0$ and $L'_j < 0$. Because of the upward sloping ABR in the EU, the sales of many small, lightweight vehicles do not help with compliance, while they help with compliance under a flat standard. Attribute basing reduces the number of products that have $L_j < 0$. Price changes become much costlier because firms have can shift sales to fewer products. Reducing fuel consumption by technology adoption or gaming helps increase the number of products that have emissions underneath the target. In the empirical section, I show that the ABR increases the stringency for French and Italian firms but reduces the stringency for German firms. The difference between firms matches the lobbying efforts by the governments described at the time of the negotiations. I also show that price changes became so expensive for firms 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 distort the demand and supply of the attribute itself. If heavier cars help with attaining the target, weight is indirectly subsidized, and producers choose to add more weight to their vehicles.[35] In this exercise, I keep the weight fixed and assume no distortions in the attribute itself. I keep weight fixed because I do not observe any changes in weight consistent with a distortion. Additionally, allowing for more choice variables for the firm is computationally costly.

The EU's political structure has also led to weak enforcement of the policy. Above, I discussed how failures in the delegation of enforcement have led to the gaming crisis in the EU market. Theoretically, the sunk cost of gaming $C(g)$ captures enforcement. When enforcement is strict and legal punishment for gaming high, the sunk costs of gaming increases. A higher sunk cost means that gaming as an abatement strategy becomes less attractive so firms resort to mix-shifting or technology adoption. In the empirical section, I discuss what happens if firms do not game EU regulation.

5 Estimation

I use a panel of 7 countries over 10 years to estimate the taste and cost parameters. I restrict the sample to 1998-2007, the years before the policy announcement. This restriction allows me to estimate a model in which firms maximize unconstrained profits as given in (6) with $\lambda = 0$.

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

It also allows me to estimate the consumer valuation of fuel costs before firms start to game emission ratings. I discuss the demand estimation that allows for the 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 project marginal costs out of sample to model technological improvements. With the estimates of consumer preferences, marginal costs, and the functional form assumption, I 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 to be proportional to income y_m in market m, so that $\alpha = \alpha_{my} = \alpha / y_{my}$.[36] I add a set of controls for which I only estimate the mean taste. These include weight, footprint, height, a dummy for 3 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.[37] Finally, I add 331 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_{jm}. I obtain the parameters by minimizing the GMM criterion:

$$\min_{\theta} \xi(\theta)' ZAZ' \xi(\theta) \tag{11}$$

where ξ is a column vector of the J demand unobservables stacked over each market. I obtain ξ after inverting the system of market shares in (4). The matrix of instruments is Z, and A is a weighting matrix. I use the two-step efficient GMM estimator so that the second 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.[38]

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 exclusion restriction allows for the correlation between prices and ξ but assume that all the vehicle characteristic choices of firms are mean independent of ξ. The counterfactual considers strategic choices in response to

[36] This specification allows wealthier markets to be less price elastic. When prices are small relative to the household budget over the lifetime of the durable good, this approximates an indirect utility function with $log(y - p_j)$ as an argument as in Berry, Levinsohn, and Pakes (1995). See Griffith, Nesheim, and O'Connell (2018) for a discussion on how the functional form of the indirect utility in price and income matters for distributional effects of price changes.

[37] 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.

[38] More specifically: (i) I use a nested-fixed point (NFP) algorithm, BLP's contraction mapping with a tight convergence criterion (1e-12) to solve for ξ_{jm}. I compute the contraction mapping in parallel per market, (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).

regulation. However, even before the regulation, firms design products to maximize profits. The concern when estimating demand is that firms know more about consumer tastes than the econometrician when designing the product. This information asymmetry 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. I construct instruments for three sets of parameters: the taste parameter for the price, the taste parameters for endogenous characteristics, and the nonlinear parameters. First, I consider instruments Z^1 for the model with price endogeneity. Second, I 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 ξ_j. 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. More mutable characteristics can be adjusted when information about the demand unobservable is revealed, while fixed characteristics are chosen before and cannot be adjusted. In the first engineering step, firms set the dimensions of the vehicle (footprint and height). In following steps, firms fit the engines in the design. I 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 I need excluded instruments for the following four variables: the three mutable characteristics and prices.

Because the mutable characteristics potentially correlate with ξ, their sums (across own or other firm products) are not valid instruments anymore. I form additional instruments that aim to capture variations in the engine design decisions of firms. Using global production data, I exploit changes in the exposure of each vehicle model to non-EU markets over time. I compute the share of each vehicle model produced in Africa, Asia, Eastern Europe, North America, and South America. Together with the EU production, the shares add up to one. 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 the US market is 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_{jm}. 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 production shares in the different continents, and the size is the production weighted average size of all vehicles produced in a region. These sums capture competition in product space in different production regions. Next, I follow Klier and Linn (2012) and include the average footprint and height of vehicles produced in different classes but 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. I use the attributes of the Q5 as an instrument for the A5 and vice versa. The idea behind the instrument is that vehicles produced on the same platform share both fixed and mutable characteristics.[39] Finally, fuel costs equal fuel prices times fuel consumption. Fuel prices are exogenous, so I interact with fuel prices 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). These instruments are nonlinear functions of the previously described included and excluded instruments and relate to the conditional expectation of the Jacobian of ξ_{jm}: $E[\frac{\partial \xi_{jm}}{\partial \sigma} \mid z_{jm}]$. To approximate the infeasible optimal instruments, I take a two-stage approach. In the first stage, I estimate the nonlinear model and compute the approximate optimal instrument at a guess for the nonlinear parameters. In the second stage, I update the approximation of the optimal instruments at the first stage estimates.[40] This procedure generates a number of additional instruments equal to the number of standard deviations of random coefficients. Appendix A3 gives a detailed overview of all the instrumental variables.

Costs I do not observe marginal costs c_{jm} but obtain them from the first-order conditions of the firms' profit maximization. From equation (7), the demand estimates, and the fact that $\lambda = 0$ in the estimation sample, I compute c_{jm}. Similarly, we can use equation (8) to back out an estimate of the cost of changing technology c'_t. The derivative gives an estimate of the slope of the marginal cost with respect to reductions in fuel consumption before the policy announcement. Information on c'_t is useful because it reveals the cost changes that firms are willing to incur to provide fuel consumption in their products. Lowering fuel consumption further would be more costly than what the firm gains from providing the improvement. If the supply model is correctly specified, any cost function for further reductions in emissions must start at $c'_{t=0}$, otherwise, it is at odds with the revealed choices of firms and consumers. I interpret the pre-policy estimate of the slope as $c'_{t=0}$ and use it as the intercept for the cost slope in the simulations.[41]

I also need an estimate of the cost changes when the regulation requires fuel consumption

[39] Klier and Linn (2012) include all the characteristics of different class vehicles on the same platform. Their 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.

[40] The initial point for nonlinear parameters in the first stage approximate optimal instrument 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).

[41] The counterfactual results also consider a case where we start from another intercept below the estimated one. Market failures in technology adoption are a reason for the actual slope to be smaller than the estimated $c'_{t=0}$.

reductions above what we observe. I use two approaches for this problem. The first approach relies on cost estimates from TNO (2011) that specify convex cost functions for percentage reductions in emissions. I use these functions in the main results and I refer to them as the engineering cost curve. However, I shift the intercept of the engineering cost function from zero to the point where the slope of the engineering function equals $c'_{t=0}$. Appendix A5 discusses the engineering cost function in detail.

The second approach relies on marginal cost estimates from the model. When marginal costs are log-linear:

$$\log(c_{jm}) = \gamma^e e_{jm} + x_{jm}\gamma^d + \omega_{jm}, \tag{12}$$

where x_{jm} is a $1 \times L$ vector of observed product characteristics, market controls, fixed effects on the vehicle model level, and cost shifters, and ω_{jm} is unobserved. Fuel consumption enters the marginal cost, as all else equal, it is likely more expensive to produce engines with lower fuel consumption. The estimated parameter on emissions informs how marginal costs change with changes in e_{jm}. The functional form allows me to predict the costs of further reductions in fuel consumption that I cannot observe within the sample. I estimate (12) by OLS when I derive costs from the RC Logit model with price endogeneity. When fuel consumption, horsepower, and weight are endogenous in the demand, ω_{jm} is potentially correlated with these attributes. I instrument the attributes when estimating (12) using an instrument set $Z^{2'}$ that includes the same instruments as Z^2 except for the instruments based on platform as these might capture cost synergies between vehicles.[42]

Estimation Results Table 4 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 Logit and a RC Logit. The last two specifications allow for endogenous characteristics and use instrument set Z^2. Appendix Table A5 reports the first stage. 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 multiple 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 considerable variation in the taste for horsepower and weight. The estimates change 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. These coefficients matter for the welfare simulations, i.e., when consumers care less about fuel costs relative to price, they 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 only price endogeneity, and the

[42] The cost shifter is an included variable in (12) and not an excluded instrument when estimating the supply. See Appendix A3 for an overview of the instruments.

Table 4: 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 Cons.	-0.027	0.001	-0.047	0.001	-0.031	0.027		
Horsepower	0.629	0.005	0.690	0.006	-0.390	0.174		
Weight	0.742	0.009	0.682	0.011	4.051	0.408		
Base	0.070	0.002	0.028	0.002	-0.326	0.056		
Height	-0.019	0.001	0.053	0.001	-0.101	0.018		
Log Labor Cost	0.386	0.003	0.042	0.004	0.140	0.066		

The table reports the estimated parameters for the demand and marginal cost equations. The first two columns in the Demand Estimates give Logit and RC Logit estimates for the model with endogenous prices and instrument set Z^1. The last two columns give the Logit and RC Logit for the model with endogenous prices and characteristics and instrument set Z^2. The first column gives Marginal Cost estimates for perfect competition (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). Column three instruments fuel consumption, horsepower, and weight with $Z^{2'}$.

mean tastes for horsepower and weight switch sign between the Logit and RC Logit estimate. The standard deviations show considerable heterogeneity in the taste for fuel costs, horsepower, and foreign vehicles in this case. In Appendix Table A6, I discuss the fit of the demand model.

The second panel of Table 4 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. The final specification instruments for endogenous characteristics and this reduces the precision. All marginal cost regressions show that increasing the fuel efficiency of the vehicle is costly. A one-unit decrease in liters per 100km increases the cost by 2.7% to 4.7%.

In the counterfactuals, I base the main results on the engineering cost function. I favor the engineering costs for several reasons. Estimation of (12) relies on changes in fuel consumption observed in the pre-policy sample 1998-2007. The pre-policy period has limited variation in fuel consumption; see Appendix Figure A1. The slope of the estimated cost function to fuel costs is not robust to changes in control variables.[43] The specification with instrument set $Z^{2'}$ has very large standard errors, pointing out that this parameter is difficult to estimate. The engineering estimates exclusively focus on technologies to decrease fuel consumption. In the estimated cost equation, multicollinearity makes it difficult to separate changes in fuel consumption from other attribute changes as carmakers redesign the whole vehicle when releasing a new version.

Finally, I use the model to compute the cost slope that rationalizes the fuel consumption choices of firms before the policy. When consumers care less about fuel costs, it becomes less attractive for firms to provide fuel consumption reductions. Allowing for more endogenous variables in the RC Logit model lowers consumer valuation for fuel costs. The marginal benefit of fuel consumption reductions for firms decreases, so does the implied marginal cost slope in (8). The model with endogenous characteristics implies a pre-policy intercept of the cost slope that is five times less steep than what the price endogeneity model implies. Cost increases caused by the policy are lower when the initial intercept is lower, so it is more likely the emission standard increases welfare.

In summary, emissions enter the model through two channels. First, all else equal, consumers dislike vehicles that have higher actual emissions because they are more costly to operate. Second, producing vehicles with lower emissions 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.

[43] The regression of the predicted marginal costs on fuel costs results in a positive sign when I change time and market controls, so that fuel economy lowers the marginal costs. These issues could be due to multicollinearity as well as changes in markups that are correlated with fuel costs but not captured in the first-order condition, see Langer and Miller (2013).

6 Welfare effects

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. 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 matter 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. Solving for the system of equations (7), (8), (9), $C(t, g)$ and the policy target is computationally infeasible. Solving the full system would require a solution for j unknown prices, t's and g's in each country, information on sunk costs, as well as a λ for each firm. I make two simplifications. First, the first-order conditions for gaming (9) are satisfied only when the sunk cost of gaming is positive. This cost is challenging to determine as it depends on future legal fees, brand reputation, and hidden efforts. I use the reduced form results from Section 3 to fix the level of gaming. I assume that each percentage improvement in fuel consumption includes 70% gaming and 30% actual technology. I solve (7) and (8) for prices and technology while imposing that every 0.3 units of t 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 translates fully to actual emission, or technology only affect official emissions and equals gaming. Second, I reduce the number of conditions to solve for by assuming that firms implement the same amount of technology fleet-wide. Firm-wide technology reduces the number of equations in (8) from j to the number of firms.

I compute the counterfactual for a binding regulation with a target to reduce sales-weighted emissions to 130 g/km. I take the vehicle fleet of each firm in 2007 as given and improve the actual emissions by 6%. This adjusts for the reduction in emissions that would have happened without the regulation.[44] The simulation keeps all vehicles and characteristics of this fleet fixed except for price, technology, and gaming. Finally, I solve for the equilibrium so that each firm complies exactly. In reality, this does not need to be the case, as firms may not comply with the standard and pay fines.[45] Though I cannot show the computed equilibria are unique, I find no other solutions when varying starting values. Firms do not have an incentive to over-comply, as this would distort choices further from the pre-policy equilibrium.

I run the following algorithm. Step 1 is to choose a value for the firm-level technology t and

[44] The different models estimated in Section 3 show a 0.7% to 1% annual decrease in emissions in the pre-policy 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.

[45] In equilibrium, this is not a constraint on the solution as all firms prefer compliance over paying fines in the main scenario. In a scenario where I restrict firms to mix-shifting, the fines matter. I discuss this in Appendix A4.

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 $\lambda \circ L$ for each product. Step 2 solves for the Nash equilibrium in prices given the values chosen in Step 1 and solves for (7). Step 3 evaluates if (8) and the regulatory target is satisfied, 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 (8) and the policy target. This final iteration gives the solutions for t, 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 p, t and λ, I compute the changes in outcomes between the initial equilibrium (no policy and 6% reduction in emissions) and the new equilibrium (binding emission standard). All welfare changes give the total vehicle lifetime changes for one cohort of new vehicle sales; the changes represent the cost or benefit of having an emission standard for one calendar year. The direct effects of the regulation are the changes in consumer surplus, profits, and gains from correcting externalities. I calculate consumer surplus with the log sum formula of Small and Rosen (1981). I add a step when gaming fools consumers. Gaming 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.[46] I obtain changes in profits from prices, marginal costs and quantities. Emission standards do not result in monetary transfers from firms to the government when every firm complies. To compute the changes in actual 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 on the yearly gains/losses in externalities.[47] 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.[48]

Finally, I compute two additional welfare effects of the policy that were not explicit targets of the policymaker. The regulation changes 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. Together, these are estimated to be more important than the CO_2 externality. They report an externality of €0.12 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.[49] I report the gains from shrinking the size of the vehicle market using this €0.12 per kilometer number. Note that this is very optimistic, as the emission standard only targets

[46] Avoided fuel taxes are a transfer from the government to consumers of which the welfare change depends on the marginal cost of public funds. I count reduced fuel expenses as a welfare gain.

[47] Yearly mileage and vehicle lifetime match statistics reported by Eurostat.

[48] This number comes from the Interagency Working Group on Social Cost of Greenhouse Gases (2016).

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

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. Emission standards can be a more actual 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.[50] For the model with endogenous characteristics, I find that consumers are more price-sensitive and value fuel costs less. At these estimates using the annual mileage and interest rate of 6%, I find that consumers only value a €1 reduction in future fuel costs at €0.42. The emission standard results in consumers purchasing vehicles with lower fuel costs, and I include the future gains consumers obtain. I compute this by changing 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).

Welfare Effects Table 5 Column I shows the central welfare estimates for the EU emission standard.[51] The first panel gives changes in the size of the fleet and actual emissions. The second panel shows the direct welfare effects: consumer surplus, profits, and value of actual emissions. The third panel gives the indirect welfare effects of other externalities and undervaluation. The different scenarios I-VII change the design of the regulation or the abatement strategies. In Columns I-III, I assume that 70% of official emission reductions are due to gaming, the standard is attribute-based, consumers are sophisticated and do not respond to gaming, and technology increases the marginal costs through the engineering cost function starting at $\hat{c}'_{t=0}$. Column I solves for optimal firm strategy so that firms can employ mix-shifting and technology (and gaming). Column II restricts the firm abatement strategy to technology adoption, and Column III restricts it to mix-shifting, both Columns solve for optimal prices, but only in Column III does the regulation enter the pricing first-order condition.

Columns I and II are almost equivalent, showing that firms choose to abate almost entirely by technology adoption and gaming.[52] In the baseline estimate of Column I, we see that the emission standard reduces sales by 1% and actual emissions by 4.9%. The standard did not change the share of small vehicles (defined as the compact and subcompact vehicle classes). If firms had responded by mix-shifting (Column III), we would see stark decreases in sales and actual emissions and a stark increase in the market share for small vehicles. The difference between Column II and III illustrates that the abatement strategy chosen by firms is crucial for the market outcomes of the

[50] 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)

[51] The table presents the simulation outcomes at the estimated parameters and a 90% confidence interval. The C.I. is computed by taking 50 draws from the estimated parameter variance-covariance matrix (assuming a joint normal distribution). For each draw of parameters, the supply side is re-estimated and new simulation outcomes computed. I compute the bootstrap from the differences between these 50 outcomes and the outcome at the mean parameters.

[52] 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 minimal amount of mix-shifting will always be efficient.

standard.

The second panel of Table 5 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, respectively. Profit losses originate from firms taking up technology beyond the marginal benefit in the absence of the policy. 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 gaming. The actual 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 €265. This number is much higher than the estimated levels of the social cost of carbon. The effect on consumer surplus and profits is even more negative with mix-shifting. Because fuel taxes are already correcting for more than the carbon externality in the EU, the policy overshoots and requires reductions in emissions that are very costly.[53]

However, the policy has two potential indirect benefits that I report in the final panel of Table 5. First, a decrease in sales reduces other externalities from traffic. Given the assumption of €0.12 per km, the reduction in traffic increases welfare by €2.2 billion. This number is substantial, and as explained above, is most likely an upper bound. Second, correcting undervaluation reduces consumer surplus losses significantly by €1.5 billion. Because consumers undervalue fuel consumption in the RC Logit model with endogenous attributes, the reduced future fuel expenses of the new choices benefit consumers in the future more than they value today. Both these indirect welfare effects need to be added in full to find positive welfare effects.[54]

Attribute basing In the baseline result, I find that firms choose to abate almost exclusively by lowering the official emission rates. The reason for this choice is the attribute basing of the standard. Table 5 Column IV gives the welfare effects of compliance to a flat standard. This scenario leads to larger price changes and more actual emission savings. Why does a flat standard result in more mix-shifting? Because of the slope in the target, more low weight products are now above the target, as illustrated for Fiat in Figure 2. The Figure plots all the products in the fleet, scaled by sales, in the emission-vehicle weight space, and shows that many products in the lower left are below the red flat target but above the green sloped target. For all these products, the policy is an implicit subsidy under a flat standard but an implicit tax for the attribute-based standard. The attribute basing makes mix-shifting much more costly because the firm has much fewer products to which it can shift sales.

[53] See Parry, Heine, Lis, and Li (2014) for a comparison between EU fuel taxes and externalities. In Appendix Table A7, I show that a tax on carbon emissions of vehicles, chosen so that it reduces emissions with the same amount as the standard, would also cost €344 per ton of CO_2. However, the tax would increase overall welfare by €5 billion because of savings in other externalities. The costs per ton of carbon across scenarios are among the highest of alternative policies studied in the US, see Gillingham and Stock (2018).

[54] Appendix Table A7 shows the simulation results with the estimates of the RC Logit with only price endogeneity. The results imply a higher consumer valuation for fuel costs so that the simulation starts 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. The estimated overvaluation causes an additional loss for consumers.

Table 5: Simulation Outcomes

	I	II	III	IV	V	VI	VII
	Opt.	Tech	Mix Shift	Flat	Fooled	Enforce	Engin.
Solve For:	p, t, λ	$p, t, (\lambda = 0)$	$p, f, (t = 0)$	p, t, λ	p, t, λ	p, t, λ	p, t, λ
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]
	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	265	297	458	296	292	173	212
	[104,711]	[114,821]	[334,546]	[57,571]	[102,887]	[67,362]	[195,2414]
	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]

Table gives the welfare effects of policy simulations with 90% 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 (prices are the solution of (7) with $\lambda = 0$). Column III only allows for price changes. 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.

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, 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 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."* Germany proposed a steeper target with a slope of $a = 0.06$ instead of 0.04 that would require lower compliance effort from German firms. The attribute basing was the result of a political agreement between car producing countries.

In Table 6, 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 I restrict the firms to mix-shift. The right panel does the same for the flat target. When I restrict firms to mix-shift, the shadow costs of the regulation increases with attribute basing. On average, the shadow cost of the regulation triples from 0.29 to 0.94 when introducing 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 mix-shifting in response to the ABR is so costly for Fiat (and for Asian firms), it abates emissions mostly by technology and gaming. The implications for market equilibrium are substantial 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 mix-shifting in Column III to almost exclusively technology and gaming in Column I.

The changes in profits in Table 6 show that French and Italian firms benefit from a flat standard, while all the compliance costs fall on German firms. This result is in line with positions the countries took when bargaining over regulation. The policy debate in 2007 focused mainly on these distributional issues and not on the effect of the slope on the likelihood of different abatement strategies and 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 would face similar compliance efforts. However, this design made mix-shifting 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. Member states contravened their legal obligation to monitor and enforce defeat devices. France, Germany, and Italy had evidence that emission control systems did not focus on actual emissions. The report states that these countries did not take steps to

Figure 2: Position of the Fiat fleet relative to flat and attribute-based target

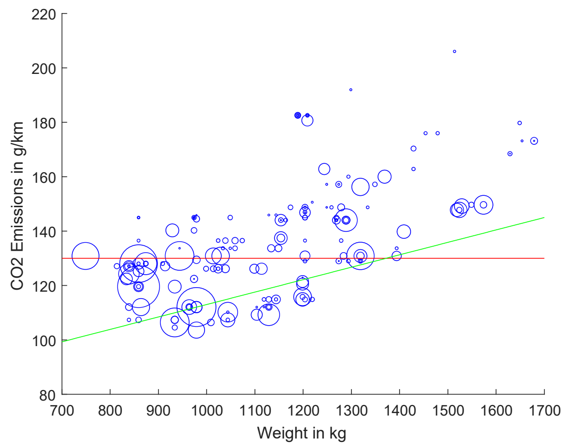

Figure plots every vehicle sold by Fiat in 2007 as it appears in the base scenario in the counterfactual. Each blue circle is a product and circles scale with sales. The red line is the flat target at 130 grams of CO_2 per km and the green line is the attribute-based target.

Table 6: Profits and Stringency per Producing State

	ABR					Flat				
	I			II		III			IV	
Solve for:	t, λ			λ		t, λ'			λ'	
	Δ Prof.	t	λ	Δ Prof.	λ	Δ Prof.	t	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. Solution I solves for equilibrium abatement to the ABR, II allows only for mix-shifting. Solution III solves for equilibrium abatement to the flat standard, IV allows only for mix-shifting.

understand the performance gap between official and actual emission, indicating maladministration. The report also blames member states for under-funding testing facilities (in practice, carmakers 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 a dissuasive penalty system. Additionally, the European Commission failed to oversee the enforcement of member states. In summary, the car-producing countries were aware of the gaming but failed to enforce the regulation. The 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 5 shed light on the economic consequences of weak enforcement. The difference between Columns I and V is the consumer's awareness of gaming. In Column I, the gaming does not affect consumer choice, while it does in Column V. When gaming fools consumers, they perceive cars to have lower fuel consumption, and discover this to be wrong while driving. This choice distortion to consumers causes 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 the standard causes almost no reductions in sales and actual emissions. Overall, the regulation now has a negative welfare effect, even when accounting for the indirect effects. Column VI shows the welfare effects under a standard with full enforcement. Enforcement increases private losses in consumer surplus and profits but leads to much higher CO_2 and other externality savings. Column VI is the simulation that shows the highest welfare numbers when including the indirect welfare effects, but the EU failed to attain this. Finally, if abatement was 100% gaming, there would be no changes in the static equilibrium with sophisticated consumers. Carmakers would report different official 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. The attribute basing increased the cost of mix-shifting as an abatement strategy and increased the likelihood of compliance by gaming and technology adoption. Next, enforcement failures on the level of the member states enabled firms to resort to gaming.

Cost function and investment inefficiencies In the main simulation, the costs for technology adoption start at the 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-actual technology for which consumers have a willingness to pay.[55] Additionally, Jaffe, Newell, and Stavins (2005) point to market failures in the supply and adoption of technology such as spillovers in technology, spillovers in adoption, and incomplete information about future returns of the investment. The

[55] 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.

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. It might have succeeded in moving the industry out of a sub-optimal equilibrium by inducing technology adoption. The framework here allows to compute welfare effects of the emission standard when the supply-side undervalues technology.[56]

By comparing Column VI and VII in Table 5, I find that supply-side failures in technology adoption are not necessarily bad for welfare when externalities are at play. Because of the inexpensive technology, the regulation makes firms reduce fuel consumption for less than the consumers' willingness to pay. Vehicles become cheaper, as consumers receive better characteristics for a price below their willingness to pay. As such, the market size increases rather than decreases. The regulation now has a rebound effect on the extensive margin.[57] The inexpensive 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. The regulation pushes firms beyond what could be explained by market failures, a policy that only requires cost-efficient technology has a less stringet target than existing policy. In general, I believe this scenario is not very credible. In line with Anderson and Sallee (2011), I expect carmakers to abate emissions with the least costly compliance strategies. If technology is so cheap and would have cost only €190 million in variable profits, then why would firms have resorted to gaming?

In Appendix Table A7 Column II, I show the results with the estimated cost function specified in (12). The results are similar to the main scenario, but the estimated cost function is less convex than the engineering cost function and results in lower consumer and profit losses.

Sunk Costs The profit changes reported in Table 5 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 stated that the technology to attain emission reductions was available at the time of the policy. The policy allowed for 8 years between announcement and enforcement. Vehicles typically go through faster redesign cycles, see Blonigen, Knittel, and Soderbery (2017). The additional sunk costs of redesigning are thus likely to be small. Because of the availability of technology and no additional redesign cycles, I 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

[56]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.

[57]See Gillingham, Kotchen, Rapson, and Wagner (2013) for an overview of the possible sources of rebound effects. A second rebound effect 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. Lastly, a decrease in the demand for fuels might lower the price of oil, causing a macro-economic rebound effect. Here, I only focus on the rebound effect on the extensive margin. The reported emission savings are an upper bound on the total savings.

be purchased or designed. Firms also face the risk of liability for noncompliance and class actions from consumers and shareholders. A complicating factor is that firms use similar defeat devices to avoid pollution standards so that the legal cases are both about misrepresenting pollutants and fuel economy. It is unclear if I 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.[58] Appendix A6 discusses the computation of the bounds. I find that the sum of the upper bounds for all firms equals €70 billion. The findings show that firms are willing to sink high amounts to comply with gaming and technology adoption. It also means that the modest positive welfare numbers presented in Table 5 are not sufficient to claim that the regulation had any positive effect. Carmakers are currently defendants in multiple legal cases, the efforts and costs related to these cases are an additional welfare loss so that the regulation has at best zero effect on welfare.

7 Conclusion

This paper evaluates the response to the 2015 EU-wide emission standard. I find that between 2007 and 2011, official sales-weighted emissions from new sales decreased by 14%. A decomposition shows that two-thirds of the decrease in emissions is attributable to firms gaming emission tests. One-third of the decrease stems from actual technology adoption. The projected carbon savings did not materialize because of noncompliance. The reasons that firms chose gaming and technology adoption are attribute basing and a lack of enforcement. Both reasons are a product of the EU political environment. A structural model of demand and supply, allowing for endogenous abatement choices, reveals that the overall effect of the regulation has been negative for consumers and producers. For the regulation to be welfare-improving, non-targeted externalities, such as accidents and local pollution, and corrections in undervaluation need to be counted.

This exercise provides important lessons for studying emission standards. First, the abatement strategy that firms use matters for policy outcomes. Policy design has large impacts on which strategies firms employ. I show that attribute basing made mix-shifting so costly that some firms had to decrease the emission ratings. Second, politics matter for the design, implementation, and enforcement of standards. Third, technology adoption is a crucial mechanism of compliance and needs consideration when evaluating emission standards. This is difficult because the evaluation depends on assumptions about the cost curve of technology adoption. The economic model in this paper estimates technology adoption costs to be higher than projections of engineers, even after accounting for consumers' undervaluation of fuel consumption reductions. It is crucial to further our understanding of why firms might face high costs when adopting greener technologies.

Overall, I want to stress that emission standards are a risky and unpredictable policy instrument

[58] In this setting, I can compute deviations from the single equilibrium with optimal compliance strategies. From observing multiple compliance choices in multiple markets, I 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.

when the goal is to reduce carbon emissions. Standards rely on emission tests that are gameable when compliance costs become very high. It is uncertain that emissions reduce because of the policy. If the engineers are correct and cost-actual technology is available, then standards might increase the size of the market rather than decrease it, causing an increase in emissions. The EU could introduce 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 for emission standards. The EU has limited fiscal authority to raise taxes. The EU parliament recently approved more stringent emission standards for 2021 to 2030. These standards are so stringent that combustion engines are not capable of reaching the target, which implies an imposed shift to alternative fuels.

The numbers derived in this paper are obtained under various assumptions and limitations. First, I focus only on the sales of new vehicles and assume no effects on prices and vehicle lifetimes in the used car market. It would be interesting to study different effects of gaming, technology adoption, and pricing on this market. Second, all welfare numbers ignore possible rebound effects on driving behavior. Third, I did not include dynamics in the analysis. Fixed costs from the standard 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 results presented here show that the welfare effects from emission standards are far from obvious and that the design of the regulation matters for which abatement strategies firms choose.

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.

Eizenberg, Alon. 2014. "Upstream Innovation and Product Variety in the U.S. Home PC Market." The Review of Economic Studies 813 (3):10031045.

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.

Gillingham, Kenneth and James H. Stock. 2018. "The Cost of Reducing Greenhouse Gas Emissions." Journal of Economic Perspectives 32 (4):53–72. URL https://www.aeaweb.org/articles?id=10.1257/jep.32.4.53.

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

Griffith, Rachel, Lars Nesheim, and Martin O'Connell. 2018. "Income effects and the welfare consequences of tax in differentiated product oligopoly." Quantitative Economics 9 (1):305–341. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/QE583.

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.

Interagency Working Group on Social Cost of Greenhouse Gases, (IWG). 2016. "Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866." Tech. rep., United States Government. URL https://www.epa.gov/sites/production/files/2016-12/documents/sc_co2_tsd_august_2016.pdf.

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.

Parry, I.W.H., D. Heine, E. Lis, and S. Li. 2014. Getting Energy Prices Right: From Principle to Practice. INTERNATIONAL MONETARY FUND. URL https://www.imf.org/en/Publications/Books/Issues/2016/12/31/Getting-Energy-Prices-Right-From-Principle-to-Practice-41345.

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.

Wollmann, Thomas G. 2018. “Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles.” American Economic Review 108 (6):1364–1406. URL http://www.aeaweb.org/articles?id=10.1257/aer.20160863.

Appendix For Online Publication Only

A1 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.

A2 Trade-Off Parameters

Table A2 presents the trade-off parameters η from estimating (2). For Model 1, I find that a 10% increase in horsepower is associated with a 1.8% increase in official 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 for introducing firm-specific trends. In Model 3, the trade-off parameters shrink somewhat, but are reasonably robust. In Model 4 and Model 5, the trade-off parameters shrink further when looking at actual emissions. However, it becomes more difficult to interpret these coefficients as a technical trade-off that firms face. actual 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 diesel 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.

A3 Excluded Instruments

For the demand 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 demand 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).

For the marginal cost specification with endogenous prices, fuel costs, horsepower and weight, the following 13 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);
- 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);

A4 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 5. 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 t, I can replace the upper-bound of λ from infinity to the level of the fine. Column III of Table A7 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 A7 also bounds technology. When firms hit the upper-bound on λ, the algorithm will resort to increasing t to attain compliance. Therefore, I use the solution of t 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 t also has an implicit per vehicle cost that should be lower than the fine. This cost depends on the equilibrium conditions and is harder to solve for (I could, in principle, restate the variable to solve for the per vehicle technology costs rather than percentage reductions).

A5 Cost Functions from Engineering Estimates

The report by TNO (2011) lists several available technologies that carmakers can use to reduce vehicle emissions. Each technological option is ordered in terms of additional marginal costs and expressed as a percentage reduction in emissions. This ordering gives a cost function for the changes in marginal costs associated with percentage reductions in emissions. Define $c_j(t)$ as the increase in marginal cost when emissions of vehicle j reduce by t%. The cost function is a polynomial of order K with base t and coefficients a_k:

$$c_j(t) = \sum_{k=1}^{K} (a_k * t_j^k). \tag{13}$$

Table A1 plots the coefficients a_k for diesel and gasoline engines. I use separate cost functions for diesel and gasoline and rely on the coefficients for medium-sized vehicles. Table 79 in TNO (2011) reports separate coefficients for small, medium, and large vehicles, but the marginal costs only differ economically for technology levels above 25% (a level that is not reached in the simulations). The engineering costs imply the following marginal cost changes: a 5% results in €66 for gasoline engines and €97 for diesel engines, 10% in €157 and €215, 15% in €283 and €387, and 20% in €461 and €620. I use this function to compute the marginal cost changes in each vehicle for each t and also compute the derivative $\partial c_j(t)/\partial t$ to find the increase in the term c'_t when $t > 0$, see (8). In Column VII of Table 5, I use this function without adjustments. In Columns I-VI, I start cost increases at the point where the engineering cost slope equals revealed willingness to pay for fuel consumption in the pre-policy period. I obtain the willingness to pay for fuel consumption reductions by computing the sample average of the solution of (8) for $c'_{t=0}$.

A6 Computation of Sunk Costs

I compute the upper bound of the sunk costs by comparing variable profits in the optimal equilibrium with variable profits in an equilibrium I restrict the firm to mix-shift. All other firms continue to play 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 not to play the pricing 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. I 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 an even more stringent regulation will replace it), I foresee a horizon of 6 years.[59] Computing this upper bound requires solving a new equilibrium for each firm and is costly. The lowest upper bound equals €500 million for Fiat, and the highest is €20 billion for Volkswagen. Appendix Table A8 shows results for all firms and smaller deviations.

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

Additional Figures and Tables

Figure A1: Sales-weighted Characteristics over Time

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

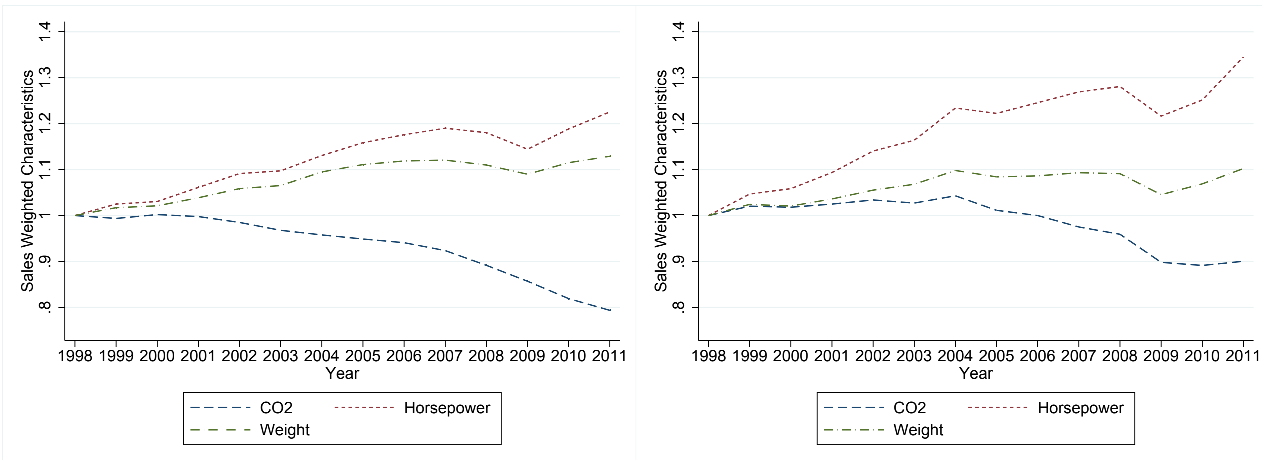

The Figure shows the evolution of sales-weighted characteristics from 1998 until 2011, indexed in 1998. The EU trends are observed in the data. The US trends are from the EPA (http://www.epa.gov/otaq/fetrends.htm).

Table A1: Parameters of the Engineering Cost Function

	a^6	a^5	a^4	a^3	a^2	a^1
Gasoline	1.207E+06	-1.386E+06	5.381E+05	-7.426E+04	9.017E+03	9.985E+02
Diesel		4.147E+05	-3.757E+05	1.308E+05	-9.708E+03	2.151E+03

Coefficients are replicated from TNO (2011) Table 79.

Table A2: 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). The estimated models correspond to Table 2. Standard errors are robust and clustered per firm. In Models 1, 2, and 3, the official $\ln(CO_2)$ rating 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(CO2) estimates as the dependent variable and has year fixed effects. Model 5 has year by firm fixed effects.

Table A3: Robustness for Table A2 and Table 2

	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 A2 and Table 2. 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 A4: 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 A2. 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 A5: 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 A6: 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 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. 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. By definition, the model has a perfect fit within the sample. Out of the estimation sample, in calendar year 2011, I cannot rely on fixed effects and demand unobservables. Prediction based on taste parameters alone shows the demand model explains attributes purchased relatively well, and not worse than the prediction error from leaving out the unobservable and fixed effects within sample.

Table A7: Simulation Outcomes

	I	II	III	IV
	RC Logit I	Estimated Tech	Flat with Fines	Tax
Solve For:	p, t, λ	p, t, λ	p, t, λ	p, t, tax
Gaming:	70%	70%	70%	70%
Consumer Soph.:	1	1	1	1
	Market Size			
Total Sales	-14.78	-0.02	-0.11	-3.85
Emissions	-19.45	-3.76	-10.66	-4.86
	Direct Welfare Effects (Δ in billion €'s)			
Consumer Surplus	-29.67	-0.72	-6.19	-7.12
Profits	-9.50	-0.05	-0.44	-1.94
CO2 Value	1.35	0.26	0.74	0.34
Gov. Revenue:				4.94
Total	-37.82	-0.51	-5.89	-3.79
Implied Value for CO2	814	82	234	344
	Indirect Welfare Effects (Δ in billion €'s)			
Other Externalities	29.66	0.05	0.22	7.83
Paternalism	-7.21	1.15	3.31	1.48
Fines			0.44	
Total:	-15.37	0.68	-1.92	5.53

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 5 but using the estimated cost function. Column III is the same as Column IV of Table 5 but introducing fines as an upperbound for the Lagrangian multiplier. Column IV introduces a tax on carbon emissions of vehicles. The level of the tax is chosen so that the emission savings equal those in Column I of Table 5. See the text for the assumptions behind the welfare calculations.

Table A8: Sunk Cost Upper Bound Estimates

Deviation:	$0.9 * (t+g)$	$0 * (t+g)$
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 $(t + g)$ is restricted to 90% of the optimal $(t + g)$. The second column is the loss in variable profits for each firm when it is restricted to fully comply with pricing while all other firms are responding optimally. Simulation I from Table 5 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%.