



MIT Center for
Energy and Environmental
Policy Research

**Technological Change, Vehicle
Characteristics, and the Opportunity
Costs of Fuel Economy Standards**

Thomas Klier and Joshua Linn

December 2013

CEEPR WP 2014-002

Technological Change, Vehicle Characteristics, and the Opportunity Costs of Fuel Economy Standards

Thomas Klier (Federal Reserve Bank of Chicago)

Joshua Linn (Resources for the Future)¹

December 2013

Abstract

Many countries are tightening passenger vehicle fuel economy standards. In assessing the welfare effects of standards, the literature has not properly accounted either for their effects on the rate of technology adoption, or for improvements in vehicle characteristics in the absence of tightening standards. A dynamic model shows that accounting for both factors has ambiguous effects on estimated welfare costs. We find that recent U.S. and European standards have affected the rate of technology adoption as well as horsepower and torque. Estimated welfare losses from reduced horsepower and torque are of similar magnitude to the welfare gains from fuel savings.

JEL codes: L62, Q4, Q5

Keywords: passenger vehicles, U.S. greenhouse gas emissions rate standards, European carbon dioxide emissions rate standards, technology adoption

¹ We thank Shanjun Li, Virginia McConnell, and seminar participants at Cornell University and the Colorado School of Mines and conference participants at the International Industrial Organization Conference. Wenfei Du provided excellent research assistance. Linn thanks the MIT Center for Energy and Environmental Policy Research and the Swedish Energy Agency for supporting this research.

1 Introduction

Because of concerns about global warming and energy security, many countries have recently adopted policies to substantially increase the average fuel economy of new passenger vehicles. The U.S. Corporate Average Fuel Economy (CAFE) standards for 2016 are about 40 percent higher than 10 years prior. New standards, extending to 2025, may increase fuel economy of new vehicles sold in the U.S. by an additional 50 percent. European standards for carbon dioxide (CO₂) emissions rates (which are inversely related to fuel economy) are scheduled to tighten by about 30 percent from 2012 to 2020. In addition, many other major developed economies, such as Japan, have similar policies, as do some developing countries, such as Mexico and China.

The tightening of standards has coincided with a growing literature on their welfare effects. A central element in this literature has been to model the possible manufacturer responses to stricter standards. The early literature allowed manufacturers to raise fuel economy by changing prices. Producers would lower prices on high fuel economy vehicles and raise prices on low fuel economy vehicles, affecting the sales mix and thereby raising average fuel economy (e.g., Greene 1991; Goldberg 1998). Subsequent analyses, particularly those conducted by the U.S. regulatory agencies—the U.S. Environmental Protection Agency (EPA) and the U.S. Department of Transportation National Highway Traffic Safety Administration (NHTSA)—assume constant market shares by vehicle size class, but allow manufacturers to adopt technology that improves fuel economy while leaving other vehicle characteristics unaffected (e.g., U.S. EPA 2012). Additional research, such as Austin and Dinan (2005) and Jacobsen (2013), allows for both fuel economy improvements and for changes in market shares. Finally, some recent studies incorporate the possibility that manufacturers can change vehicle characteristics, such as horsepower, to comply with rising standards (e.g., Whitefoot et al. 2011; Klier and Linn 2012a; Whitefoot and Skerlos 2012). The literature concludes that fuel economy standards impose much higher welfare costs per gallon of gasoline saved than a gasoline tax. That cost difference gets smaller, however, when incorporating additional manufacturer behavioral margins.

This paper takes a different approach to conceptualize welfare effects of fuel economy standards. By failing to incorporate industry dynamics, the previous literature mis-estimates welfare costs of tighter standards. In particular, the literature has failed to account properly for technology adoption in the absence of tightening standards, and for the effects of tighter

standards on the rate of technology adoption. Examining manufacturers' responses to recent standards in the United States and Europe, we show that the standards increased the rate of adoption. After accounting for improvements in characteristics that would have occurred in the absence of tighter standards, we also find substantial welfare costs from the standards' effects on vehicle characteristics.

Figure 1 motivates our key argument. The figure, which is reproduced from Klier and Linn (2012a), shows that, when U.S. fuel economy standards were constant from about 1985 to 2005, technology improved steadily. Manufacturers in turn used these improvements to raise horsepower and weight while holding fuel economy constant. This pattern suggests that manufacturers continued to improve vehicle characteristics other than fuel economy while standards were unchanged.

The steady technology adoption indicated in Figure 1 suggests that welfare analysis needs to incorporate two dynamic aspects of this industry. First, it must compare equilibria with and without tighter standards, after the standards have been tightened. In the equilibrium without tighter standards, manufacturers adopt technology and improve various characteristics; with tighter standards, manufacturers focus more on improving fuel economy. The previous literature, however, has compared equilibria before and after tightened standards, and therefore cannot account for improved vehicle characteristics in the absence of tightened standards; this leads to an understatement of welfare costs.

Second, tighter standards may encourage manufacturers to innovate and adopt technology more quickly, as suggested by the literature on profit incentives and technology (e.g., Newell et al. 1999). To date, the literature on standards has made ad hoc assumptions on technology costs that determine the effect of standards on the rate of adoption (e.g., Austin and Dinan 2005).

We argue in this paper that the previous literature has omitted two important factors of new vehicles markets: technology adoption that improves characteristics in the absence of tightened standards, and the effects of standards on the rate of adoption. The objective of the paper is to demonstrate—in theory and in practice—the welfare consequences of accounting for these factors. We depart from the structural approach to estimating welfare costs and instead base welfare estimates on observed manufacturer responses to recently tightened standards.

To provide a framework for the statistical analysis, we begin with a simple dynamic model in which manufacturers can adopt technology *and* choose vehicle characteristics. Technology

adoption increases powertrain efficiency, which represents the amount of useful mechanical energy per unit energy contained in the fuel. Following Knittel (2011) and Klier and Linn (2012a), we define a technology frontier. A specific frontier represents a certain powertrain efficiency; moving along the frontier, a manufacturer trades off fuel economy (miles per gallon [mpg]), and vehicle characteristics such as weight and horsepower. Thus, powertrain efficiency is represented by the distance of the frontier from the origin, and a specific point along the frontier establishes the mix of fuel economy and other characteristics.

In the equilibrium without tightened fuel economy standards, manufacturers adopt technology to improve efficiency. Because consumers prefer horsepower improvements to proportional fuel economy improvements (Klier and Linn 2012a), most of the technology adoption is used to increase horsepower. A tightening of standards, however, causes manufacturers to increase the rate of technology adoption *and* to move along the frontier, so that more of the adoption is used to increase fuel economy than in the absence of the tightened standards. We refer to the resulting movement along the frontier as a change in the *direction* of technology adoption. Because the tightening of fuel economy standards affects the direction of technology adoption, the resulting mix of vehicle characteristics is different from that when standards are not tightened. Reductions in characteristics other than fuel economy represent an *opportunity cost* of the tighter fuel economy standards.

This analysis demonstrates two reasons why previous welfare estimates of tightened standards are incorrect. First, by holding all characteristics at the pre-standard level in the no-policy counterfactual, previous studies understate welfare costs. Second, the literature imposes essentially assumes the effects of standards on the rate of adoption.

The remainder of the paper focuses on the following two questions: (a) Have recent fuel economy standards affected the rate of technology adoption and vehicle characteristics other than fuel economy? (b) What are the welfare consequences of those changes? First, we use detailed engine and vehicle characteristics data to estimate technical tradeoffs among fuel economy and other characteristics. We estimate these tradeoffs separately for the U.S. and European vehicle markets. This analysis builds on Knittel (2011) and Klier and Linn (2012a), both of which estimate tradeoffs using cross-sectional and time series variation in vehicle characteristics. We extend their analyses by using matched engine and vehicle model production data to distinguish between medium-run and long-run tradeoffs among fuel economy, weight, and power. We make

this distinction because engine design cycles typically last 8–10 years. Technological tradeoffs between fuel economy and other characteristics across design cycles may be different from tradeoffs within design cycles. Failing to distinguish between within-cycle (medium-run) and cross-cycle (long-run) tradeoffs can overstate manufacturers' ability to trade off weight and power for fuel economy in the medium run and understate this ability in the long run. Therefore, the distinction is important for assessing how easily manufacturers can meet a particular standard at a given time. We compare tradeoffs for existing and newly redesigned engines using an approach similar to that of Linn (2008). We further improve on the literature by estimating a separate frontier by engine, model, and model-year, rather than by model-year (Knittel 2011).

We use the estimated frontiers to examine whether recent standards have affected the rate or direction of technology adoption. Knittel (2011) and Klier and Linn (2012a) provide suggestive evidence that the introduction of the CAFE standards in 1978 affected both the rate and direction of technology adoption. Knittel (2011) finds that the rate of adoption was faster in the early 1980s than in later years but does not control for other factors, such as import competition. Klier and Linn (2012a) show that falling weight and horsepower explains about half of the overall fuel economy increase in the early 1980s (see Figure 1), but they do not establish a causal connection between fuel economy standards and weight and horsepower.

This paper analyzes four recent changes in standards in the United States and Europe. The United States tightened fuel economy standards for light trucks in 2003 and for both cars and light trucks between 2007 and 2009. Europe adopted mandatory CO₂ emissions rate standards between 2007 and 2009. This system replaced a voluntary standard, which, incidentally, manufacturers did not meet (Klier and Linn 2012b).

We identify the effect of standards on the rate and direction of technology adoption using the variation in regulatory stringency across manufacturers and over time. This variation allows us to control for other factors that affect technology, the two most important of which are the rising gasoline prices in the mid- to late 2000s and the subsequent recession. The recession affected brand market shares in the United States and Europe and dramatically reduced manufacturer profits (Li et al. 2013; Busse et al. 2013). Both factors would likely encourage consumers to purchase less expensive vehicles with higher fuel economy, which could affect manufacturers' technology choices. However, we report several pieces of evidence that the identification strategy controls for these factors.

Regarding the four cases of tightening standards we find that the change in U.S. light truck standards in 2003 and 2007 affected both the rate and direction of technology adoption. The 2007 U.S. car standards affected the rate of technology adoption, although not as much as for light trucks; the evidence regarding whether the 2007 car standards affected the direction is mixed. The European standards affected the rate of adoption and had a small, but statistically significant, effect on the direction of technology adoption.

Finally, we use the empirical results to estimate the opportunity costs of the standards—that is, the value of characteristics given up for improved fuel economy. We focus on opportunity costs because the previous literature has either assumed them to be zero or has not properly defined the baseline from which opportunity costs should be measured. We estimate the opportunity costs of a hypothetical 10 percent increase in fuel economy for both the United States and Europe. The estimated opportunity costs for U.S. light trucks are similar in magnitude to the value of the improved fuel economy. For U.S. and European cars, we find that opportunity costs are smaller than for light trucks. We also find that comparing equilibria before and after the tightening of standards, which is the comparison made in the previous literature, results in estimated opportunity costs that are close to zero. We conclude that, on balance, the previous literature has significantly understated welfare costs by improperly estimating opportunity costs.

This paper bridges the literature on technology and profit incentives with the literature analyzing vehicle standards. The technology literature has demonstrated that profit and market forces affect product characteristics, but it has not analyzed welfare consequences of such policy-induced changes. For example, Newell et al. (1999) show that characteristics of air conditioners respond to regulatory and market pressures. Popp (2002) and Linn (2008) find that the rates of innovation and technology adoption in the manufacturing sector respond to energy prices. In contrast, the extensive literature on passenger vehicle standards—including our own work—has not considered the welfare consequences of changing vehicle characteristics in a dynamic setting. Also related is the literature on consumer valuation of product characteristics and product design (e.g., Mazzeo et al. 2013 and Sweeting forthcoming). Our approach applies more generally to other industries in which technology choices shape multiple product attributes—such as trucks and many home appliances.

Our paper is most closely related to Knittel (2011), and our paper differs along several dimensions: it (a) focuses on the welfare consequences of incorporating dynamics rather than

focusing on technical feasibility of tightening standards; (b) improves on the frontier estimation; (c) estimates the effects of recent standards on the rate and direction of technology adoption; and (d) quantifies the welfare consequences of failing to account for improved vehicle characteristics in the absence of tightening standards and for the effect of standards on the rate of adoption.

2 A Simple Model of Standards

2.1 Equilibrium in the Absence of a Standard

The market consists of multiple manufacturers. Within the market, we analyze a single manufacturer that sells a single type of vehicle. We include multiple time periods, indexed by t . The set of consumers is large, and their demand depends on the vehicle's price, p_t ; its fuel economy, m_t ; and its horsepower, h_t (to simplify the notation, we omit manufacturer and vehicle subscripts). Quantity demanded, q_t , is $q_t = q(p_t, m_t, h_t)$, where the function is decreasing in p_t and increasing in both m_t and h_t . In the model, horsepower serves as a proxy for power train characteristics that consumers may care about, other than fuel economy, such as engine size, maximum torque, and 0–60 time; we omit vehicle weight from the simple model.

The manufacturer chooses the price of the vehicle as well as its horsepower, fuel economy, and power train efficiency, η_t . The efficiency describes the amount of mechanical energy available from a given amount of fuel. Starting from a particular power train, which has a certain fuel economy, horsepower, and efficiency, the manufacturer can increase fuel economy in two ways. First, the manufacturer can increase fuel economy by decreasing horsepower, as given by $m_t = m(h_t, \eta_t)$. (1)

Fuel economy is decreasing in h_t , which reflects the fact that, for a power train with a given η_t , the manufacturer can design the power train to have a higher horsepower at the expense of fuel economy. For example, the manufacturer can retune the engine. Second, the manufacturer can adopt technologies that increase η_t . For example, starting with a six-cylinder engine with a five-speed automatic transmission, the manufacturer could increase efficiency by replacing the five-speed transmission with a six-speed transmission. Increasing the efficiency raises the cost of producing the vehicle. The marginal cost, c_t , is a function of efficiency, $c_t = c_t(\eta_t)$, where the

first and second derivatives are positive. Note that the function has a time index, the reason for which we discuss below.

We refer to the fuel economy frontier as the maximum fuel economy that can be achieved for a particular horsepower and efficiency. As the manufacturer moves along the frontier and trades off fuel economy for horsepower, marginal costs do not change. Increasing efficiency causes the frontier to shift out, as Figure 2 shows.

The manufacturer's profit maximization problem is

$$\begin{aligned} & \max_{p_t, m_t, h_t, \eta_t} [p_t - c_t(\eta_t)]q_t(p_t, m_t, h_t) \\ \text{s.t. } & m_t = m(h_t, \eta_t). \end{aligned}$$

The manufacturer chooses the price, fuel economy, horsepower, and efficiency subject to the frontier constraint.

After substituting the frontier constraint into the objective function, there are three first-order conditions, for p_t , h_t , and η_t . The first-order condition for price is the standard monopoly markup equation and the first-order condition for η_t yields

$$(p_t - c_t) \frac{\partial q}{\partial m} \frac{\partial m}{\partial \eta} = \frac{\partial c}{\partial \eta} q_t. \quad (2)$$

The left-hand side is the difference between price and marginal costs multiplied by the increase in sales that would arise from raising efficiency. The right-hand side is the increase in marginal costs from raising efficiency multiplied by the number of vehicles sold. Thus, the manufacturer equates the marginal benefit and the marginal cost of raising efficiency.

The first-order condition for h_t yields

$$\frac{\partial q / \partial h}{\partial q / \partial m} = - \frac{\partial m / \partial h}{\partial m}. \quad (3)$$

Equation (3) shows that the manufacturer equates the ratio of the marginal benefit of raising horsepower and fuel economy (in terms of the sales increase) with the technological tradeoff between the two characteristics.

We present the equilibrium graphically in Figure 3. Indifference curves for fuel economy and horsepower represent consumer preferences for those characteristics. Consumers prefer horsepower to fuel economy in the sense that the willingness to pay for an increase in

horsepower is greater than for a proportional increase in fuel economy. The indifference curve plotted in Figure 3 represents the set of points such that consumers have equal utility from the vehicle, holding its price fixed.

The figure shows the equilibrium for time $t = s$. According to equation (3), the manufacturer chooses point X_s to maximize profits at time s such that the slope of the indifference curve is equal to the slope of the technological constraint.

Next, we introduce dynamics. To focus on the welfare consequences of technology adoption, we assume that innovation occurs exogenously over time. Marginal costs associated with producing a vehicle with a particular efficiency, η' , decrease over time. Thus, comparing marginal costs at time s to time $s+1$, $c_s(\eta') > c_{s+1}(\eta')$.

From the first-order condition for efficiency, equation (3), we see that because of innovation, the manufacturer increases the efficiency over time. Figure 3 shows the outward shift of the frontier from time $t = s$ to time $t = s+1$. Nearly all of the efficiency increase is devoted to raising horsepower rather than fuel economy; fuel economy at X_{s+1} is only slightly higher than fuel economy at X_s . The steepness of the indifference curve explains this result, which is consistent with the aggregate patterns in the U.S. market from 1985 to 2005 (Figure 1).

2.2 Equilibrium with a Fuel Economy Standard

Suppose that at time $t = s$, the regulator unexpectedly sets a fuel economy standard of m^* , for all $t > s$. The standard applies at the beginning of the next time period, $t = s+1$; the timing reflects the situation in the United States and elsewhere, in which the standard is announced before it is enforced. Also consistent with recent history, the standard is set above the manufacturer's time $s+1$ fuel economy from the no-policy case.

The regulator introduces flexibility in meeting the standard by allowing manufacturers to trade credits; manufacturers that exceed the standard generate credits in proportion to the amount by which they exceed the standard. Such manufacturers can sell credits to other manufacturers that fall short of the standard. Because of this flexibility, a manufacturer can choose to (a) exactly meet the standard, (b) fall short of the standard and purchase credits from other manufacturers, or (c) exceed the standard and sell credits. Let the market-clearing credit price at

time t be λ_t , which is measured in dollars per mpg per vehicle. We assume that the credit market is perfectly competitive and treat λ_t as exogenous to the manufacturer.

The manufacturer's profit maximization problem is:

$$\max_{p_t, m_t, h_t, \eta_t} [p_t - c_t(\eta_t) - \lambda_t(m^* - m_t)]q_t(p_t, m_t, h_t)$$

s.t. $m_t = m(h_t, \eta_t)$.

The credit market price, λ_t , creates an implicit tax or subsidy proportional to the difference between the standard and the vehicle's fuel economy.

The first-order conditions for efficiency and horsepower are:

$$(p_t - c_t - \lambda_t(m^* - m_t)) \frac{\partial q}{\partial m} \frac{\partial m}{\partial \eta} + \lambda_t \frac{\partial m}{\partial \eta} = \frac{\partial c}{\partial \eta} q_t$$

$$\frac{\partial q / \partial h}{\partial q / \partial m} = -\left\{ \frac{\lambda_t q_t}{\partial q / \partial m [p_t - c_t - \lambda_t(m^* - m_t)]} + 1 \right\} \frac{\partial m}{\partial h} \text{ for } t \geq s+1. \quad (4)$$

We first consider the situation in which the profit-maximizing fuel economy is below the standard, so that $m_{s+1} < m^*$ (i.e., the manufacturer elects to purchase credits to comply).

Comparing the first-order conditions in equation (4) with the corresponding equations from the no-policy case, the standard causes the manufacturer to adopt higher powertrain efficiency (η) and then move along that new frontier toward higher fuel economy and lower horsepower.

Figure 4 depicts the equilibriums with and without the standard. With the standard, the manufacturer chooses point X_{s+1}^* , which has higher fuel economy and lower horsepower than the no-policy equilibrium X_{s+1} . Because of credit trading, the equilibrium fuel economy m_{s+1} may differ from the level of the standard. Not shown in the graph is the fact that the price of the vehicle is higher at X_{s+1}^* than in the no-policy equilibrium.

In an alternative case, the manufacturer increases fuel economy enough to exceed the standard and sell excess credits (i.e., $m_{s+1} > m^*$). Compared to the no-policy scenario, the manufacturer increases efficiency more and moves along the frontier toward higher fuel economy. Thus, we observe that in both cases the fuel economy standard affects the direction (movement along the technology frontier) and rate (outward shift of the technology frontier).

We note that this analysis includes several assumptions that simplify the exposition. Most importantly, each manufacturer sells a single type of vehicle and innovation is exogenous. Relaxing both assumptions does not affect the main conclusions.

2.3 Welfare Analysis

The previous literature—including the analysis by the regulatory agencies for the U.S. standards—has not allowed for technology adoption that improves characteristics in the absence of the standards. We now discuss the welfare implications of this assumption.

We briefly summarize the approach used in the analysis by the regulatory agencies. The EPA/NHTSA analysis begins by considering the no-policy equilibrium. Then, it uses a simulation model to estimate the increase in η and the associated costs such that (a) all manufacturers meet the standard in the next time period (subject to upper bounds on available technology and manufacturer costs) and (b) characteristics other than fuel economy do not change from their initial levels. The change in marginal costs is estimated by comparing costs before and after the standards are tightened. An assumed manufacturer markup to translate the production cost increases to price increases. The resulting price increases are used to estimate the change in manufacturer profits and the lost income for vehicle consumers.

By failing to account for technology adoption in the absence of standards, this approach yields incorrect welfare estimates. Figure 1 shows the extent of actual technology adoption between 1985 and 2005— as manufacturers increased horsepower and weight without a change in the fuel economy standard. The EPA/NHTSA analysis does not account for this technology adoption—and resulting consumer welfare improvements. In other words, it fails to account for the outward shift of the indifference curve in the no-policy equilibrium in Figure 3. The EPA/NHTSA comparison also fails to properly account for vehicle price changes in the absence of standards. The proper comparison is between two equilibria in the same time period, i.e., $t = s + 1$, one with and one without the standard.

In short, the EPA/NHTSA approach underestimates welfare costs because it does not account for improved characteristics in the absence of stricter standards. That analysis also assumes that there is no innovation and that manufacturers cannot move along the technology frontier; allowing for either would reduce the estimated consumer welfare costs. Whether the approach, on balance, under- or overestimates overall welfare costs is therefore an open question.

As the introduction notes, none of the previous literature includes dynamics; therefore, these studies hold characteristics other than fuel economy equal to their initial levels in the no-policy case. Thus, they underestimate welfare costs for the same reason as EPA/NHTSA.² Furthermore, in the literature, accounting for the effect of standards on the rate of technology adoption is based on engineering estimates of technology costs, rather than on observed manufacturer behavior (e.g., Austin and Dinan 2005 and Klier and Linn 2012a). In principle, these assumptions could yield welfare cost estimates that are too high or too low. At the end of the paper, we quantify the welfare implications of accounting for a) improved vehicle characteristics in the absence of the standards and b) the effect of standards on the rate of technology adoption.

3 Estimating the Technical Tradeoffs among Vehicle Characteristics

3.1 Empirical Strategy

In this section we estimate the shape of technology frontiers as well as shifts of the frontiers over time using data on U.S. passenger vehicles and European cars. Because the United States has historically regulated fuel economy and Europe has regulated CO₂ emissions rates, we estimate a fuel economy frontier for the United States and an emissions rate frontier for Europe.

We define the location of the fuel economy frontier at year t as the change in log fuel economy between the initial year of the sample and year t . The location is measured along the fuel economy axis (see Figure 2), and represents the hypothetical case in which all efficiency improvements between the initial year and year t were used to increase fuel economy. For a given fuel type, a vehicle's fuel economy and its CO₂ emissions rate are inversely proportional to one another. The location of the emissions rate frontier is defined in a manner analogous to that of the fuel economy frontier: it reports the reduction in the log emissions rate assuming all technology adoption is used to reduce the emissions rate.

A vehicle model version and year define a unique observation in our data. As explained in Section 3.2, the definition of a model version differs between the U.S. and European data, but in both cases the data reflect within-model variation in engines and model trims. Similar to Knittel (2011) and Klier and Linn (2012a), we begin with a simple equation describing the fuel economy or emissions rate as a linear function of horsepower, weight, and other characteristics:

² Austin and Dinan (2005) allow for adoption in the absence of the standards but do not allow for innovation or movement along the frontier. Failing to account for these margins overstates welfare costs.

$$\ln e_{it} = \beta_0 + \beta_h \ln(h_{it}) + \beta_w \ln(w_{it}) + \tau_t + X_{it} \delta + \varepsilon_{it}, \quad (5)$$

where e_{it} is the fuel economy (for the U.S. analysis) or CO₂ emissions rate (for the European analysis) of model version i in year t ; h_{it} and w_{it} are horsepower and weight; τ_t is a set of model-year fixed effects (see Section 3.2 for the definition of a model-year); X_{it} contains a set of vehicle characteristics, including the transmission type, fuel type (gasoline, diesel fuel, or 85 percent ethanol [E85]), and number of engine cylinders; ε_{it} is an error term; and $\beta_0, \beta_h, \beta_w$, and δ are parameters to be estimated.

Equation (5) can be estimated separately for the United States and Europe. For the U.S. analysis, the dependent variable is fuel economy; for the European analysis, the dependent variable is the CO₂ emissions rate.

The coefficients on weight and horsepower capture the tradeoffs among fuel economy/emissions, weight, and horsepower. The coefficients are expected to be negative if the dependent variable is fuel economy and positive if the dependent variable is the emissions rate. If the technology frontiers for European and U.S. vehicles have the same curvature, the coefficients in equation (5) would have the same magnitude but opposite signs.

The model-year fixed effects capture fuel economy increases or emissions rate decreases that are possible without reducing weight or power, and correspond to η_t from the model in Section 2. More precisely, the increase in time fixed effects between two years equals the shift of the frontier away from the origin. Importantly, because we estimate equation (5) by ordinary least squares (OLS), we interpret the frontier shift as the potential change in the *average* log fuel economy across all models. We estimate equation (5) by OLS to maintain consistency with Knittel (2011) and Klier and Linn (2012a).

Equation (5) makes an implicit assumption about the underlying technology: the frontier shifts out proportionately over time. For two reasons, this assumption is unlikely to hold in practice. First, manufacturers may adopt power train technology at different rates. For example, manufacturers may differ in their ability to improve or adopt power train technology between one time period and the next, or they may choose to improve other vehicle attributes, such as safety, instead of power train technology. To allow for these possibilities, we replace the year fixed effects, τ_t , with model by model-year interactions, τ_{mt} . The interactions also address a

concern raised in Whitefoot et al. (2011) about a potential correlation among weight, horsepower, and unobserved model-level characteristics.

Regularities in engine design are the second reason this assumption is unlikely to hold in practice. Engines are produced in well-defined models, several of which are part of a specific engine program (see Section 3.2.1). Manufacturers often provide a single engine program for multiple vehicle models, and many vehicle models have multiple versions that contain engines represented by different programs (Klier and Linn 2012a). Furthermore, manufacturers typically stagger the design cycle for vehicle models and engine programs, so that redesigns are completed for a subset of their models and engines in a particular year. Because of the regular design cycles, the practice of selling an engine program in multiple vehicle models and vice versa, and the staggering of the design cycles, the frontier shift is likely to vary across versions of a model in a particular model-year. Therefore, estimating equation (5) by OLS would likely yield biased estimates of the parameters.

We use engine production data to address the second point. In particular, for each version of a vehicle model, we match the set of engine programs corresponding to engine models sold with that version. The variable, r_{it} , is equal to one if the model version is sold with an engine program that has been redesigned or if the engine program was not previously sold with this version. The final estimating equation is

$$\ln e_{it} = \beta_0 + \beta_h \ln(h_{it}) + \beta_w \ln(w_{it}) + r_{it}\tau_{mt} + X_{it}\delta + \varepsilon_{it}, \quad (6)$$

where $r_{it}\tau_{mt}$ is the interaction of the redesign variable by model and model-year. These interaction terms relax the assumption in equation (5), which stated that the frontier shifts out proportionately over time for all versions of a model.³ We assume that within-model and redesign variation in unobserved characteristics is uncorrelated with observed characteristics.

Several main hypotheses are to be tested using equation (6) for the U.S. analysis, in which the dependent variable is fuel economy. First, the coefficients on weight and horsepower are expected to be negative, reflecting the tradeoffs among fuel economy, weight, and horsepower along the frontier. Second, the interactions of redesign, model, and model-year, which measure the distance between the frontier and the origin, increase over time as manufacturers adopt

³ Vehicle models are also designed at regular intervals, and the model design cycles do not always coincide with the engine design cycles. We focus on engine design cycles because the relationship between fuel economy and other characteristics depends largely on the power train and weight, and not on other vehicle characteristics.

technology that causes the frontier to shift away from the origin. The hypotheses are analogous for the European analysis, in which the dependent variable is the emissions rate.

In summary, equation (6) has several important features. First, we allow the tradeoffs between fuel economy/emissions rates and other characteristics to depend on whether a powertrain has been redesigned. Second, we allow the frontier to shift out by different amounts for each model. Third, and importantly for Section 4, we do not impose assumptions on the effect of the standards on the direction or rate of technology adoption.

3.2 Data

The U.S. data come from several sources. Vehicle sales are from Wards Auto Infobank. Monthly sales data are aggregated to the model by model-year, where a model-year begins in September of the previous calendar year and ends in August of the current year. The vehicle sales data are measured at the vehicle model level. The sales data distinguish different power sources, such as gasoline/diesel, hybrid, and electric. We merged to the sales data other engine characteristics—such as engine displacement, number of cylinders, horsepower, torque, and fuel economy—from Wards annual yearbooks. Those characteristics were measured at the model version level. The characteristics data distinguish diesel fuel from gasoline versions.

Finally, we merge to the Wards data the engine data by model, fuel type, and number of cylinders. The engine data distinguish three levels of engine aggregation: an engine platform combines related engine programs, which may consist of multiple engine models. The data, which originated from IHS Global Insight, allow us to determine when a vehicle is sold with a redesigned engine model and when an engine program is first introduced in a vehicle.⁴

Table 1 provides some summary statistics for the U.S. data for the years 2005 and 2010. The table shows unweighted averages across model versions. There are more than 1,300 observations per year. Between 2005 and 2010, fuel economy increased 6 percent, weight increased 5 percent, and horsepower increased 13 percent. Panel A of Figure 5 shows the trends over the entire sample period, 2000–2012. Horsepower and weight increased steadily in the first half of the

⁴ The production data are worldwide for 2000–2007 and cover North America for 2008–2012. This introduces some measurement error in identifying redesign years for engines that are produced only outside North America but are sold in the United States. On average, about 25 percent of vehicles sold in the United States have engines produced outside North America. Restricting the sample to models with engines produced within North America does not appreciably affect the estimated frontiers; this suggests that any measurement error in the redesign variable does not significantly bias the estimates.

sample and then leveled off (more so for weight than horsepower), whereas fuel economy was constant in the first half and then increased; these patterns foreshadow the results in Section 4.

The European data were obtained from R.L. Polk and cover the years 2005–2010. The data include all new cars sold in Sweden and the countries with the eight largest markets in Europe: Austria, Belgium, France, Germany, Italy, the Netherlands, Spain, and the United Kingdom. Observations are by country, year, and model version, where a version denotes a unique model name, model trim, number of doors, engine displacement, horsepower, transmission type (manual or automatic), and fuel type (gasoline or diesel fuel). We pool data across European countries so that the final data set contains about 47,000 observations per year. Thus, a model version in the European data is much more disaggregated than in the U.S. data. A European model-year corresponds to a calendar year (Klier and Linn 2013).

Table 1 reports summary statistics for the European data for comparison with the U.S. data. Fuel economy is much lower and horsepower is much higher in the United States than in Europe. The reported weight is larger in Europe, but that is because the European data include the gross vehicle weight, and the U.S. data include the curb weight (gross vehicle weight includes the weight of passengers and cargo, which curb weight excludes). The table also shows that fuel economy increased nearly twice as much (in percentage terms) in Europe as in the United States, whereas increases in weight and horsepower were about the same. Panel B of Figure 5 shows that horsepower, weight, and fuel economy increased in the first half of the sample, but in the second half fuel economy increased more quickly while weight and horsepower were flat overall.

3.3 Estimation Results

Table 2 shows the estimates of equation (6) for the United States, with column 1 showing results for cars and column 2 for light trucks. We could include horsepower and torque in all regressions, but in practice they are extremely highly correlated with one another. Our regressions for U.S. and European cars use horsepower; our regressions for U.S. light trucks use torque, which, for light trucks, is more highly correlated with fuel economy than is horsepower.

Fuel economy, horsepower, and weight are in logs, and the reported horsepower and weight coefficients represent elasticities. The regressions include dummy variables for whether the vehicle uses diesel fuel, has a hybrid power train, is a flex-fuel vehicle (capable of using E85), or has a manual transmission; the coefficients on the indicator variables approximately equal the percentage change in fuel economy associated with having these characteristics. Besides the

reported variables, regressions include fixed effects for the number of cylinders and doors and interactions of redesign, model, and model-year.

The estimates in column 1 suggest that a 1 percent increase in horsepower decreases log fuel economy by about 0.24, which is significant at the 1 percent level. The estimate is significantly larger than Klier and Linn (2012a) because the latter focuses on within-engine program variation, whereas these estimates reflect both cross-engine and within-engine program variation. The weight coefficient in column 1 is smaller than Klier and Linn (2012a) for the same reason. The horsepower and weight coefficients also differ from Knittel (2011), but the sample periods and data sources differ as well.

The diesel fuel coefficient implies that the log fuel economy of diesel fuel cars is about 0.34 larger than gasoline-powered vehicles. The coefficient on the manual transmission dummy, which is expected to be positive, is in fact negative, but it is quite small and is not statistically significant. The coefficient on the hybrid power train dummy indicates that the log fuel economy of hybrid cars is about 0.26 higher than comparable gasoline-powered vehicles.

Compared to cars, the light truck estimate for the torque coefficient is smaller than the horsepower coefficient, and the estimate for the weight coefficient is larger in magnitude. The light truck and car hybrid coefficients are essentially the same. The coefficient on flex-fuel vehicles is negative, reflecting the lower energy content of E85 compared to gasoline.

The differences between the coefficients for cars and light trucks motivate our estimation of a separate frontier for the two vehicle categories. Appendix Table 1 separates the categories further, reporting results by market segment. Cars have three market segments (small, medium, and large/luxury), and light trucks have four segments (crossovers, sport utility vehicles, vans, and pickup trucks). Coefficients vary substantially across segments; for example, weight and horsepower have larger effects on fuel economy for small cars than for other car segments.

Table 3 reports results for Europe. Because the dependent variable is the emissions rate rather than fuel economy, the signs of the coefficients are opposite from the corresponding U.S. coefficients. Besides the reported variables, column 1 includes fixed effects for the number of engine cylinders and interactions of redesign, model, and model-year. The European regressions do not include vehicles with hybrid power trains or vehicles that use flex fuel, but column 1 is otherwise comparable to the U.S. car regression.

The European regressions include only passenger cars and we compare the European results with the U.S. car results. The magnitudes of the European horsepower and weight coefficients are very similar to those of the U.S. estimates. The European diesel fuel coefficient is smaller than the U.S. coefficient (in magnitude), but this is because diesel fuel has a higher carbon content than gasoline; if we use fuel economy rather than the emissions rate as the independent variable for the European regressions, the magnitude of the European diesel fuel coefficient is very similar to that of the U.S. coefficient.

A model trim is defined as a unique model name, body type, number of doors, driven wheels, and trim level; different model trims may have different engine models. The greater disaggregation of the European data allows us to estimate a separate frontier for each model trim. For consistency with the U.S. analysis, we focus below on the estimates using redesign by model and model-year interactions, but column 2 of Table 3 reports the redesign by model trim and year results for comparison. The coefficient estimates are quite similar in columns 1 and 2 of Table 3. Appendix Table 2, which reports separate regressions by car market segment, shows that the coefficients vary somewhat across segments, but less so than in the U.S. segment-level regressions in Appendix Table 1.

4 Have Standards Affected the Direction and Rate of Adoption?

In this section, we use the estimates of equation (6) to investigate whether the recent U.S. and European standards affected the rate and direction of technology adoption. We first report qualitative aggregate results followed by quantitative cross-sectional results, in which we control for potentially confounding factors.

4.1 Hypotheses for Aggregate Direction and Rate

We consider whether the average rate or direction of technology adoption changed after the standards were first adopted. We define the *rate* of fuel economy technology adoption in a particular year as the change between the current and previous years in the market-wide average estimate of $r_{it}\tau_m$ from equation (6). The change represents the increase in average log fuel economy, relative to the previous year, if all of the adopted technology were used to increase fuel economy—that is, if manufacturers held fixed other vehicle characteristics. We define the *direction* of technology adoption as the log of the ratio of fuel economy to horsepower or weight (i.e., there are two direction variables).

In the aggregate analysis, we do not attempt to control for potentially confounding factors that affect rate and direction. Instead, we ask simply whether the average rate and direction changed after the standards changed. We consider the U.S. light truck fuel economy standards adopted in 2003, the U.S. car and light truck fuel economy standards adopted in 2007 (and tightened in 2009), and the European CO₂ emissions rate standards adopted in 2007 (and finalized in 2009). In each case we ask whether the average rate and direction of technology adoption changed after the standards were adopted. Note that we look for changes after the standards were adopted rather than when they first had to be met, which is usually two to three years after adoption. In the context of the vehicle and engine design cycles noted above, we would expect manufacturers to make changes as soon as the standards have been adopted.

4.2 Aggregate Results

Figure 6 shows the aggregate results for the United States and Europe. Vertical lines indicate the adoption years of the standards. The solid black curve is the cumulative frontier shift since the year 2000. The curve indicates that the average fuel economy of U.S. cars would have been 12 percent higher in the year 2010 than in 2000 if all new technology had been used to raise fuel economy and if all other vehicle characteristics had remained unchanged from their 2000 levels. The red line is the change in actual average fuel economy compared to the year 2000. The other lines in the figures are the counterfactual changes in fuel economy that would have occurred had the corresponding characteristic been held fixed and the frontier not shifted; that is, they represent the fuel economy increase by moving along the frontier. For example, the horsepower curve indicates that if horsepower had been held fixed from 2000 to 2004, cars would have had about 3 percent higher fuel economy in 2004 than they actually did. The curve is computed using the actual horsepower change and the horsepower coefficient reported in Table 2. By construction, in the figure the sum of the change in characteristics is equal to the frontier shift—that is, the estimated model-redesign fixed effect.

The figure shows that the average rate and (in most cases) the direction changed soon after the standards changed. Regarding the rate, for U.S. cars (Panel A), the frontier shifted out twice as quickly from 2008–2012 as compared to 2000–2007. For U.S. light trucks (Panel B), the frontier shifted out twice as quickly after 2003 as compared to 2000–2003. The earlier timing for the light trucks is consistent with the fact that the light truck standards tightened before the car

standards. For European cars, the frontier also shifted out more quickly after 2007 compared to 2005–2007.

There is also clear evidence that the direction changed, particularly for U.S. cars and light trucks. Until about 2007, average car fuel economy was flat, as manufacturers used the outward shifts of the frontier to improve other characteristics, particularly horsepower. After 2007, on the other hand, fuel economy began increasing at about the same rate as the frontier. The pattern is similar for light trucks; fuel economy was roughly flat until about 2004, after which it began increasing.

Figure 6 shows the market-wide average patterns, and Figures 7–9 provide company or brand-level detail. The figures are constructed similarly to Figure 6, except that each panel represents a different company (in the United States) or brand (in Europe).⁵ The figures illustrate considerable cross-firm heterogeneity in the rate and direction of technology adoption, but most firms exhibit similar patterns to those shown in Figure 6.

4.3 Hypotheses for Cross-Sectional Rate and Direction

Although the aggregate results suggest that the standards affected the rate and direction of technology adoption, there may be confounding influences. For example, gasoline prices began rising in 2003. Given vehicle design lags of three years or more, rising gasoline prices may have affected the rate and direction of adoption as early as 2006. We next discuss our approach to control for such potential confounding effects.

The main feature of our identification strategy is that we exploit cross-sectional variation in the stringency of the standards. Although the adoption of each of the four standards (U.S. light trucks in 2003, U.S. cars and light trucks in 2007, and European cars in 2007) affects the entire market, the incentives for changing the direction and rate of technology adoption vary across manufacturers, depending on how close they are to achieving the new standard.

This strategy would seem to run counter to the ability of manufacturers to trade credits to meet compliance (see section 2). While credit trading simplified the exposition of our model, in the US market it has only been possible since 2011 and no cross-firm trades have been observed to date. If we drop credit trading from the model in Section 2, first-order conditions analogous to

⁵ For Ford, General Motors, and Nissan, fuel economy dropped noticeably in 2010. Starting in 2010, the Wards fuel economy for flex-fuel vehicles corresponds to the fuel economy using 85 percent ethanol rather than gasoline. The flex-fuel indicator variable in equation (6) controls for this change when we estimate the frontier. For that reason, the company frontiers did not shift toward the origin when the reporting change occurred.

equation (2) show that the standards cause greater changes in rate and direction for manufacturers whose vehicles, prior to the standards, had fuel economy further below the standards than for other manufacturers.

Given the absence of trading, we define the *stringency* of standards as the difference between a manufacturer's pre-standard fuel economy or emissions rate and the level of the new standard. Stringency can vary across manufacturers because of different levels of pre-standard fuel economy or because of differences in the mix of vehicles offered. In both the United States and Europe, standards are administered based on the physical characteristics of vehicles. The U.S. standards for cars and light trucks are calculated based on vehicle footprint, which roughly corresponds to the rectangle defined by the four wheels. Larger vehicles are subject to lower fuel economy standards. The European standards are based on weight, such that lighter vehicles are subject to lower emissions rate standards. Accordingly, in a footprint-based system, manufacturers with larger vehicles are subject to lower fuel economy standards, and similarly for the European weight-based system.

To identify the effects of the standards on the rate and direction of technology adoption, we use a differences-in-differences framework. We estimate separate regressions for U.S. cars, U.S. light trucks, and European cars. For the rate of technology adoption, each regression is a variation of the following equation:

$$r_{it} \hat{\tau}_{m} = \gamma_R S_F Post_t + Seg_m Post_t \ln(e_m) + \theta_t + \omega_m + \nu_{it}. \quad (7)$$

Observations are by redesign, model, and model-year; that is, there are two observations for a model that was redesigned in a particular year. The dependent variable is the redesign by model and model-year interaction term estimated in equation (6). The variable $Post_t$ is a dummy variable equal to one after the standard has been adopted (e.g., post-2007 for Europe), and S_F measures stringency by manufacturer, F . The variable is the difference between the log of the manufacturer's average fuel economy or emissions rate in the first year of the sample and the log of the manufacturer's standard; γ_R is the coefficient on the interaction of $Post_t$ with S_F . For the U.S. light truck standards, we allow for the possibility that the 2003 and 2007 standards differed from one another in their effects on the direction and rate of adoption and estimate a separate γ_R for each time period. The term $Seg_m Post_t \ln(e_m)$ represents the triple interaction of market segment fixed effects with $Post_t$ and the log of the average fuel economy or emissions rate of

the corresponding model in the initial year of the sample. Note that when estimating equation (7), we include all lower-order terms for the triple interaction; we omit these terms in the expression for brevity. Later in the section, we discuss how the triple interactions address concerns about gasoline prices and other possibly confounding factors.

Equation (7) includes both year fixed effects (θ_t) and model fixed effects (ω_m). The year fixed effects control for the average level of the frontier each year and for any unobserved factors that affect the dependent variable proportionately. The vehicle fixed effects control for the average frontier of the corresponding model over the sample. Because of the presence of vehicle fixed effects, a vehicle's frontier shift is measured relative to its average frontier.

To illustrate the differences-in-differences interpretation of γ_R , suppose the average frontier for manufacturer A shifts out at the same rate as the frontier for manufacturer B prior to the adoption of tighter light truck standards. Assume further that the stringency variable is more negative for A than for B, meaning the standard is more stringent for manufacturer A. The coefficient γ_R is negative if the average frontier for A shifts out more quickly than the average frontier for B after the light truck standards were adopted. Note that the approach cannot distinguish between a case in which the standards caused a one-time outward frontier shift and a case in which the standards caused the frontier to shift out at a faster rate for multiple years. Either case would result in a negative coefficient, but we lack enough years of post-standards data to distinguish them.

There are two main issues that threaten the identification of γ_R . First, unobserved manufacturer or model-level characteristics may be correlated with the stringency-time period variable, $S_F * Post_t$. The model fixed effects partially mitigate this concern by controlling for time-invariant manufacturer heterogeneity. For example, the estimates are unbiased if the difference between the fuel economy of General Motors' and Toyota's cars that existed prior to 2000 would have persisted after 2000 in the absence of the new standards. This assumption cannot be tested directly, but it is supported by the long period of time, prior to 2005, during which the standards were constant and manufacturers' relative fuel economy was quite stable (Jacobsen 2013). Constructing S_F from the vehicle fuel economies at the beginning of the sample further mitigates the first concern because the new standards do not directly affect this variable.

The second potential concern is that other factors, such as fuel prices or the recession, may also affect incentives for technology adoption. The triple interaction in equation (7) controls for such factors to the extent that they are common within a market segment or are proportional to the vehicle's initial fuel economy. The underlying assumption is that, after including the triple interaction, the stringency is uncorrelated with the effects of the recession and gasoline prices. Section 4.5 documents strong evidence supporting this assumption.

Next, we turn to the direction of technology adoption. We define a set of direction variables, dir , at the vehicle level. The fuel economy–horsepower direction, for example, is the log of the ratio of fuel economy to horsepower. Direction variables for fuel economy–torque and fuel economy–weight are defined similarly. The hypothesis to be tested is that an increase in the stringency of the fuel economy standard causes the direction to shift to fuel economy and away from torque, horsepower, and weight. We estimate the equation

$$dir_{it} = \gamma_D S_F Post_t + Seg_m Post_t \ln(e_m) + \theta_t + \omega_m + \nu_{it}. \quad (8)$$

For the United States we estimate four regressions: two for cars and two for light trucks, where the dependent variables for the car regressions are the horsepower and weight direction variables, and the dependent variables for the light truck regressions are the torque and weight direction variables. For Europe we estimate two regressions, for the horsepower and weight direction variables. Observations are by model version and year.

The interpretation of the coefficient γ_D is similar to that of γ_R . For horsepower, for example, the coefficient is negative if manufacturers with a lower value of S_F shift toward higher fuel economy and away from horsepower, and if this change is greater for manufacturers for which the standard is more stringent. Thus, a negative coefficient suggests that the standards cause manufacturers to change the direction toward fuel economy. For Europe, the coefficient γ_D is negative for the horsepower regression if manufacturers with a higher initial emissions rate reduce emissions rates at the expense of horsepower more than do other manufacturers.

4.4 Cross-Sectional Results

For the United States, we allow the effects of the standards to vary across four time periods: 2000–2002, in which light truck and car standards were unchanged; 2003–2006, in which higher light truck standards were first adopted; 2007–2009, in which car and light truck standards were adopted; and 2010–2012. The last time period allows for the possibility that manufacturers

responded more strongly as the tighter standards took effect. The key independent variables are the interactions between time period fixed effects. We test whether (a) the direction and rate of technology adoption for light trucks differed between the first time period and the subsequent periods and (b) the direction and rate of adoption for cars differed between the last two periods and the first two periods.

Panel A of Table 4 shows results from estimating equation (8), in which we assess the effect of the standards on the *direction* of technology adoption. Columns 1 and 2 show results for cars, and columns 3 and 4 for light trucks. We find no evidence that the standards affected the direction for cars, but we find strong evidence for light trucks that the standards caused the direction to shift toward fuel economy and away from torque.

To interpret the magnitudes for light truck torque, we consider a manufacturer for which the stringency is one standard deviation below the mean in the second time period (i.e., a manufacturer with low initial fuel economy compared to its standard). The size of the estimated effect implies that that manufacturer decreased torque and increased fuel economy 5 percent (about 1 mpg) compared to a manufacturer with mean stringency; given that the average light truck fuel economy increased 3 mpg during the sample period, the movement along the frontier represents a significant fuel economy increase.

Panel B of Table 4 shows results from equation (7), which focuses on the *rate*. The results suggest that the standards increased the rate of adoption for cars in 2010–2012 but not in the earlier periods. The truck results are consistent with the hypothesis that the standards affected the rate of adoption, as companies facing more stringent standards increased their rates of adoption more than other companies in the middle two time periods (see columns 3 and 4); however, the coefficient in the final time period is smaller and is only marginally statistically significant. In comparing the rate and direction estimates, we find them to be larger and more precise for light trucks than for cars. This difference could be explained by the fact that the standards for cars only tightened in the last two years of the sample; the aggregate analysis in Section 4.2 suggested that the rate of adoption increased after 2007, but the statistical evidence suggests that this response was correlated with gasoline prices, the recession, or other factors. We conclude that

the rate and direction changed first for light trucks and then for cars. This timing is consistent with the timing of the U.S. standards.⁶

Interpreting the magnitudes in Panel B, we consider the same hypothetical manufacturer facing stringency one standard deviation below the mean. For cars in 2010–2012, the rate of adoption for this manufacturer is 0.5 percentage points higher than the observed average rate of 1.4 percent per year. The light truck results for 2007–2009 suggest that the manufacturer increases the rate of adoption 1.5 percentage points above the mean of 1.5 percent per year. Thus, the magnitudes imply substantial increases in the rate of adoption.

Table 5 reports the results for Europe. The key independent variable is the interaction of a dummy variable equal to one for 2008–2010 and the difference between the log emissions rate of the manufacturer and the log of the 2015 standard. As with Table 4, Panel A of Table 5 focuses on the direction of technology adoption (equation [8]) and Panel B on the rate (equation [7]). If the coefficients are negative, manufacturers with higher initial emissions rates shift direction toward lower emissions rates and raise the rate of adoption, compared to other manufacturers.

The European standards had a statistically significant effect on horsepower and weight, but the magnitudes of both effects are small. A one standard deviation increase in stringency causes a shift along the frontier that reduces emissions rates by 2 percent (recall that the corresponding estimate for U.S. light trucks was 5 percent). We also find that the rate of adoption increased. The magnitude implies that a one-standard-deviation increase in stringency increases the rate of adoption by 0.3 percentage points, compared to the mean rate of adoption of 2 percent. Thus, the magnitude is noticeable, but smaller than for the United States. We conclude that (a) the European standards had a relatively small, but statistically significant, effect on horsepower, weight, and the rate of adoption.

4.5 Potential Omitted Variables Bias

The fact that from 2003 to 2009 the adoption rate increased for U.S. light trucks but not for U.S. cars supports the validity of our identification strategy; confounding factors would yield spurious results only if they affected light trucks and not cars during that time period. However, gasoline prices and the recession may have differentially affected cars and light trucks; these

⁶ Above we noted that the frontier estimates in equation (6) differ from Knittel (2011). However, using the data from that paper, we find that the rate of technology adoption increased for light trucks after 2003 but not for cars, which is consistent with the results reported in this paper.

factors represent the primary threats to the validity of equations (7) and (8). As noted, we control for these factors using triple interactions of fuel economy by market segment by year. Because the reported regressions include model fixed effects, the main concern would be time-varying shocks that differentially affect vehicles in the same market segment or with the same fuel economy. This section reports four approaches to assess the validity of the research design.

First, in the main regressions, we assume that the stringency variable is exogenous after controlling for segment-level demand and supply shocks. While this approach controls for segment-level shocks that affect technology adoption, if segment shocks happen to be correlated with stringency we would be concerned that subsegment shocks may also be correlated with stringency. We can assess whether segment shocks are correlated with stringency by omitting the triple interactions in equations (7) and (8). Appendix Tables 3 and 4 report the same specifications as in Tables 4 and 5, except that we omit the triple interactions. The magnitudes in the appendix tables are similar to those reported in the main tables, and the qualitative conclusions are the same: there is strong evidence that the standards affected the rate and direction for U.S. light trucks, weaker evidence for the rate and direction for U.S. cars, and evidence that the European standards affected the rate and direction. These results support the assumed exogeneity of the stringency variable.

Second, the estimates would be biased if fuel prices or the recession (or other factors) reduced demand for vehicles with low efficiency sufficiently for them to exit the market. We construct an indicator variable that is equal to one if a vehicle version exits between the current and next year. We regress the exit variable on the fuel price and on total market sales interacted with the version's fuel economy. Changes in total market sales serve as a proxy for the effect of the recession on the aggregate market. (For the European regressions we use the emissions rate and registrations instead of fuel economy and sales.) Importantly, the regressions also include the triple interaction of time period, market segment, and initial model fuel economy. If either the fuel price or market sales interaction is statistically significant, we would be concerned that gasoline prices or the recession cause exit and thereby bias the results. Panel A of Table 6 reports the coefficients on the interactions. None of the interaction terms is large and statistically significant at conventional levels. A one-standard-deviation increase in fuel economy and a one-standard-deviation increase in fuel prices or aggregate sales cause a very small (less than 2 percentage points) change in exit probability.

Third, gasoline prices or the recession could affect technology via market shares. For example, if gasoline prices raise the market share of vehicles with high fuel economy, manufacturers would have greater incentive to adopt technology that raises the fuel economy of those vehicles (Acemoglu et al. 2012). Of particular concern is the possibility that fuel prices or the recession differentially affect market shares of vehicles sold by firms for which the standards are more stringent. In that case, the coefficients on the stringency interactions in Tables 4 and 5 could reflect the effects of fuel prices or the recession on technology. Panel B of Table 6 reports regressions similar to Panel A, except that (a) the dependent variable is the log of sales or registrations rather than the exit indicator and (b) the key independent variables are the interaction of stringency, time period, and either fuel economy or aggregate sales. Statistically significant or large interaction coefficients would raise concerns that the other variables in equations (7) and (8) do not control adequately for the effect of fuel prices or the recession on market shares. Only the fuel price coefficient for European cars is statistically significant, and in that case the point estimate is small; a 20 percent increase in the fuel price (as occurred during the European estimation sample) affects market shares by less than 1 percent.

Finally, we control directly for gasoline prices by adding to equations (7) and (8) the interactions of gasoline prices with the stringency–time period interaction. The main results (not reported but available upon request) are unaffected.

5 The Opportunity Costs of Standards

The empirical results suggest that tighter standards increased the rate of adoption in both the United States and Europe and affected the direction for U.S. light trucks and European cars. In this section, we use these results to estimate the opportunity cost of tightening the standards in either market and to show the welfare implications of failing to account for technology adoption in the absence of the standards.

We focus on opportunity costs because they had not been included in previous analyses of the welfare effects of fuel economy standards. A complete welfare analysis would require a dynamic model of manufacturer technology adoption, the choice of vehicle characteristics, as well as of consumer demand. Because of the many challenges of estimating such a model's parameters, a full welfare analysis of the standards is beyond the scope of this paper.

Instead, to derive rough estimates of opportunity costs, we make some simplifying assumptions regarding manufacturers' responses to hypothetical standards. First, the standards

do not affect vehicle prices or market shares. Klier and Linn (2012a) suggest that, over periods of three to five years, it is less costly to manufacturers to adjust vehicle characteristics than to change vehicle prices and market shares; over the five-year time horizon considered here, the assumption of constant market shares may therefore not be very strong. Our second assumption is that manufacturers increase the fuel economy of all vehicles by the amount the standard requires. These assumptions allow us to focus on opportunity costs while using simulations that do not contain too many moving parts. Because of these simplifying assumptions, we treat the welfare estimates as approximations.

Although we could base the simulations on the actual standards, to compare results across the United States and Europe, we use the same hypothetical standard for the two regions. The initial year for the simulations is 2007, which is roughly the midpoint in both the U.S. and European data sets. The analysis spans five years, 2007–2012, and we estimate the opportunity costs of raising fuel economy 10 percent over that time period; such a rate of increase is greater than that required by the European standards but not as great as the requirement under the U.S. light truck standards.

We first consider a scenario in which standards are unchanged from the 2007 levels; we refer to that as the no-policy scenario. For this scenario, we set the adoption rate equal to the average frontier shift estimated in equation (7) prior to the tightening of the standards. Based on these estimates, we assume that all efficiency improvements in the United States are used to increase horsepower or torque, whereas efficiency improvements in Europe increase fuel economy and horsepower in equal proportions. These assumptions allow us to estimate the 2012 fuel economy and horsepower of every vehicle that was sold in 2007. To simplify the analysis, we assume that the technology adoption does not affect weight.

We consider two scenarios in which the standards raise fuel economy 10 percent. In the first scenario, we assume that the rate of adoption is the same as in the no-policy scenario. In the second, we use the average rate of adoption estimated from equation (6) over the years 2010–2012 for U.S. cars and light trucks, and over the years 2008–2010 for European cars. We compute the movement along the frontier needed to meet the new standards.

Table 7 presents the results from these simulations. The two rows in each panel show results from the two scenarios, for low and high rates of technology adoption. The first column shows the assumed rate of technology adoption. The remaining columns show the results of the

simulations, including the percentage change in horsepower relative to the no-policy case, the consumer willingness-to-pay for the lost horsepower, and the value of the fuel savings.⁷ The willingness to pay for the lost horsepower is computed using the lower estimates in Klier and Linn (2012a) of \$10 per horsepower per ton (in 2007 dollars). This value corresponds to the lower value in the vehicle demand literature (Whitefoot and Skerlos 2012); the opportunity costs reported in Table 7 should therefore be considered conservative estimates.⁸

For the U.S. simulations, the opportunity costs—as measured by the willingness to pay for lost horsepower—are similar to the fuel savings for light trucks, but are smaller for cars.⁹ The costs are about 2-3 times higher in the low-technology case than in the high-technology case. The difference between the two cases demonstrates the importance of accounting for the effect of standards on the rate of adoption when estimating welfare costs. The European opportunity costs are lower than for the United States, but are still sizeable compared to the fuel savings. The European opportunity costs are lower because, without tightened standards, the European technology adoption rate is higher and a higher fraction of U.S. technology adoption is devoted to raising horsepower.

Section 2.3 shows that ignoring technology adoption in the absence of tightening standards yields underestimates of opportunity costs of the standards. Many previous studies, including recent analysis by the U.S. regulatory agencies, do not account for such technology adoption. Instead, they typically compare the equilibria before and after the standards and in doing so underestimate opportunity costs. We approximate the magnitude of the under-estimate by using our model to compare the 2007 equilibria against the 2012 equilibria with the standards, rather than comparing the 2012 equilibria with and without the standards (which is the comparison

⁷ We assume a maximum 35-year vehicle lifetime, adjusting for survival probabilities for cars and light trucks, and we use the estimated annual vehicle miles traveled by age from U.S. EPA (2012). Consumers value fuel savings at a 10 percent discount rate, which lies within the range of estimates in Busse et al. (2013). Real fuel prices are held constant at their 2007 levels over the lifetime of the vehicle, which is consistent with recent fuel price variation (Klier and Linn 2010). To maintain comparability across regions, assumptions are the same for European and U.S. consumers except for fuel prices.

⁸ We do not allow for heterogeneous preferences for vehicle characteristics, which could affect the welfare analysis (Bento et al. 2012).

⁹ The estimation results in Section 4.4 showed no effect of the standards on horsepower for cars. It may seem surprising that estimated opportunity costs in Table 7 are nonzero. However, the actual standards for cars were less stringent than the standards modeled in the simulations. The simulations are performed assuming that the standards increase the rate of adoption by the same amount as observed in response to the actual standards. More stringent standards could increase the rate of adoption further than observed, in which case the results in Table 7 would over-estimate opportunity costs for U.S. cars.

made in Table 7). A simple before-and-after comparison results in opportunity cost (in the high-technology case) for U.S. light trucks of -\$69 and for Europe of -\$129; negative costs arise because the 2012 horsepower with standards is higher than the 2007 horsepower (these numbers are not reported in the table). Thus, whereas our preferred estimates of opportunity costs are large and positive, ignoring technology adoption in the absence of the standard yields much smaller—and possibly even negative—opportunity costs.

6 Conclusion

Fuel economy and greenhouse gas emissions rate standards will substantially increase passenger vehicle fuel economy in the United States, Europe, as well as other regions. Because vehicle manufacturers choose multiple vehicle characteristics and because technical tradeoffs exist across some of these characteristics, the tightening of fuel economy standards will likely affect other vehicle characteristics besides fuel economy. This paper suggests a new approach to conduct welfare analysis of standards.

We use a simple model of technology adoption to illustrate the connection between fuel economy standards and both the rate *and* direction of technology adoption. Standards increase the rate of technology adoption and cause manufacturers to trade off fuel economy for other characteristics. Reductions in other vehicle characteristics represent opportunity costs of increased fuel economy standards that the previous literature has not estimated properly because of a failure to account for technology adoption in the absence of tighter standards.

Consistent with the model's prediction, we find that recently tightened fuel economy standards in the United States and Europe have increased the rate of technology adoption. We also find strong evidence that the standards reduced light truck torque in the United States and horsepower in Europe. Using these results to simulate the imposition of hypothetical standards, we find that the opportunity costs for U.S. light trucks are similar to the value of the fuel savings. The opportunity costs for U.S. and European cars are economically significant.

The simulations yield three conclusions. First, opportunity costs are smaller after accounting for effects of standards on the rate of adoption (i.e., comparing the low and high cases in Table 7). Second, even after accounting for this effect, opportunity costs are large compared to fuel savings. Finally, failing to account for improved vehicle characteristics in the absence of increased fuel economy standards substantially lowers the estimated opportunity costs.


We leave for future work the incorporation of the opportunity costs in a fully dynamic model of the vehicles market. In such an analysis, it would be possible to relax the assumption maintained in this paper that consumers have homogeneous willingness to pay for vehicle characteristics and that consumers fully value fuel savings. Undervaluation of fuel savings is a commonly used justification for fuel economy standards (Allcott 2013), and future work should consider the welfare and policy implications of this possibility in a dynamic context.

7 References

1. Acemoglu, D., P. Aghion, L. Bursztyn, and D. Hemous (2012). The Environment and Directed Technical Change. *American Economic Review* 102: 131–166.
2. Allcott, H. (2013). The Welfare Effects of Misperceived Product Costs: Data and Calibrations from the Automobile Market. *American Economic Journal: Economic Policy* 5: 1–29.
3. Austin, D. and T. Dinan (2005). Clearing the Air: The Costs and Consequences of Higher CAFE Standards and Gasoline Prices. *Journal of Environmental Economics and Management* 50: 562–582.
4. Bento, A. M., S. Li, and K. Roth (2012). Is There an Energy Paradox in Fuel Economy? A Note on the Role of Consumer Heterogeneity and Sorting Bias. *Economics Letters* 115: 44–48.
5. Busse, M., C. Knittel, and F. Zettelmeyer (2013). Are Consumers Myopic? Evidence from New and Used Car Purchases. *American Economic Review* 103(1): 220–256.
6. Goldberg, P. K. (1998). The Effects of the Corporate Average Fuel Efficiency Standards in the U.S. *The Journal of Industrial Economics* 46: 1–33.
7. Greene, D. L. (1991). Short-Run Pricing Strategies To Increase Corporate Average Fuel Economy. *Economic Inquiry* 29: 101–114.
8. Jacobsen, M. (2013). Evaluating U.S. Fuel Economy Standards in a Model with Producer and Household Heterogeneity. *American Economic Journal: Economic Policy* 5(2): 148–187.
9. Klier, T. and J. Linn (2010). The Price of Gasoline and New Vehicle Fuel Economy: Evidence from Monthly Sales Data. *American Economic Journal: Economic Policy* 2(3): 134–153.
10. Klier, T. and J. Linn (2012a). New Vehicle Characteristics and the Cost of the Corporate Average Fuel Economy Standards. *RAND Journal of Economics* 43: 186–213.
11. Klier, T. and J. Linn (2012b). Using Vehicle Taxes To Reduce Carbon Dioxide Emissions Rates of New Passenger Vehicles: Evidence from France, Germany, and Sweden. Discussion paper 12-34. Washington, DC: Resources for the Future.
12. Klier, T. and J. Linn (2013). Fuel Prices and New Vehicle Fuel Economy—Comparing the United States and Western Europe. *Journal of Environmental Economics and Management* 66: 280–300.
13. Knittel, C. (2011). Automobiles on Steroids: Product Attribute Trade-offs and Technological Progress in the Automobile Sector. *The American Economic Review* 107: 3368–3399.
14. Li, S., J. Linn, and E. Spiller (2013). Evaluating “Cash-for-Clunkers”: Program Effects on Auto Sales and the Environment. *Journal of Environmental Economics and Management* 65: 175–193.

15. Linn, J. (2008). Energy Prices and the Adoption of Energy-Saving Technology. *The Economic Journal* 118: 1986–2012.
16. Mazzeo, M., K. Seim, and M. Varela (2013). The Welfare Consequences of Mergers with Product Repositioning.
17. Newell, R. G., A. B. Jaffe, and R. N. Stavins (1999). The Induced Innovation Hypothesis and Energy-Saving Technological Change. *Quarterly Journal of Economics* 114: 941–975.
18. Popp, D. (2002). Induced Innovation and Energy Prices. *American Economic Review* 92: 160–180.
19. Porter, M. (1999). America’s Green Strategy. *Scientific American* 264: 168.
20. Sweeting, A. (forthcoming). Dynamic Product Repositioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry. *Econometrica*.
21. U.S. Department of Transportation (2006). *Corporate Average Fuel Economy and CAFE Reform for MY 2008–2011 Light Trucks: Final Regulatory Impact Analysis*. Washington, DC: U.S. Department of Transportation.
22. U.S. EPA and NHTSA (2012). *Joint Technical Support Document: Proposed Rulemaking for 2017–2025 Light-Duty Vehicle Greenhouse Gas Emissions Standards and Corporate Average Fuel Economy Standards*. Washington, DC: U.S. EPA and NHTSA.
23. Whitefoot, K., M. Fowlie, and S. Skerlos (2011). Product Design Response to Policy Intervention: Evaluating Fuel Economy Standards Using an Engineering Model of Product Design. Working paper.
24. 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: 402–411.


Figure 1. Fraction Change in Fuel Economy, Weight, and Power, 1975–2008 for Cars Sold by U.S. Manufacturers



Notes: The figure reports the fraction change of the sales-weighted mean fuel economy (in mpg), weight (in pounds), and horsepower, relative to 1975.


Source: Klier and Linn (2012a).

Figure 2. Technology Frontier




Notes: The solid curve shows the maximum fuel economy (m) as a function of horsepower (h) for a power train with efficiency η' . The dashed curve shows the maximum fuel efficiency for a power train with a higher efficiency, η'' .

Figure 3. No-Policy Equilibrium




Notes : The figure plots the vehicle's fuel economy against its horsepower. The solid curve represents the technological tradeoff between fuel economy and horsepower at time $t = s$ for the chosen efficiency η_s . The dashed curve represents the technological tradeoff at time $t = s + 1$ for the chosen efficiency $\eta_{s+1} > \eta_s$. The downward-sloping lines represent indifference curves at time $t = s$ and time $t = s + 1$. The points X_s and X_{s+1} are the equilibria at time s and time $s + 1$.

Figure 4. Equilibrium with a Fuel Economy Standard




Notes : The figure shows the equilibrium in period $s + 1$ caused by imposing a fuel economy standard of m^* , in which the firm may purchase or sell credits to comply. The equilibrium is at the point X_{s+1}^* , which is the intersection of the frontier and the diagonal line representing the indifference curve at the equilibrium. The frontier labeled η''' is the frontier chosen with the standard.

Figure 5. Fuel Economy, Horsepower, and Weight




Notes : Panel A plots the fraction change in sales-weighted fuel economy, weight, and power since 2000 for the United States, using the same data set as Table 1. Panel B plots fraction changes in registration-weighted fuel economy, weight, and power since 2005 for Europe, using the same data set as Table 1.

Figure 6. Technology Adoption in the United States and Europe




Notes : In Panel A, the frontier plots the change in redesign, model, and model-year interactions estimated in column 1 of Table 2; in Panel B, the frontier plots the interactions estimated in column 2 of Table 2; and in Panel C, the frontier plots the interactions estimated in column 1 of Table 3. In Panels A and B, fuel economy is the change since 2000 in average log fuel economy across specifications. Horsepower, weight, and diesel represent the increase in fuel economy that would have been possible if these characteristics had remained at their 2000 levels. The variables are computed using the log change in the characteristic since 2000, multiplied by the negative of the coefficient on the corresponding characteristic from the regressions in Table 2. The curves in Panel C are constructed similarly, and represent changes since 2005. Vertical lines indicate the adoption of the higher standards.


Figure 7. U.S. Technology Adoption by Company, Cars



█ Frontier
 █ Fuel economy
 █ Horsepower
 █ Weight
 █ Hybrid
 █ Auto/manual, cylinders, diesel fuel




█ Frontier
 █ Fuel economy
 █ Horsepower
 █ Weight
 █ Hybrid
 █ Auto/manual, cylinders, diesel fuel




█ Frontier
 █ Fuel economy
 █ Horsepower
 █ Weight
 █ Hybrid
 █ Auto/manual, cylinders, diesel fuel

Panel D: Honda




Frontier Fuel economy Horsepower Weight Hybrid Auto/manual, cylinders, diesel fuel

Panel E: Nissan




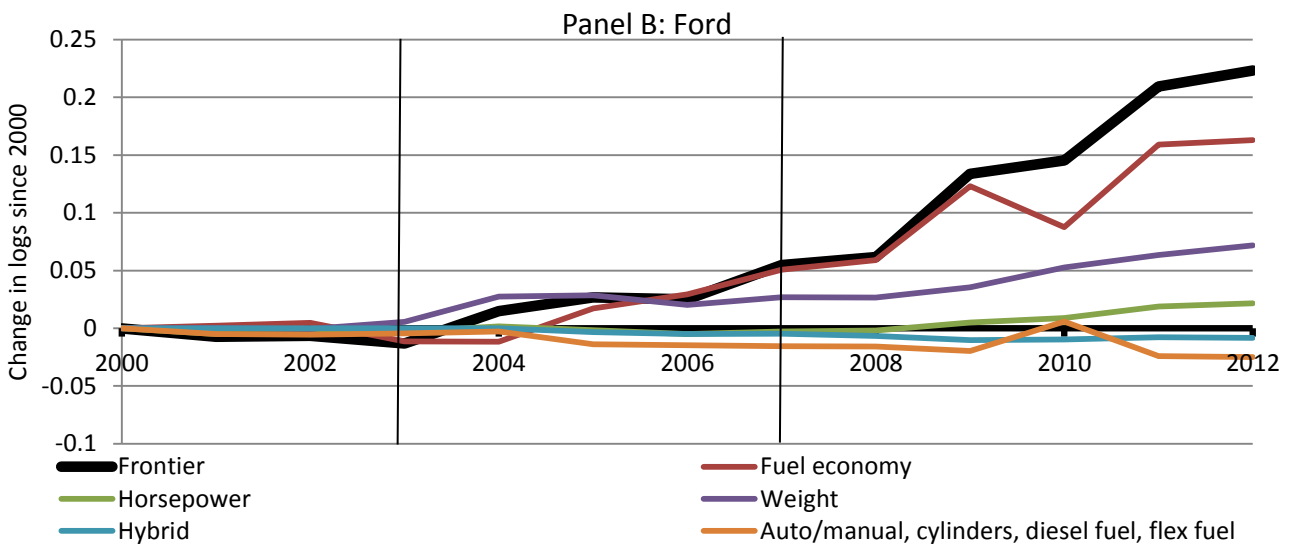

Frontier Fuel economy Horsepower Weight Hybrid Auto/manual, cylinders, diesel fuel

Panel F: Chrysler



Frontier Fuel economy Horsepower Weight Hybrid Auto/manual, cylinders, diesel fuel

Figure 8. U.S. Technology Adoption by Company, Trucks











Figure 9. European Technology Adoption by Brand








Table 1. Summary Statistics for the United States and Europe, 2005 and 2010

	<u>United States</u>		<u>Europe</u>	
	<u>2005</u>	<u>2010</u>	<u>2005</u>	<u>2010</u>
Sales or registrations	12,586.85 (11,894.34)	6,855.51 (8,051.22)	268.19 (829.14)	203.79 (628.92)
Fuel economy (mpg)	25.21 (6.10)	26.78 (7.01)	34.47 (8.94)	38.16 (9.27)
Horsepower	228.21 (64.64)	258.12 (77.63)	134.44 (56.01)	150.01 (68.34)
Weight (tons)	1.99 (0.41)	2.09 (0.46)	2.09 (0.37)	2.21 (0.42)
Number of obserations	1,352	1,546	46,521	47,884

Notes : The table reports the means of the indicated variables, with standard deviations in parentheses, for model-years 2005 and 2010. The United States data set includes observations by model version from 2000 to 2011, and the European data set includes observations by model version for 2005 to 2010. For the United States, weight is the curb weight; for Europe, weight is the gross vehicle weight. See text for details on the construction of the data sets.

Table 2. United States: Tradeoffs between Fuel Economy and Other Vehicle Characteristics

	(1)	(2)
	<u>Dependent variable: log fuel economy</u>	
Log horsepower or torque	-0.237 (0.015)	-0.156 (0.016)
Log weight	-0.336 (0.044)	-0.430 (0.047)
Diesel fuel	0.344 (0.019)	0.269 (0.020)
Hybrid	0.260 (0.020)	0.293 (0.010)
Flex fuel		-0.282 (0.014)
Manual transmission	-0.002 (0.005)	-0.005 (0.004)
Number of observations	6,856	12,208
R ²	0.957	0.937
Sample includes	Cars	Light trucks
Regression includes	Interactions of redesign, model, and model-year, and fixed effects for number of cylinders and number of doors	Interactions of redesign, model, and model-year, and fixed effects for number of cylinders and number of doors

Notes : The table reports coefficient estimates, with standard errors in parentheses, clustered by redesign, model, and model-year. Observations are by model version and model-year, and the dependent variable is log fuel economy. Besides the reported variables, the regressions include the variables indicated at the bottom of the table. The sample includes cars in column 1 and light trucks in column 2. Column 1 uses the log of horsepower as an independent variable, and column 2 uses the log of torque.

Table 3. Europe: Tradeoffs between CO₂ Emissions Rate and Other Vehicle Characteristics

	(1)	(2)
	<u>Dependent variable: log CO₂ emissions rate</u>	
Log horsepower	0.190 (0.002)	0.158 (0.002)
Log weight	0.307 (0.007)	0.241 (0.012)
Diesel fuel	-0.174 (0.001)	-0.172 (0.001)
Manual transmission	-0.071 (0.001)	-0.076 (0.001)
Number of observations	276,376	276,376
R ²	0.916	0.944
Regression includes	Number of cylinders and interactions of redesign, model, and model-year	Number of cylinders and interactions of redesign, model-trim, and model-year

Notes : The table reports coefficient estimates, with standard errors in parentheses, clustered by redesign, model trim, and model-year. Observations are by model version and model-year, and the dependent variable is log of the CO₂ emissions rate. Besides the reported variables, the regressions include the variables indicated at the bottom of the table. A model trim includes all specifications with the same model, body type, number of doors, driven wheels, and trim level.

Table 4. Effect of U.S. Standards on Direction and Rate of Technology Adoption

	(1)	(2)	(3)	(4)
<u>Panel A: direction</u>				
Dependent variable	Log (fuel economy / horsepower)	Log (fuel economy / weight)	Log (fuel economy / torque)	Log (fuel economy / weight)
Stringency X 2003–2006	-0.029 (0.113)	-0.090 (0.069)	-0.516 (0.152)	0.048 (0.097)
Stringency X 2007–2009	0.122 (0.119)	0.060 (0.072)	-0.643 (0.178)	-0.012 (0.111)
Stringency X 2010–2012	0.014 (0.125)	0.101 (0.078)	-0.792 (0.175)	-0.137 (0.123)
Number of observations	6,856	6,856	11,966	11,966
R ²	0.832	0.874	0.795	0.855
<u>Panel B: rate</u>				
Stringency X 2003–2006	0.017 (0.040)	0.006 (0.040)	-0.226 (0.063)	-0.241 (0.065)
Stringency X 2007–2009	-0.024 (0.047)	-0.056 (0.046)	-0.269 (0.065)	-0.261 (0.066)
Stringency X 2010–2012	-0.091 (0.051)	-0.101 (0.050)	-0.142 (0.067)	-0.123 (0.068)
Number of observations	1,749	1,749	1,425	1,425
R ²	0.768	0.766	0.847	0.838
Sample includes	Cars	Cars	Light trucks	Light trucks
Frontier estimated by	Entire market	Market segment	Entire market	Market segment

Notes : The table reports coefficient estimates, with standard errors in parentheses, clustered by redesign, model, and model-year. Observations are by model version and model-year in Panel A and by redesign, model, and model-year in Panel B. Regressions in columns 1 and 2 include cars, and regressions in columns 3 and 4 include light trucks. In Panel A, the dependent variable is the log of the ratio of fuel economy to horsepower in column 1, the log of the ratio of fuel economy to weight in columns 2 and 4, and the log of the ratio of fuel economy to torque in column 3. The dependent variable in Panel B is the estimated redesign–model-year interaction from equation (7). Columns 1 and 3 use the estimated redesign–model-year interactions from Table 2 and columns 2 and 4 use estimates from Table 3. Stringency is the difference between the log sales-weighted fuel economy in 2000 and the log sales-weighted standard for the corresponding company and vehicle type. The calculation uses the 2016 standards. All regressions include model fixed effects and triple interactions between model-year, market segment, and the fuel economy of the model in 2000, along with all associated main effects and double interaction terms.

Table 5. Effect of European Emissions Rate Standards on Direction and Rate of Technology Adoption

	(1)	(2)
<u>Panel A: direction</u>		
Dependent variable	Log (fuel economy / horsepower)	Log (fuel economy / weight)
Stringency X post 2007	-0.030 (0.008)	-0.014 (0.007)
Number of observations	275,675	275,675
R ²	0.765	0.586
<u>Panel B: rate</u>		
Stringency X post 2007	-0.029 (0.004)	-0.022 (0.006)
Number of observations	63,824	63,824
R ²	0.952	0.964

Notes : The table reports coefficient estimates, with standard errors in parentheses, clustered by redesign, model trim, and model-year. Observations are by model version and model-year in Panel A and by model trim and model-year in Panel B. The dependent variable in Panel A is the log of the ratio of fuel economy to horsepower in column 1 and the log of the ratio of fuel economy to weight in column 2. In Panel B, the dependent variable in column 1 is the estimated model-year interaction from column 1 of Table 4, and the dependent variable in column 2 is the estimated model trim by model-year interaction from column 2 of Table 4. Stringency is the difference between the log registration-weighted brand emissions rate in 2005 and the log registration-weighted 2015 standard for the corresponding brand, adjusted by weight. The variable is interacted with a dummy variable equal to one for years 2008–2010. Regressions include model-year fixed effects, the interactions between a post-2007 dummy variable and market segment fixed effects, the interaction between the model's 2005 emissions rate and a post-2007 dummy variable, the interaction between the model's 2005 emissions rate and a set of segment fixed effects, and the interaction between the log of the model's 2005 emissions rate interacted with the post-2007 by segment interactions. All regressions include model trim fixed effects.

Table 6. Potential Omitted Variables Bias: Fuel Prices and the Recession

	(1)	(2)	(3)
		<u>Panel A: exit</u>	
Log mpg/emissions rate X log fuel price	-0.162 (0.223)	-0.403 (0.242)	0.040 (0.036)
Log mpg/emissions rate X log aggregate sales / registrations	0.238 (0.316)	-0.260 (0.306)	0.017 (0.109)
Number of observations	5,850	5,535	228,492
R ²	0.274	0.391	0.009
Sample includes	U.S. cars	U.S. light trucks	European cars
		<u>Panel B: sales and registrations</u>	
Stringency X 2003–2006 X log gas price	2.705 (3.583)	-0.850 (3.093)	
Stringency X 2007–2009 X log gas price	6.253 (8.877)	7.406 (6.565)	
Stringency X 2010–2012 X log gas price	21.801 (68.811)	3.653 (9.256)	
Stringency X 2003–2006 X log aggregate sales	-35.060 (45.445)	10.022 (43.094)	
Stringency X 2007–2009 X log aggregate sales	-3.173 (5.884)	-7.151 (4.557)	
Stringency X 2010–2012 X log aggregate sales	1.067 (8.949)	-1.257 (3.178)	
Stringency X post 2007 X log fuel price			4.433 (7.076)
Stringency X post 2007 X log aggregate registrations			9.885 (4.636)
Number of observations	1,506	1,206	3,590
R ²	0.347	0.224	0.219
Sample includes	U.S. cars	U.S. light trucks	European cars

Notes : Standard errors are in parentheses, clustered by redesign, model, and model-year in Panel A and by model and model-year in Panel B. Observations are by model version and model-year in Panel A and by model and model-year in Panel B. The dependent variable in Panel A is an indicator equal to one if the vehicle exits between the current and next years. The dependent variable in Panel B is the log sales in columns 1 and 2 and the log registrations in column 3. Log mpg X log fuel price is the interaction between fuel economy and the log of the fuel price. Log mpg X aggregate sales is the interaction between the vehicle's log fuel economy and the log of the total annual sales in the market. In Panel, A columns 1 and 2 include these variables along with the main effects of the interaction terms. In Panel A, column 3 uses the same main effects and interaction terms, except that the emissions rate replaces fuel economy and registrations replaces sales. Instead of these variables, Panel B includes the interactions of stringency, time period, and gas price or aggregate sales. Stringency and time periods are defined as in Tables 6 and 7. All regressions include triple interactions of year, market segment, and initial fuel economy, along with lower-order main effects and interactions, as in Tables 6 and 7.

Table 7. Effects on Consumer Welfare of a 10 Percent Fuel Economy Increase

	(1)	(2)	(3)	(4)	(5)
	Frontier shift (2007–2012)	Unregulated fraction horsepower / torque change	Regulated fraction horsepower / torque change	WTP for horsepower / torque change (2005 \$)	WTP for mpg change (2005 \$)
<u>Panel A: U.S. Cars</u>					
Low technology adoption	0.053	0.180	-0.199	-757	707
High technology adoption	0.120	0.180	0.085	-219	707
<u>Panel B: U.S. Light Trucks</u>					
Low technology adoption	0.038	0.331	-0.398	-1,659	1,262
High technology adoption	0.097	0.331	-0.022	-953	1,262
<u>Panel C: Europe</u>					
Low technology adoption	0.078	0.230	-0.115	-324	957
High technology adoption	0.121	0.230	0.110	-125	957

Notes : The table reports results of a hypothetical standard that raises fuel economy of all vehicles by 10 percent over five years. Panel A shows results for U.S. cars, Panel B for U.S. light trucks, and Panel C for Europe. Each row represents a separate simulation performed on all vehicles in the 2007 data. The low technology adoption scenario assumes that the technology frontier shifts out by the average amount observed over 2000–2009 for U.S. cars, 2000–2002 for U.S. trucks, and 2005–2007 for European cars. The high technology adoption scenario assumes that the frontier shifts out at the average rate over 2009–2012 for U.S. cars, over 2003–2012 for light trucks, and over 2008–2010 for European cars. The frontier shift is the change in the technology frontier. The unregulated scenarios allow the frontier to shift out at the same rate as the low technology adoption case. The unregulated fuel economy for each vehicle is computed using the frontier shift and the fraction of technology adoption used to increase fuel economy in the absence of regulation, computed using the ratio of the percentage change in average fuel economy to the percentage frontier shift during the time periods. The unregulated horsepower or torque increase is calculated assuming that the remaining technology adoption is used to increase horsepower or torque (horsepower for cars and torque for U.S. light trucks). Horsepower or torque in each scenario is computed using the corresponding frontier shift and assuming that all technology adoption not used to increase fuel economy by 10 percent is used to increase horsepower or torque. The fraction change in horsepower or torque is the difference between the horsepower or torque in the scenario and the unregulated horsepower or torque, where observations are weighted by 2007 registrations. The willingness to pay (WTP) for the horsepower or torque changes is computed using a value of \$10 per horsepower or torque per ton in column 3. In column 5, the WTP for the fuel economy increase is calculated using a fuel price of \$6.15 per gallon in Europe, a 10 percent discount rate, and the vehicle miles traveled and survival estimates in U.S. EPA (2012). All WTP estimates are in 2005 dollars, and all observations are weighted by 2007 registrations.

Appendix Table 1. Tradeoffs by Segment: United States

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	<u>Dependent variable: log fuel economy</u>						
Log horsepower or torque	-0.245 (0.028)	-0.188 (0.028)	-0.200 (0.016)	-0.165 (0.022)	-0.154 (0.026)	-0.183 (0.051)	-0.122 (0.021)
Log weight	-0.654 (0.107)	-0.186 (0.060)	-0.275 (0.053)	-0.571 (0.105)	-0.587 (0.035)	-0.358 (0.051)	-0.399 (0.058)
Diesel fuel	0.375 (0.017)	0.201 (0.028)	0.309 (0.022)	0.287 (0.018)	0.228 (0.019)		
Hybrid	0.350 (0.030)	0.293 (0.023)	0.105 (0.026)	0.319 (0.015)	0.286 (0.025)		0.298 (0.007)
Flex fuel				-0.352 (0.003)	-0.290 (0.011)	-0.225 (0.003)	-0.280 (0.017)
Manual transmission	-0.002 (0.010)	0.006 (0.006)	-0.012 (0.007)	0.008 (0.008)	0.003 (0.005)		-0.004 (0.004)
Number of observations	1,798	2,188	2,870	2,416	2,826	1,105	5,861
R ²	0.943	0.945	0.907	0.933	0.910	0.959	0.867
Sample includes	Small cars	Medium cars	Large/luxury cars	Crossovers	Sport utility vehicles	Vans	Pickups

Notes : The table reports coefficient estimates, with standard errors in parentheses, clustered by model and model-year. Columns 1–3 report similar regressions to those reported in column 1 of Table 2, except that the sample is restricted to included observations from the market segment indicated at the bottom of the table; likewise, columns 4–7 correspond to column 2 in Table 2.

Appendix Table 2. Tradeoffs by Segment: Europe

	(1)	(2)	(3)	(4)	(5)	(6)
	<u>Dependent variable: log CO₂ emissions rate</u>					
Log horsepower	0.163 (0.009)	0.182 (0.003)	0.252 (0.002)	0.145 (0.003)	0.061 (0.005)	0.021 (0.019)
Log weight	0.246 (0.040)	0.207 (0.018)	0.263 (0.012)	0.360 (0.014)	0.375 (0.016)	0.206 (0.035)
Diesel fuel	-0.140 (0.004)	-0.197 (0.001)	-0.173 (0.001)	-0.174 (0.001)	-0.163 (0.002)	-0.134 (0.005)
Manual transmission	-0.058 (0.004)	-0.065 (0.001)	-0.076 (0.001)	-0.076 (0.001)	-0.059 (0.001)	-0.009 (0.009)
Number of observations	8,263	47,425	91,430	90,206	35,170	3,882
R ²	0.812	0.820	0.796	0.868	0.909	0.892
Sample includes	Mini	Small	Lower medium	Medium	Upper medium	Large

Notes : The table reports coefficient estimates, with standard errors in parentheses, clustered by model and model-year. Each column reports a similar regression to column 1 of Table 4 except that the sample is restricted to included observations from the market segment indicated at the bottom of the table.

Appendix Table 3. Effect of U.S. Standards on Direction and Rate of Technology Adoption, Omitting Other Controls

	(1)	(2)	(3)	(4)
<u>Panel A: direction</u>				
Dependent variable	Log (fuel economy / horsepower)	Log (fuel economy / weight)	Log (fuel economy / torque)	Log (fuel economy / weight)
Stringency X 2003–2006	-0.086 (0.115)	-0.102 (0.067)	-0.264 (0.137)	0.062 (0.100)
Stringency X 2007–2009	0.040 (0.114)	0.089 (0.069)	-0.338 (0.163)	0.100 (0.117)
Stringency X 2010–2012	-0.006 (0.124)	0.125 (0.078)	-0.441 (0.165)	0.031 (0.134)
Number of observations	6,856	6,856	11,966	11,966
R ²	0.830	0.871	0.791	0.848
<u>Panel B: rate</u>				
Stringency X 2003–2006	-0.034 (0.046)	-0.053 (0.049)	-0.270 (0.056)	-0.276 (0.061)
Stringency X 2007–2009	-0.046 (0.050)	-0.051 (0.051)	-0.371 (0.055)	-0.337 (0.062)
Stringency X 2010–2012	-0.140 (0.064)	-0.091 (0.062)	-0.281 (0.056)	-0.225 (0.063)
Number of observations	1,749	1,749	1,425	1,425
R ²	0.753	0.754	0.829	0.821
Sample includes	Cars	Cars	Light trucks	Light trucks
Frontier estimated by	Entire market	Market segment	Entire market	Market segment

Notes : The table reports coefficient estimates, with standard errors in parentheses, clustered by redesign, model, and model-year. Regressions are the same as in Table 6 except that the independent variables include only the reported variables, model-year fixed effects, and model fixed effects.

Appendix Table 4. Effect of European Emissions Rate Standards on Direction and Rate of Technology Adoption, Omitting Other Controls

	(1)	(2)
<u>Panel A: direction</u>		
Dependent variable	Log (fuel economy / horsepower)	Log (fuel economy / weight)
Stringency X post 2007	-0.076 (0.007)	-0.046 (0.006)
Number of observations	275,675	275,675
R ²	0.764	0.585
<u>Panel B: rate</u>		
Stringency X post 2007	-0.028 (0.003)	-0.022 (0.005)
Number of observations	63,824	63,824
R ²	0.950	0.963
Frontier estimated by	Entire market	Market segment

Notes : The table reports coefficient estimates, with standard errors in parentheses, clustered by redesign, model, and model-year. Regressions are the same as in Table 7 except that the independent variables include the reported variables, model-year fixed effects, and model trim fixed effects.