DISCUSSION PAPER

Dec. 2013; second revision Sep. 2015 ■ RFF DP 13-40-REV2

The Effect of Vehicle Fuel Economy Standards on Technology Adoption

Thomas Klier and Joshua Linn

1616 P St. NW
Washington, DC 20036
202-328-5000 www.rff.org

RESOURCES FOR THE FUTURE

The Effect of Vehicle Fuel Economy Standards on Technology Adoption

Thomas Klier and Joshua Linn

Abstract

Many countries are tightening passenger vehicle fuel economy standards. This paper provides the first empirical evidence on the effects of fuel economy standards on technology adoption. We investigate changes in the rate and direction of technology adoption, that is, the extent to which technology is used to increase fuel economy at the expense of other vehicle attributes. We find that recent changes in US and European standards have both increased the rate of technology adoption and affected the direction of technology adoption. Producers reduced horsepower and torque compared to a counterfactual in which fuel economy standards remained unchanged. We estimate opportunity costs from reduced horsepower and torque to be economically significant relative to the gains from fuel savings.

Key Words: passenger vehicles, US greenhouse gas emissions rate standards, European carbon dioxide emissions rate standards, technology adoption

JEL Classification Numbers: L62, Q4, Q5

Contents

The Effect of Vehicle Fuel Economy Standards on Technology Adoption

Thomas Klier and Joshua Linn*

1. Introduction

Due to concerns about global warming and energy security, a number of countries have recently adopted policies to substantially increase the average fuel economy of new passenger vehicles. The current US Corporate Average Fuel Economy (CAFE) standards, to be met by 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 United States 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 between 2012 and 2020. In addition, many other countries, representing developed economies, such as Japan and Canada, as well as developing countries, such as Mexico and China, have put in place similar policies.

The growing literature on fuel economy standards has used structural models of the new vehicles market to characterize the welfare effects of such standards.[1] Following conventional microeconomic theory, these models allow for the possibility that manufacturers raise prices on vehicles with low fuel economy and reduce prices on vehicles with high fuel economy in order to meet standards in the short term (Green 1991; Goldberg 1998). Several models, such as in Austin and Dinan (2005), also include a longer-term perspective by allowing for the adoption of fuel-saving technology. A number of recent studies include a third margin along which manufacturers can respond to standards; it consists of trading off fuel economy for other vehicle characteristics such as horsepower (e.g., Whitefoot et al. 2011; Knittel 2011; Klier and Linn 2012; Whitefoot

* Thomas Klier, Federal Reserve Bank of Chicago. Joshua Linn, Resources for the Future. We thank Shanjun Li, Virginia McConnell, seminar participants at the Colorado School of Mines, Cornell University, and the University of Maryland, and conference participants at the International Industrial Organization Conference. Wenfei Du and Sam Goldberg 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] Analysis by the US regulatory agencies—the US Environmental Protection Agency (EPA) and the US Department of Transportation National Highway Traffic Safety Administration (NHTSA)—tends to be more favorable to fuel economy standards than the economics literature (see, e.g., Jacobsen 2012). Although there are many differences between the analysis of the regulatory agencies and the studies in the economics literature, assumptions about technology costs likely play an important role in explaining the differing conclusions.

and Skerlos 2012). Thus, the literature suggests that manufacturers respond to tightened fuel economy standards by adjusting vehicle prices, adopting fuel-saving technology, and trading off fuel economy for other vehicle characteristics.

The broader literature on technology adoption and innovation provides additional context on the effect of fuel economy standards. It suggests that tightened standards can influence the rate of technological change, resulting in manufacturers innovating and adopting technology more quickly, in addition to trading off fuel economy for other vehicle characteristics. That literature has demonstrated that profit and regulatory forces can affect product characteristics. 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.[2] While fuel economy standards can affect the rate of technology adoption and cause manufacturers to trade off fuel economy for other characteristics, neither the CAFE nor the technology literature has tested such predictions. Specifically, Knittel (2011) and Klier and Linn (2012) provide suggestive evidence that the introduction of the CAFE standards in 1978 affected technology adoption. Knittel (2011) finds a positive correlation between the market-wide average rate of adoption and changes in the standards during the 1980s. Klier and Linn (2012) show that falling weight and horsepower explain about half of the overall fuel economy increase in the early 1980s. Yet neither study establishes a causal connection between fuel economy standards and technology.

Our paper provides the first such evidence for the new vehicles market. We analyze four recent changes in standards in the United States and Europe and examine manufacturers' response to those changes. 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, replacing a voluntary standard that, incidentally, manufacturers did not meet (Klier and Linn, forthcoming). The two regions' standards contrast in several dimensions: the European standards are more stringent than the US standards, requiring about 20 percent higher fuel economy in 2015; and the US standards depend on the vehicle's footprint (roughly, the area defined by its axles) whereas the European standards

[2] There is a related literature on consumer valuation of product characteristics and product design (e.g., Mazzeo et al. 2013 and Sweeting forthcoming).

depend on the vehicle's weight. Other market conditions, such as fuel prices, differ considerably across the two regions.

Departing from most of the structural CAFE literature, but in the spirit of the technology literature, we take a reduced-form approach and ask whether recent changes in US and European standards have increased the rate at which manufactures adopt technology. In light of the recent focus in the CAFE literature on tradeoffs between fuel economy and other vehicle characteristics, we also ask whether this margin of adjustment is important in practice.

Our empirical strategy consists of two stages. The major empirical objective in the first stage is to distinguish between a) technology adoption that trades off fuel economy for other vehicle characteristics and b) technology adoption that raises fuel economy without sacrificing other characteristics. The approach builds on Linn (2008), Knittel (2011) and Klier and Linn (2012). We begin by defining a frontier that corresponds to a vehicle's minimum fuel consumption rate (which is inversely proportional to fuel economy), given its horsepower, weight and other characteristics. In the short run vehicle characteristics are fixed and in the medium run manufacturers can trade off characteristics without redesigning the vehicle. Long-run tradeoffs between fuel economy and characteristics occur when the vehicle is redesigned. We make this distinction because engine design cycles typically last 8–10 years, and technological tradeoffs between fuel economy and other characteristics across design cycles may be different from tradeoffs within design cycles. To simplify our analysis we assume that the magnitude of the tradeoffs along the frontier—for example, the percent change in horsepower needed to increase fuel economy by 1 percent—does not vary across vehicles or over time. This assumption allows us to use a simple linear regression technique that simultaneously estimates the shape of the frontier and shifts of the frontier over time (we partially relax this assumption subsequently).

We use detailed engine and vehicle characteristics data to estimate tradeoffs between fuel economy and vehicle characteristics separately for the US and European vehicle markets. Previously, Knittel (2011) and Klier and Linn (2012) estimated tradeoffs using cross-sectional and time series variation in vehicle characteristics. We extend their analyses by matching engine data to vehicle model production data. That allows us to distinguish between medium-run and long-run tradeoffs among fuel economy, weight, and power. 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. We further improve on the literature by estimating separate frontiers by engine, model, and model-year, rather than by just model-year, as in Knittel (2011).

In the second stage, we use the estimated frontiers to examine whether recent changes in standards have affected the rate or direction of technology adoption, where a change in direction refers to movement along the frontier. The previous literature has been unable to identify these effects because it has treated the standards as an aggregate shock to the industry. In contrast, 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.[3]

Regarding the four cases of tightening standards, we find that the change in US light truck standards in 2003 and 2007 affected both the rate and direction of technology adoption. The 2007 US car standards affected the rate of technology adoption, although not as much as for light trucks; there is mixed evidence whether the 2007 car standards also affected the direction of technology adoption. The results for US cars are consistent with the fact that the car standards tightened at the end of the sample whereas the light truck standards were tightening over a longer time period. The European standards affected the rate of adoption and also had a small, but statistically significant, effect on the direction of technology adoption. Thus, despite the regional differences in the stringency and form of the standards, as well as in market conditions such as fuel prices, we find strong empirical evidence that both US and European standards affected the rate and direction of technology adoption.

To assess the economic significance of these effects, we perform simulations that yield rough estimates of the value of changes in fuel economy and other characteristics caused by tightened standards. To do that, we define the opportunity costs of the standards as the effect of the standards on a consumer's willingness to pay for vehicle characteristics other than fuel economy. We estimate the opportunity costs of a hypothetical 10 percent increase in fuel economy standards for both the United States and Europe. The estimated opportunity costs are smaller than the value of the fuel savings but are nonetheless economically significant. Thus, structural models used to estimate welfare costs of tightened standards that do not include tradeoffs between vehicle characteristics miss a quantitatively important aspect of the welfare

[3] 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. In the results section we report several pieces of evidence that suggest that our identification strategy can control for these factors.

costs. The results also underscore the need to simultaneously endogenize the rate and direction of the standards when estimating their welfare costs; the opportunity costs would be much larger if one did not account for the increase in the rate of technology adoption caused by tighter standards.

Our paper is most closely related to Knittel (2011). Using US data from 1980-2006, Knittel estimates the tradeoffs between fuel economy and other characteristics as well as the cross-model annual average rate of adoption. Knittel uses these estimates to assess the technical feasibility rather than welfare costs of the 2016 standards, and whether weight or horsepower reductions would be needed to meet the standards, assuming the frontiers continue to shift at their historical rates. Considering the initial phase-in of the US standards during the late 1970s and 1980s, he also reports a positive correlation between changes in the standards and the market-wide average rate of technology adoption.

In four ways, our analysis extends Knittel's analysis. First, we compare the effects of changes in US and European standards on the rate and direction of technology adoption. We estimate similar technical tradeoffs across the regions and we find that both regions' standards affect technology, albeit by different amounts. Second, by estimating the effects of standards on the rate and direction of technology adoption, rather than just the rate, we can gauge the opportunity costs of technology adoption, whereas the previous empirical literature has not accounted for both margins. Third, instead of computing time series correlations between the level of the standards and average adoption rates, we identify the effects of the standards using cross-sectional variation in regulatory stringency. This allows us to control for unobserved shocks that may be correlated with changes in the standards, such as business cycle fluctuations. Finally, rather than focusing on the initial US phase-in period, we analyze the recent changes in US standards for cars and light trucks and European standards for cars. Many observers believe that, given existing technology, the new standards are harder to meet than the initial standards and thus may have different effects on the rate and direction of technology adoption. In the concluding section we discuss three policy implications of our results: a) because standards increase the rate of adoption, opportunity costs are lower than those suggested in the literature; b) the European results are suggestive of future potential technology improvements in the US market; and c) opportunity costs of the standards have been economically significant and should be included in welfare analysis of future standards.

2. Data and Overview of the US and European Standards

2.1 Data

The US 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, and hybrid. We merged to the sales data other vehicle characteristics—such as engine displacement, number of cylinders, horsepower, torque, and fuel economy—from Wards annual yearbooks (fuel economy is harmonized to reflect changes over time in EPA rating systems). Those characteristics were measured at the model version level, for example distinguishing versions of the Honda Accord that have 4 and 6 cylinders. A "model version" includes one or more versions of the same model and trim, which may differ because of certain options such as body type (for example, two-door or hatchback). The characteristics data distinguish diesel fuel from gasoline versions.[4]

Finally, we merge to the Wards data additional engine data by vehicle model, fuel type, and number of cylinders. These engine data distinguish three levels of engine aggregation: an engine platform combines related engine programs, each of which in turn may consist of multiple engine models. The data, which originated with 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.[5] On average, models contain redesigned engines about once every five years in the United States, and less frequently in Europe. There exists a fair amount of cross-company variation. Models sold by GM contain redesigned engines nearly twice as frequently as

[4] More specifically, the sales are reported by model, fuel type, and month. The characteristics are reported by "model version", which refers to a unique model, fuel type, trim, body type, etc. We aggregate monthly sales to annual sales and merge the sales data to characteristics data by model, year, and fuel type. In the analysis below, the first stage estimation uses only the characteristics data, whereas some of the robustness checks for the second-stage analysis uses the sales data. In those cases, we allocate sales of a particular model, year, and fuel type evenly across corresponding model versions.

[5] The production data available to us have worldwide coverage for 2000–2007 but only 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 estimation results; this suggests that any measurement error in the redesign variable does not significantly bias the estimates.

models sold by Honda, Nissan, and Toyota, whereas models sold by Ford are more similar to those sold by the Japanese producers. Chrysler lies in between the two extremes. This variation arises from underlying variation in the length of design cycles and the number of engines sold with a particular model.

Table 1 provides some summary statistics for the US 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 1 shows the trends over the entire sample period, 2000–2012. Horsepower and weight increased steadily in the United States in the first half of the sample and then leveled off (more so for weight than horsepower, which resumed growth in 2009), whereas fuel economy was constant in the first half and then increased steadily in the second half; 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 the eight largest markets in Europe–Austria, Belgium, France, Germany, Italy, the Netherlands, Spain, and the United Kingdom–and in Sweden. 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 US data. A European model-year corresponds to a calendar year (Klier and Linn 2013). Panel B of Figure 1 shows that horsepower, weight, and fuel economy in Europe 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. The summary statistics reported in Table 1 allow comparisons between the European and the US data. Fuel economy is much lower and horsepower is much higher in the United States than in Europe. The fuel consumption rate (measured in gallons per 100 miles) is inversely related to fuel economy, and is much lower in Europe than the United States. The reported weight is larger in Europe, but that is because the European data report the gross vehicle weight, while the US data report the curb weight (gross vehicle weight includes the weight of passengers and cargo, and curb weight does not). 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.

2.2 Overview

We briefly summarize the US and European standards and discuss the empirical implications of the contrasts between the two regions' policies. The US CAFE standards were first phased in by the mid-1980s and set standards of 27.5 mpg for cars and 20.7 mpg for light trucks. In 2003 NHTSA finalized new, tighter, standards for light trucks that increased fuel economy to 24.1 mpg by 2011. The 2007 Energy Independence and Security Act set standards for 2020 that would require a combined 35 mpg for cars and light trucks; in 2009 the compliance deadline for these new standards was moved to 2016. Thus, tighter standards were announced in 2003 (light trucks), 2007 (cars and light trucks) and 2009 (cars and light trucks). In those three cases the new standards are based on the footprint of the vehicle, which is roughly the vehicle's width multiplied by the distance between the two axles.

Europe has traditionally taxed fuels at much higher rates than the United States (Klier and Linn 2013). European cars also tend to be much smaller and have higher fuel economy than US cars. In 1998 and 1999 manufacturers agreed to increase fuel economy to about 39 mpg within 10 years. However, they did not increase fuel economy quickly enough to meet the targets and in 2007 the European Commission issued requirements for about 42 mpg by 2015. Whereas the initial agreement was purely voluntary, compliance with the 2015 standards is backed by fines (Klier and Linn 2014). The standard is weight-based, depending on the difference between a specific car's weight and a benchmark weight.

The US and European standards thus differ in several important ways. First, the European standards are more stringent, which may make it more challenging for manufacturers to increase the rate of technology adoption. Second, one of the reasons for adopting footprint rather than weight-based standards in the United States was to provide stronger incentives for weight reduction. For that reason the tighter standards might cause more weight reduction in the United States than in Europe.[6] Third, high fuel prices in Europe incentivize manufacturers to raise fuel economy when they adopt fuel-saving technology, even in the absence of tighter standards. Therefore we might expect tighter standards to affect the direction of technology adoption to a smaller extent in Europe than the United States, all else equal.

[6] The European standards are a linear function of the vehicle's weight, which contrasts with the Japanese weight-based standards that jump discreetly with the vehicle's weight. The linearity of the European standards prevents the bunching around the weight cutoffs that Ito and Sallee (2014) observe in Japan.

3. Estimating the Frontier: Technical Tradeoffs among Vehicle Characteristics

3.1 Empirical Strategy

In this section we estimate technology frontiers for the United States and Europe. Because the United States has historically regulated the harmonic mean of fuel economy and Europe regulates the CO_2 emissions rate, we estimate a fuel consumption rate frontier for the United States and an emissions rate frontier for Europe.[7]

Before defining the frontier, we briefly summarize the typical vehicle design process. Engines are redesigned every 5-10 years and vehicle models every 4-6 years (Klier and Linn 2012). During an engine or model redesign, the manufacturer can implement large changes. For example, the manufacturer could add technologies to the engine such as cylinder deactivation, which allows an engine with 6 or 8 cylinders to effectively use only 3 or 4 cylinders when the engine is not under a heavy load. Similarly, during a model redesign, the manufacturer could trade off a sedan's cabin space for trunk space. Between redesigns, however, only smaller changes are possible, such as changing the number of transmission speeds or re-tuning the engine.

We focus on the tradeoffs between weight, horsepower and fuel consumption rate that are possible within and across design cycles of both engine and vehicle. We define the frontier as the minimum fuel consumption rate (or emissions rate) given a particular horsepower and weight, and holding fixed other vehicle characteristics that cannot be altered without a redesign (such as the number of engine cylinders). The curvature of the frontier represents medium-run tradeoffs—for example, between the fuel consumption rate and horsepower—that are possible by adopting technology used in other versions of the same vehicle model, without redesigning it. The long-run shifts of the frontier include technology adoption that is feasible between redesigns and that does not impose the same tradeoffs between characteristics as in the medium run. Shifts of the frontier thus represent technology adoption during redesigns. Note that our definition of the frontier differs from Knittel (2011), who defines the frontier as the fuel economy given horsepower, weight, and other characteristics, and holding production costs constant. Instead, we define the frontier based on design cycles because a) we do not observe production costs, which

[7] We compute the harmonic mean using the reciprocal of fuel economy—i.e., the fuel consumption rate. Thus, using the fuel consumption rate, rather than fuel economy, is consistent with the form of the US standards. It also fosters comparability with the European analysis, as the fuel consumption rate is proportional to the CO_2 emissions rate.

makes it difficult to control for costs in the estimation; and b) we do observe redesigns, which allows us to distinguish between movements along the frontier before a redesign, and shifts of the frontier across redesigns.

We define the efficiency of the power train as the amount of useful energy produced by the engine. The fuel consumption rate depends on the power train efficiency as well as vehicle characteristics such as weight and horsepower. That is, holding efficiency fixed, the manufacturer can trade off the fuel consumption rate for other characteristics. Our objective is to estimate both the magnitudes of these tradeoffs—that is, the shape of the frontier—as well as the efficiency of the power train at any point in time. We use a linear regression equation that is similar to Knittel (2011) and Klier and Linn (2012):

$$\ln e_{it} = \beta_0 + \beta_h \ln(h_{it}) + \beta_w \ln(w_{it}) + \tau_{imt} + X_{it}\delta + \varepsilon_{it}, \quad (1)$$

where e_{it} is the fuel consumption rate (for the US analysis) or CO_2 emissions rate (for the European analysis) of model version i in model-year t; h_{it} and w_{it} are horsepower and weight; τ_{imt} is a set of redesign by model by model-year interactions; 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 β_0, β_h, β_w, and δ are parameters to be estimated.

Equation (1) can be estimated separately for the United States and Europe. For the US analysis, the dependent variable is the log fuel consumption rate; for the European analysis, the dependent variable is the log CO_2 emissions rate.

Model by model-year interactions, τ_{mt}, represent the average power train efficiency across different versions of the same vehicle model in a particular model-year (more specifically, the change in efficiency relative to the first year of the sample). A change in τ_{mt} between consecutive years for a particular model is interpreted as the change in the fuel consumption rate if all of the technology adopted between the two years was used to decrease the fuel consumption rate, and all other variables in equation (1) are held fixed. Because all other characteristics are held fixed, changing the frontier corresponds to a change in efficiency. Importantly, because we estimate equation (1) by ordinary least squares (OLS), we interpret the frontier shift as the potential change in the *average* log fuel consumption rate across all versions of a vehicle model. We estimate equation (1) by OLS to maintain consistency with Knittel (2011) and Klier and Linn (2012).

Versions of a particular model often are sold with different engines. Because engine redesign cycles do not always correspond to vehicle model redesign cycles, it is possible that the frontier shifts by different amounts between two years for different versions of the same model. To allow for this possibility, equation (1) includes the triple interaction of τ_{mt} with a dummy variable that is equal to one for versions that have a redesigned engine. To construct this dummy variable, for each version of a vehicle model we match the set of engine programs sold with that version. A version is considered to be redesigned if it is sold with an engine program that is redesigned in year t. The triple interaction, τ_{imt}, allows for the possibility that the frontier shifts more for versions that are sold with a redesigned engine than for versions sold with an engine that has not been redesigned. We expect the interactions of redesign, model, and model-year to decrease over time as manufacturers adopt technology that causes the frontier to shift. The hypotheses are analogous for the European analysis, in which the dependent variable is the emissions rate.

The coefficients on weight and horsepower capture the tradeoffs among vehicle fuel consumption/emissions rates, weight, and horsepower across versions of the same model. Because of the redesign by model by model-year interactions, the coefficients are interpreted as the possible tradeoffs between these characteristics in a particular year without a redesign. For example, the manufacturer could replace steel parts with aluminum parts to reduce weight, or retune the engine to trade off horsepower for fuel economy across versions of the same model. The weight and horsepower coefficients are expected to be positive. If the technology frontiers for European and US vehicles have the same curvature, the coefficients in equation (1) would be equal. The variables in X_{it} control for other determinants of the fuel consumption rate or emissions rate that are fixed between redesigns.

We note that identification of equation (1) differs fundamentally from identifying a consumer demand or hedonic price model. Equation (1) characterizes a technical relationship between fuel economy and certain other vehicle attributes. Equation (1) omits many vehicle attributes that consumers may care about, such as the quality of the sound system. Omitting sound system quality would not bias the estimated coefficients even if the quality is correlated with included variables such as horsepower, however. The reason is that sound system quality only affects fuel economy via weight, and weight is included in equation (1). The estimates in equation (1) would be biased if omitted variables are a) correlated with included variables and b) affect fuel economy independently from the effects of the included variables. In practice, the regression models appear to include most determinants of fuel economy, as indicated by the very high R-squared values reported below.

In summary, equation (1) has several important features. First, we allow the tradeoffs between fuel consumption/emissions rates and other characteristics to depend on whether an engine 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 either the direction or rate of technology adoption.

3.2 Estimation Results

Table 2 shows the estimates of equation (1) 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 US and European cars use horsepower; our regressions for US light trucks use torque, which, for light trucks, is more highly correlated with the fuel consumption rate than is horsepower.

Fuel consumption rate, 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 the fuel consumption rate 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 increases the log fuel consumption rate by about 0.24, which is significant at the 1 percent level. The estimate is significantly larger than in Klier and Linn (2012) 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 in Klier and Linn (2012) for the same reason. The horsepower and weight coefficients also differ from Knittel (2011), but so do the sample periods, independent variables, and data sources.

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

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. The car and light truck hybrid coefficients are essentially the same. The coefficient on flex-fuel vehicles is positive, reflecting the lower energy content of E85 compared to gasoline.

Table 3 reports results for Europe; here the dependent variable is the log emissions rate rather than the log fuel consumption rate. 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 flex-fuel vehicles, but column 1 is otherwise comparable to the US car regression.

Because the European regressions include only passenger cars, we compare the European results with the US car results. The magnitudes of the European horsepower and weight coefficients are very similar to those of the US estimates. The European diesel fuel coefficient is smaller than the US coefficient, but this is because diesel fuel has a higher carbon content than gasoline; if we use the fuel consumption rate 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 US 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 position for each model trim and year. For consistency with the US 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 reported in column 2 are quite similar to those in columns 1 of Table 3.

3.3 Robustness

How confident are we in translating these estimates into a production possibilities frontier? Perhaps the main threat to the identification of possible shifts of the frontier is that the shape of the frontier—as captured by the coefficients on the vehicle characteristics—is assumed to be the same across models as well as over time. If the shape does vary across models or over time, the estimated model by model-year by redesign interactions would be biased. To address this possibility, we allow for variation in the shape of the frontier in several ways. First, in the main analysis we estimate technology adoption separately for cars and light trucks; the differences between columns 1 and 2 in Table 2 illustrate the importance of doing so. In addition, we investigate a further breakdown of each of those two categories, as they, in turn, encompass vehicles with rather different characteristics. Appendix Table 1 separates the categories further, reporting results by market segment. For cars we distinguish three market segments (small, medium, and large/luxury), and for light trucks four segments (crossovers, sport utility vehicles, vans, and pickup trucks). Coefficients are found to vary substantially across segments; for example, weight and horsepower have larger effects on the fuel consumption rate for small cars than for other car segments. Appendix Table 2, which reports separate regressions by European car market segment, shows that the coefficients vary somewhat across segments, but less so than in the US segment-level regressions in Appendix Table 1. Because of this variation in technical tradeoffs across segments in both the United States and Europe, we continue with the more detailed segment analysis in the second stage of our estimation (see section 4.3).

We also examine whether the tradeoffs between vehicle characteristics along the frontier vary by company or over time. We find some variation in the tradeoffs by company, but allowing for this variation (e.g., by estimating equation (1) by company) does not affect the main results in the next section. Likewise, allowing for changes in the tradeoffs over time (e.g., by interacting the vehicle characteristics in equation (1) with a linear time trend) does not affect the results.

As a final robustness check, following Knittel (2011), we introduce higher-order interactions of the vehicle characteristics to equation (1), such as the interaction between log horsepower and log weight. This partially relaxes the functional form assumptions in equation (1) about the tradeoffs, but it does not affect the second-stage results.

4. Have Standards Affected the Rate and Direction of Technology Adoption?

In this section, we use the estimates of equation (1) to investigate whether the recent changes in US and European standards have affected the rate and direction of technology adoption. We first report qualitative aggregate results followed by quantitative cross-sectional results.

4.1 Hypotheses for Aggregate Rate and Direction

We first consider whether the market-wide average rate or direction of technology adoption changed subsequent to the adoption of the new standards. We define the aggregate *rate* of adoption of fuel consumption technology in a specific year as the change between the current and previous year in the market-wide average estimate of τ_{imt} from equation (1). The change represents the decrease in the average log fuel consumption rate, relative to the previous year, if all of the adopted technology were used to decrease the fuel consumption rate—that is, if manufacturers held fixed other vehicle characteristics. We define the *direction* of technology adoption as the log of the ratio of the fuel consumption rate to horsepower or weight, respectively (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 simply ask whether the average rate and direction changed after the standards changed. We consider the US light truck fuel economy standards adopted in 2003, the US car and light truck fuel economy standards adopted in 2007 (and tightened in 2009), and the European CO_2 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. Mostly because of the fact that models and engines are redesigned on regular cycles, manufacturers often begin to respond to future standards before they are fully implemented. For example, they may introduce vehicles with lower fuel consumption rates prior to the actual increase in the standards, either because the regulator offers credits for pre-standard increases in fuel consumption, or because the manufacturer would otherwise have to wait until several years after the new standards are implemented for the next opportunity to redesign the vehicle and reduce its fuel consumption rates. Consequently, we start looking for evidence of manufacturer responses once the tighter standards are adopted.

4.2 Aggregate Results

Figure 2 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 dashed curves indicate the 95 percent confidence intervals). The frontier line indicates that the average fuel consumption rate of US cars would have been 12 percent lower in the year 2010 than in 2000 if all vehicle characteristics besides fuel consumption rate had remained at their 2000 levels, in which case the new technology would have entirely been used to reduce fuel consumption rates. The red line is the change in the actual average fuel consumption rate compared to the year 2000. The other lines in Figure 2 represent a series of simple counterfactuals. To construct each line, we suppose that the corresponding characteristic—for example, horsepower for the green curve—remains fixed at its initial level rather than changes over time. For example, horsepower increased by about 25 percent for US cars between 2000 and 2012 (see Figure 1). In the counterfactual, we suppose that horsepower had instead remained fixed at its 2000 level and that the foregone horsepower increase had been used to reduce the fuel consumption rate. This calculation uses the horsepower coefficient in equation (1) and shows that, if horsepower of US cars had remain constant at its 2000 level rather than increasing 25 percent, fuel consumption rates of US cars would have been about 6 percent lower in 2012 than they actually were. By construction, in the figure the sum of the change in characteristics is equal to the frontier shift.

The figure shows that the average rate and (in most cases) the direction of technology adoption changed soon after the standards changed. Regarding the rate, for US cars (Panel A), the frontier shifted out twice as quickly after 2007 as compared to 2000–2007. For US 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. For reference, for each time period the figure notes report the average rates at which the frontier increased.

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

Figure 2 illustrates the market-wide average patterns; Appendix Figures 1–3 supplement that with company or brand-level detail. The figures are constructed similarly to Figure 2, except that each panel represents a different company (in the United States) or brand (in Europe). [8] The figures illustrate considerable cross-firm heterogeneity in the rate and direction of technology adoption, but most firms exhibit patterns similar to those shown in Figure 2.

4.3 Hypotheses for Cross-Sectional Rate and Direction

Although the aggregate results suggest that the introduction of tighter standards affected both the rate and direction of technology adoption, there may have been confounding influences. For example, gasoline prices began rising in 2003. Given vehicle design lags of three years or

[8] The fuel consumption rate decreased noticeably in 2010 for several manufacturers. Starting in 2010, the fuel consumption rate calculated from Wards data for flex-fuel vehicles corresponds to the fuel consumption rate using 85 percent ethanol rather than gasoline. Those manufacturers sell a large share of flex-fuel vehicles.

more, rising gasoline prices may have affected the rate and direction of adoption as early as 2006. This subsection presents an 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 (US light trucks in 2003, US 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 are likely to vary across manufacturers, depending on how close they are to achieving the new standard. In both the United States and Europe, the new standards require all manufacturers to reduce fuel consumption and emissions rates. We define the *stringency* of standards as the difference between a manufacturer's pre-standard fuel consumption or emissions rate and the level of the new standard. A manufacturer with a high fuel consumption rate (and low fuel economy) therefore has a higher value of the stringency variable.[9] Stringency can vary across manufacturers because of different levels of pre-standard fuel economy or because of differences in the mix of vehicles offered (recall that the US standards depend on the vehicle's footprint whereas the European standards depend on weight). Accordingly, in a footprint-based system and similarly for the European weight-based system, manufacturers with larger vehicles are subject to higher fuel consumption rate requirements.[10]

Cross-manufacturer variation in stringency enables a differences-in-differences strategy for identifying the effects of standards on the rate and direction of technology adoption. We estimate separate regressions for US cars, US light trucks, and European cars. For the <u>rate</u> of technology adoption, each regression is a variation of the following equation:

$$\hat{\tau}_{imt} = \gamma_R S_F Post_t + Seg_m Post_t \ln(e_m) + \theta_t + \omega_m + \upsilon_{it} \qquad (2)$$

---

[9] We assume that all firms elect to meet the standards. Historically, in the US market, several firms, such as Mercedes, have elected to pay fines instead of meeting the standards. However, beginning in 2011 the EPA and NHTSA jointly regulate greenhouse gas emissions rates and fuel economy, and the EPA fines under the Clean Air Act are orders of magnitude higher than the historical NHTSA noncompliance fines. Hence all firms are expected to comply with the new standards, which is consistent with observed behavior in the early years of the 2016 standards.

[10] One might be concerned that provisions in the US and European regulations that equalize marginal compliance costs across firms, such as cross-firm credit trading, would reduce or even eliminate the cross-firm variation in stringency. However, the identification strategy is valid even if marginal compliance costs do not vary across firms. Even in that case, the stringency variable would be proportional to the regulatory pressure caused by the regulations (Roth 2014). In addition, the very limited credit trading to date suggests that firms are unlikely to overshoot the standards, and that marginal costs are unlikely to be equated across firms (Leard and McConnell 2015).

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 (1). 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 consumption 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 US 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 consumption or emissions rate of the corresponding model in the initial year of the sample. Note that when estimating equation (2), we include all lower-order terms for the triple interaction; we omit these terms in the expression for brevity. Later in the subsection, we discuss how the triple interactions address concerns about gasoline prices and other possibly confounding factors.[11]

Equation (2) 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 over the sample.

The central hypothesis to be tested is that γ_R is negative. To illustrate the differences-in-differences interpretation of γ_R, suppose the average frontier for manufacturer A shifts 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 greater 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 more quickly than the average frontier for B after the new light truck standards were adopted (recall that a decrease in the fuel consumption rate corresponds to an increase in fuel economy). Note that this approach cannot distinguish between a case in which the standards

[11] Rather than estimating equations (1) and (2) separately, we could replace the triple interactions of model by model-year by redesign in equation (1) with the independent variables in equation (2). This would increase the efficiency of the estimates, but the two-stage approach allows for an unobserved component (essentially, a random effect) of the model by model-year by redesign triple interaction.

caused a one-time frontier shift and a case in which the standards caused the frontier to shift 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.

In this differences-in-differences research design, the identifying assumption is that stringency is uncorrelated with deviations from pre-policy trends in rate or direction. That is, we assume that pre-trends in rate or direction are unbiased estimates of the counterfactual trends that would have occurred in the absence of the changes in stringency. In principle, unobserved demand or supply-side shocks could cause the estimate of γ_R to be biased or spurious. On the demand side, changes in consumer preferences that are contemporaneous with the stringency changes could cause manufacturers to change rate or direction, and the manufacturer responses could be correlated with stringency. For example, an increase in consumer demand for fuel economy in the mid-2000s could cause manufacturers to trade off horsepower and weight for fuel economy. Ford, which has a higher measure of stringency than Toyota and which typically produces cars with lower fuel economy than Toyota, could respond more to such a demand shock. On the supply side, changes in technology costs could affect rate or direction. For example, innovation that reduces the cost of improving efficiency of large cars but does not affect the cost of improving efficiency of small cars would cause the rate to increase, and this increase could be greater for manufacturers that sell more large cars than small cars.

The triple interaction term in equation (2), $Seg_m Post_t \ln(e_m)$, controls for such demand or supply shocks that occur in the post-policy periods and that are correlated with stringency, to the extent that these shocks affect rate or direction proportionately for all vehicles in the same market segment, or proportionately to the vehicle's initial fuel economy. Returning to the innovation example, if the cost-reducing innovation affects rate and direction similarly for all large cars, or in proportion to those cars' initial fuel economy, the estimate of γ_R would be unbiased. Likewise, the estimate would be unbiased if changes in consumer demand for fuel economy affect rate and direction in proportion to initial fuel economy, or affect rate and direction proportionately for all vehicles in the same market segment. After presenting the main results in Section 4.4, we present evidence in Section 4.5 supporting the identifying assumption.

Next, we turn to the direction of technology adoption. A change in standards could cause manufacturers to move along the frontier by reducing the fuel consumption rate and reducing horsepower or weight.[12] We define a set of direction variables, dir_{it}, at the model version level. The fuel consumption rate–horsepower direction, for example, is the log of the ratio of fuel consumption rate to horsepower. Direction variables for fuel consumption rate–torque and fuel consumption rate–weight are defined similarly. The hypothesis to be tested is that an increase in the stringency of the fuel economy standard causes a change in the direction of technology adoption towards a reduction in the fuel consumption rate relative to torque, horsepower, or weight. We estimate the equation

$$dir_{it} = \gamma_D S_F Post_t + Seg_m Post_t \ln(e_m) + \theta_t + \omega_m + \upsilon_{it} \, . \quad (3)$$

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, one each 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 higher value of S_F move toward a lower fuel consumption rate 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 a lower fuel consumption rate. 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 new car and light truck standards were adopted and the tighter light truck standards, adopted in the previous

[12] The standards could cause manufacturers to trade fuel consumption for any vehicle characteristics in equation (1). We focus on weight and horsepower (torque) primarily because these two variables are continuous, making it simpler to define the dependent variable in equation (3). In principle we could define the direction variables in terms of fuel economy rather than fuel consumption rates, but this would not be consistent with the form of the standards as noted in the text.

period, took effect; and 2010–2012 as these standards took effect (subsequent tightening of the standards through 2025 occurred after the end of the sample period). 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 stringency and the 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 first two periods and the last two periods.

Panel A of Table 4 shows results from estimating equation (3), 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 for periods 2, 3 and 4: there the standards caused the direction of technology adoption to shift toward reducing the fuel consumption rate at the expense of torque (and, to a lesser extent, of weight).

To interpret the magnitudes for light truck torque, we consider a manufacturer for which the stringency is one standard deviation above the mean in the second time period (i.e., a manufacturer with a high initial fuel consumption rate compared to its standard). For example, Mitsubishi is near the mean stringency for US cars, whereas the standard is about one standard deviation more stringent for Volvo and one standard deviation less stringent for Toyota. The size of the estimated effect implies that Volvo decreased torque and decreased the fuel consumption rate 5 percent (about 1 mpg) compared to Mitsubishi (the estimate of 5 percent is obtained by multiplying one standard deviation of the stringency variable by the light truck coefficient for 2007-2009 in Table 4). Given that the average light truck fuel consumption rate decreased 3 mpg during the sample period, the movement along the frontier represents a substantial fuel consumption rate decrease.

Panel B of Table 4 shows results from equation (2), 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); the coefficient in the final time period is smaller and is only marginally statistically significant.

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

In comparing the rate and direction estimates, we find them to be larger and more precise for light trucks than for cars. Differences in technological opportunities or preferences for fuel economy could explain this result. Alternatively, this difference could be explained by the fact that the standards for cars only tightened toward the end of the sample; the aggregate analysis in Section 4.2 suggested that the rate of adoption increased after 2007, but the statistical evidence in the cross-sectional estimation suggests that this response was correlated with potentially confounding factors omitted from the aggregate analysis. 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 change in US standards.[13]

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 [3]) and Panel B on the rate (equation [2]). 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 the direction of technology adoption away from horsepower and weight. 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 US 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 the European standards had a relatively small, but statistically significant, effect on the direction (horsepower and weight) as well as on the rate of technology adoption.

[13] Above we noted that the frontier estimates in equation (1) differ from those by Knittel (2011). Using the data from that paper and applying our estimation strategy, 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 Table 4.

4.5 Potential Omitted Variables Bias

The fact that from 2003 to 2009 the adoption rate increased for US light trucks but not for US 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 factors represent the primary threats to the validity of equations (2) and (3). As noted, we control for these factors using triple interactions of fuel consumption rate 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 consumption rate. For example, firms may change the rate or direction for reasons other than the standards, such as a firm-level technology cost shock or a strategic change in product positioning (we note that changes in product positioning *caused* by the standards would not constitute bias). This section reports four approaches to assess the magnitude of any biases in this research design.

First, in the main regressions, we assume that the stringency variable is exogenous after controlling for segment-level shocks to technology adoption. While this approach controls for segment-level shocks that affect technology adoption, one might be concerned that segment 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, without 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 US light trucks, weaker evidence for the rate and direction for US cars, and evidence that the European standards affected the rate and direction.

Furthermore, firm-level changes in rate or direction for strategic reasons or other reasons uncorrelated with the standards are likely to occur gradually over time, in contrast to the discrete changes in stringency of the standards. If we add company-specific time trends γ_R is identified off of the linear time trends. The US results are similar if we add interactions of linear time trends with company fixed effects (Appendix Table 5), which supports 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 fuel economy 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 consumption rate. 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 the fuel consumption rate and sales.) Importantly, the regressions also include the triple interaction of time period, market segment, and initial model fuel consumption rate. 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.

Third, gasoline prices or the recession could affect technology via market shares. For example, if gasoline prices raise the market share of vehicles with low fuel consumption rates, manufacturers would have greater incentive to adopt technology that reduces the fuel consumption rate of such 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 either 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 consumption rate or aggregate sales. Statistically significant or large interaction coefficients would raise concerns that the other variables in equations (2) and (3) do not adequately control for the effect of fuel prices or the recession on market shares. We find only the fuel price coefficient for European cars to be statistically significant. 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 (2) and (3) the interactions of gasoline prices with the stringency–time period interaction. The main results (not reported but available upon request) are unaffected.[14]

[14] The Appendix reports further robustness checks on the main rate and direction results. Panel B of Appendix Table 5 adds interactions of a linear time trend to segment fixed effects to allow for segment-level deviations from the pre-policy time trends. Appendix Table 6 reports US results if we add interactions of a linear time trend with weight and horsepower in equation (1), which allows for the possibility that the frontier tradeoffs change smoothly over time. Appendix Tables 7 and 8 report direction results by market segment.

5. The Opportunity Costs of Standards

The empirical results suggest that tighter fuel economy standards increased the rate of technology adoption in both the United States and Europe and also affected the direction of technology adoption for US light trucks and European cars. In that context, one could quantify the opportunity cost of technology adoption by comparing equilibria with and without tighter standards and computing the reduced willingness to pay for characteristics other than the fuel consumption rate. In this section, we use our results to estimate the opportunity cost of tightening the standards.

We focus on opportunity costs because a full welfare analysis of the standards is beyond the scope of this paper, as it would require the specification of a dynamic model of manufacturer technology adoption, the choice of vehicle characteristics, as well as consumer demand. Instead, we assume that the standards do not affect vehicle prices or market shares. Klier and Linn (2012) suggest that, over a period 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 multi-year time horizon considered here, the assumption of constant market shares may therefore not be very strong. This assumption avoids the need to introduce a structural model of vehicle demand and supply and estimate additional parameters. The assumption of constant market shares may cause us to overestimate costs, and we therefore treat our results as rough estimates of the opportunity costs. On the other hand, imposing the assumption of constant market shares allows us to isolate the implications of technology adoption for opportunity costs by eliminating manufacturers' price responses from the simulations.

Although we could base the simulations on the actual change in standards, we consider the same change in standards relative to the observed standards to compare results across the United States and Europe. We define the baseline as the observed level of the standards for US cars, US light trucks, and European cars. For the baseline scenario we use the predicted values of the frontier, horsepower, weight, and fuel consumption rate or emissions rate, as obtained by equations (1) – (3), referencing the levels of the standards in the final time period of the sample (2010 – 2012 for the United States and 2008 – 2010 for Europe). Against this baseline we compare the outcomes of a situation where the standards require fuel consumption rates to be 10 percent lower instead. The level of the hypothetical standard is chosen so that the resulting stringency is equivalent to a one-standard deviation increase in stringency—consistent with the in-sample variation used to identify the coefficients. We use the estimating equations, combined with the counterfactual stringency levels to predict the weight, horsepower, and fuel economy in both the observed and counterfactual levels of the standards.

Table 7 presents the results from these simulations, with each row showing results for the indicated set of vehicles. The first column shows the effect of the tighter standards on the frontier, representing the percent efficiency increase relative to the baseline. The remaining columns show the percentage changes in horsepower (or torque for US light trucks) and weight relative to the no-policy case, the consumer willingness-to-pay for the lost horsepower, and the value of the fuel savings.[15] The willingness to pay for the lost horsepower is computed using estimates in Klier and Linn (2012) of $10 per horsepower per ton (in 2007 dollars). This value corresponds to the lower-end value of willingness to pay as reported in the vehicle demand literature (Whitefoot and Skerlos 2012); the opportunity costs reported in Table 7 should therefore be considered conservative estimates.[16]

For the US simulations, the opportunity costs—as measured by the willingness to pay for lost horsepower—are smaller than the value of the fuel savings (see column 5) but are nonetheless economically significant.[17] 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, due to the lower horsepower levels in the European than the US market, even a large percent reduction in horsepower translates to a relatively small reduction in levels.

6. Conclusion

Recent increases in fuel economy and greenhouse gas emissions rate standards require a substantial increase in passenger vehicle fuel economy in the United States, in Europe, as well as in other regions. Economic theory suggests that tighter performance standards increase the rate

---

[15] 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 US 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 US consumers except for fuel prices.

[16] We do not allow for heterogeneous preferences for vehicle characteristics, which could affect the welfare analysis (Bento et al. 2012). In addition, we assume that consumers care about absolute horsepower rather than horsepower relative to other vehicles on the road; the empirical literature on consumer demand has not made this distinction.

[17] 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 US cars.

of technology adoption. Because vehicle manufacturers choose multiple vehicle characteristics in designing a vehicle and because technical tradeoffs exist across some of these characteristics, theory also suggests that the tightening of fuel economy standards will likely affect other vehicle characteristics besides fuel economy.

Consistent with these expectations, we find evidence 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 affected the direction of technology adoption by reducing light truck torque in the United States and both vehicle weight and horsepower in Europe. The results are robust to controlling for other, potentially confounding, influences on technology.

Thus, we observe that even the European standards, which are more stringent than the US standards, have a statistically and economically significant effect on the rate of technology adoption, suggesting that as US standards tighten over time, they will continue to increase the rate of technology adoption. The European standards are found to have smaller effects on weight and horsepower than the US standards. In principle, differences in tradeoffs could explain this result, but that is not likely as we estimate similar technical tradeoffs between vehicle characteristics. Instead, the result could arise a) from the fact that European fuel prices are much higher than US fuel prices, which creates greater incentives to trade off horsepower and weight for fuel economy in the United States than in Europe, even if standards had remained unchanged (i.e., in the counterfactual); and b) from the fact that European standards are weight-based, which might create less of an incentive for trading off weight for fuel economy than the footprint-based standards in effect in the United States. Testing these alternative hypotheses is left for future research.

This paper is the first to document the effects of fuel economy standards on the rate of technology adoption. Furthermore, the previous literature has not quantified the opportunity costs of actual standards due to changes in other vehicle characteristics besides fuel economy. We use the empirical results to derive back-of-the-envelope estimates of these opportunity costs. In simulating the imposition of hypothetical standards, we find that the opportunity costs are smaller in size than the value of the fuel savings but nonetheless economically significant. Furthermore, accounting for the effect of the standards on the rate of technology adoption reduces the estimated opportunity costs; the increase in the rate of adoption reduces by about 15 percent the fuel economy improvement attained by trading off weight and horsepower.

We leave for future work further investigation of the mechanisms behind the observed changes in rate and direction, as well as the costs to manufacturers of meeting the standards. For example, manufacturers facing greater stringency may adopt technology that other manufacturers already employ, and there may be international spillovers in which manufacturers adopt technologies in one market that they already use in another market. Future work may also incorporate the opportunity costs into 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 the introduction of fuel economy standards (Allcott 2013); future work could consider the welfare and policy implications of this possibility in a dynamic context.

References

Acemoglu, D., P. Aghion, L. Bursztyn, and D. Hemous (2012). The Environment and Directed Technical Change. *American Economic Review* 102: 131–166.

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.

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.

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.

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.

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

Greene, D. L. (1991). Short-Run Pricing Strategies To Increase Corporate Average Fuel Economy. *Economic Inquiry* 29: 101–114.

Jacobsen, M. (2013). Evaluating US Fuel Economy Standards in a Model with Producer and Household Heterogeneity. *American Economic Journal: Economic Policy* 5(2): 148–187.

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.

——— (2012). New Vehicle Characteristics and the Cost of the Corporate Average Fuel Economy Standards. *RAND Journal of Economics* 43: 186–213.

——— (forthcoming). Using Vehicle Taxes To Reduce Carbon Dioxide Emissions Rates of New Passenger Vehicles: Evidence from France, Germany, and Sweden. *American Economic Journal: Economic Policy*.

——— (2013). Fuel Prices and New Vehicle Fuel Economy—Comparing the United States and Western Europe. *Journal of Environmental Economics and Management* 66: 280-300.

Knittel, C. (2011). Automobiles on Steroids: Product Attribute Trade-offs and Technological Progress in the Automobile Sector. *The American Economic Review* 107: 3368–3399.

Leard, B. and V. McConnell (2015). New Markets for Pollution and Energy Efficiency: Credit Trading Under Automobile Greenhouse Gas and Fuel Economy Standards. RFF Discussion Paper 15-16.

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.

Linn, J. (2008). Energy Prices and the Adoption of Energy-Saving Technology. *The Economic Journal* 118: 1986–2012.

Mazzeo, M., K. Seim, and M. Varela (2013). The Welfare Consequences of Mergers with Product Repositioning.

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.

Popp, D. (2002). Induced Innovation and Energy Prices. *American Economic Review* 92: 160–180.

Porter, M. (1999). America's Green Strategy. *Scientific American* 264: 168.

Roth, K. (2014). The Unintended Consequences of Uncoordinated Regulation: Evidence from the Transportation Sector.

Sweeting, A. (forthcoming). Dynamic Product Repositioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry. *Econometrica*.

US Department of Transportation (2006). *Corporate Average Fuel Economy and CAFE Reform for MY 2008–2011 Light Trucks: Final Regulatory Impact Analysis*. Washington, DC: US Department of Transportation.

US 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: US EPA and NHTSA.

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.

Whitefoot, K. S. and S. J. Skerlos (2012). Design Incentives To Increase Vehicle Size Created from the US Footprint-Based Fuel Economy Standards. *Energy Policy* 41: 402–411.

Figures and Tables

See following pages.

Figure 1. Fuel Consumption Rate, Emissions Rate, Horsepower, and Weight

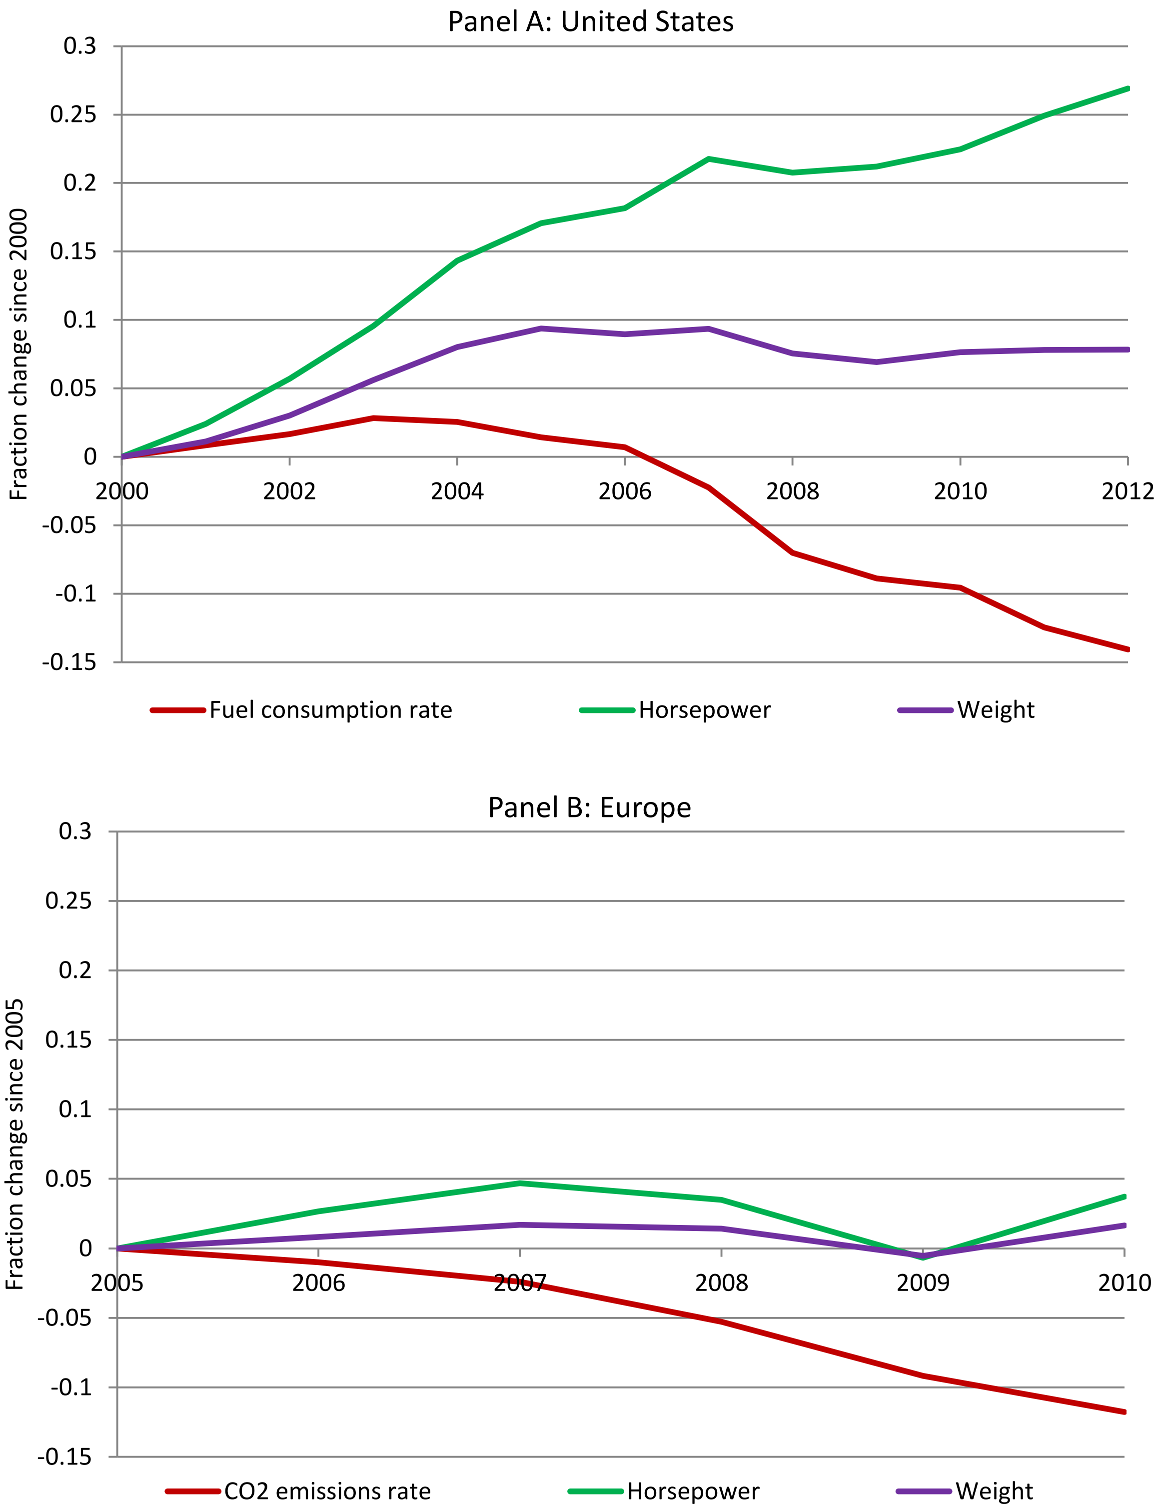

Notes : Panel A plots the fraction change in sales-weighted fuel consumption rate (Panel A), CO_2 emissions rate (Panel B), weight, and power since 2000 for the United States, using the same data set as Table 1. Panel B plots fraction changes in registration-weighted fuel economy, weight, and power since 2005 for Europe, using the same data set as Table 1.

Figure 2. Technology Adoption in the United States and Europe

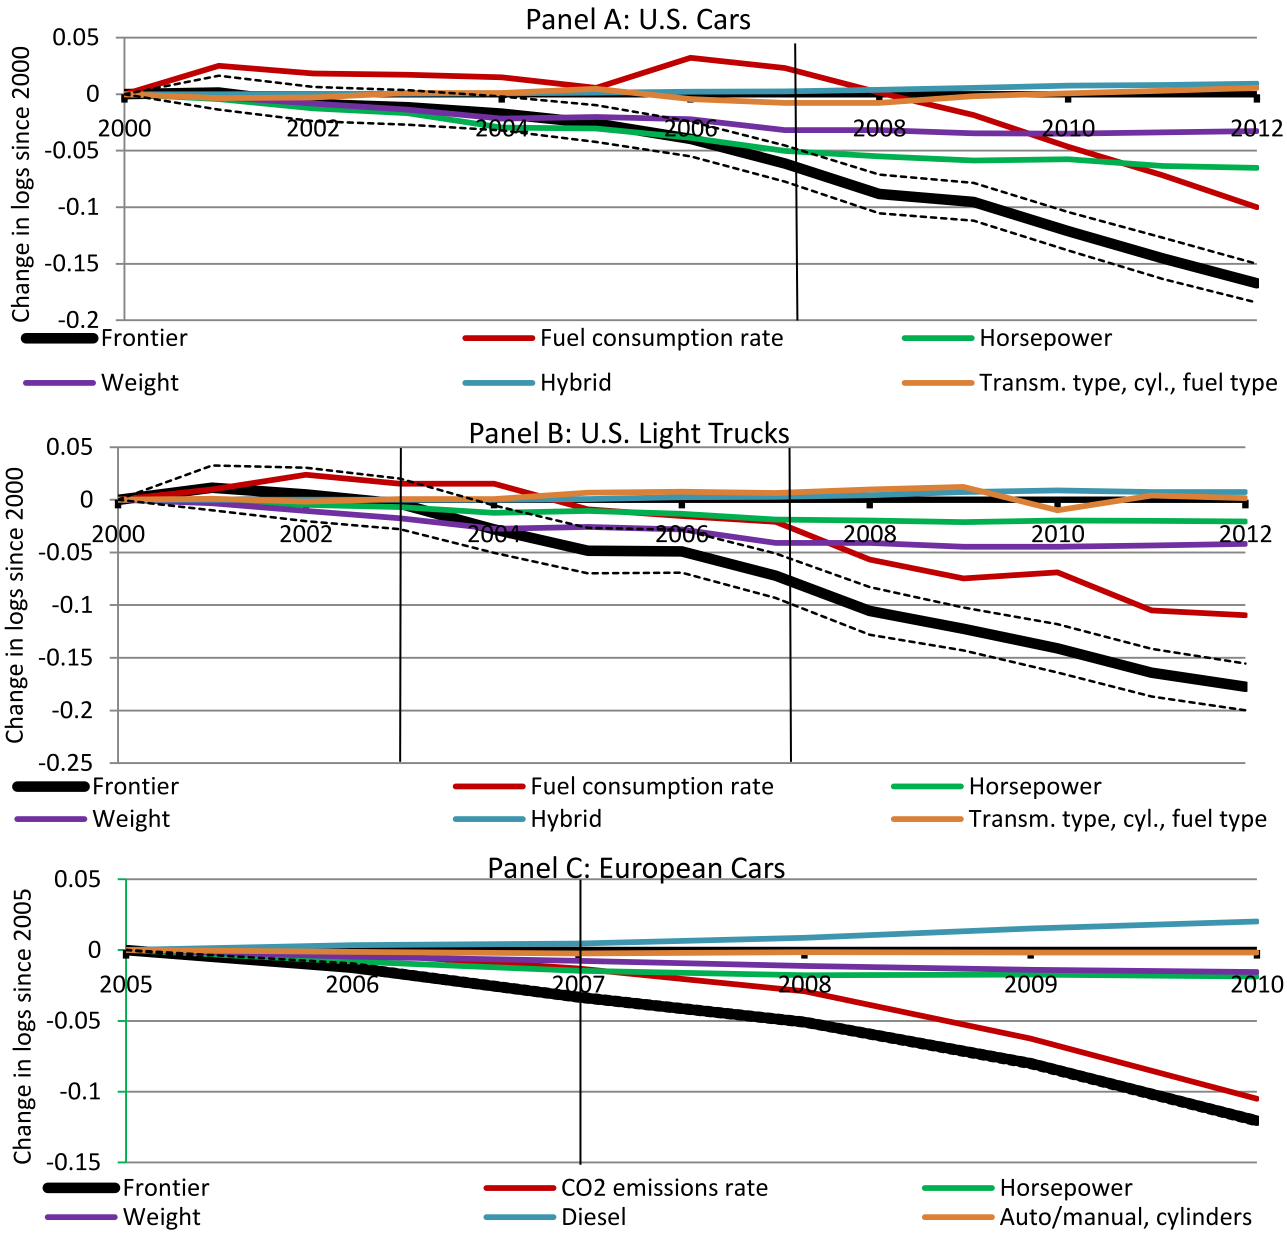

Notes : In Panel A, the frontier plots the change in redesign, model, and model-year interactions estimated in column 1 of Table 2; in Panel B, the frontier plots the interactions estimated in column 2 of Table 2; and in Panel C, the frontier plots the interactions estimated in column 1 of Table 3. In Panels A and B, fuel consumption rate is the change since 2000 in the average log fuel consumption rate across specifications. Horsepower, weight, and diesel represent the decrease in the fuel consumption rate that would have been possible if these characteristics had remained at their 2000 levels; see text for details. The curves in Panel C are constructed similarly, and represent changes since 2005. Vertical lines indicate the adoption of the higher standards. The dotted lines indicate the 95 percent confidence intervals of the estimated frontiers. In Panel A the frontier shifts 0.7 percent per year 2000-2006 and 2.1 percent per year 2007-2012. In Panel B the frontier shifts 0.1 percent per year 2000-2002, 1.5 percent per year 2003-2006, and 2.1 percent per year 2007-2012. In Panel C the frontier shifts 1.7 percent per year 2005-2007 and 2.9 percent per year 2008-2010.

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

	United States		Europe	
	2005	2010	2005	2010
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)
Fuel consumption rate (gallons per 100 miles)	4.16 (0.84)	3.97 (0.99)	3.11 (0.86)	2.79 (0.76)
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 Consumption Rate and Other Vehicle Characteristics

	(1)	(2)
	Dependent variable: log fuel consumption rate	
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 the log fuel consumption rate. Besides the reported variables, the regressions include the variables indicated at the bottom of the table. Units are the same as in Table 1; torque is measured in Newton-meters. 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 CO_2 Emissions Rate and Other Vehicle

	(1)	(2)
	Dependent variable: log CO_2 emissions rate	
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 CO_2 emissions rate. Units are the same as in Table 1. 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)
		Panel A: direction		
Dependent variable	Log (fuel cons rate / horsepower)	Log (fuel cons rate / weight)	Log (fuel cons rate / torque)	Log (fuel cons rate / weight)
Stringency X 2003–2006	0.131 (0.078)	0.070 (0.059)	-0.824 (0.136)	-0.260 (0.085)
Stringency X 2007–2009	0.084 (0.090)	0.022 (0.065)	-0.939 (0.148)	-0.308 (0.094)
Stringency X 2010–2012	-0.029 (0.093)	0.057 (0.071)	-0.749 (0.157)	-0.094 (0.106)
Number of observations	6,856	6,856	11,966	11,966
R^2	0.783	0.623	0.625	0.633
		Panel B: rate		
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 consumption rate to horsepower in column 1, the log of the ratio of fuel consumption rate to weight in columns 2 and 4, and the log of the ratio of fuel consumption rate to torque in column 3. Units are the same as in Table 1; torque is measured in Newton-meters. 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 standard for the corresponding company and vehicle type and the log sales-weighted fuel economy in 2000. The calculation uses the 2016 standards. All regressions include model fixed effects and triple interactions between model-year, market segment, and log fuel consumption rate 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)
	Panel A: direction	
Dependent variable	Log (emissions rate / horsepower)	Log (emissions rate / 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
	Panel B: rate	
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 the CO_2 emissions rate to horsepower in column 1 and the log of the ratio of fuel economy to weight in column 2. Units are the same as in Table 1. 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 2015 standard and the log registration-weighted brand emissions rate in 2005. 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)
		Panel A: exit	
Log fuel cons rate / emissions rate X log fuel price	0.104	0.409	0.040
	(0.219)	(0.257)	(0.036)
Log fuel cons rate / emissions rate X log aggregate sales /	-0.103	0.299	0.017
	(0.324)	(0.349)	(0.109)
Number of observations	5,850	5,524	228,492
R^2	0.162	0.237	0.009
Sample includes	U.S. cars	U.S. light trucks	European cars
		Panel B: sales and registrations	
Stringency X 2003–2006 X log gas price	-2.705	0.850	
	(3.583)	(3.093)	
Stringency X 2007–2009 X log gas price	-6.253	-7.406	
	(8.877)	(6.565)	
Stringency X 2010–2012 X log gas price	-21.801	-3.653	
	(68.811)	(9.256)	
Stringency X 2003–2006 X log aggregate sales	35.060	-10.022	
	(45.445)	(43.094)	
Stringency X 2007–2009 X log aggregate sales	3.173	7.151	
	(5.884)	(4.557)	
Stringency X 2010–2012 X log aggregate sales	-1.067	1.257	
	(8.949)	(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 fuel cons rate X log fuel price is the interaction between log fuel consumption rate and the log of the fuel price. Log fuel cons rate X aggregate sales is the interaction between the vehicle's log fuel consumption rate 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 the fuel consumption rate 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 consumption rate, along with lower-order main effects and interactions, as in Tables 6 and 7. Units are the same as in Table 1; fuel prices are measured in dollars per gallon (United States) and euros per liter (Europe).

Table 7. Effects on Consumer Welfare of a 10 Percent Fuel Consumption Rate Decrease

	(1) Frontier shift (percent fuel consumption rate change)	(2) Percent change in horsepower / torque	(3) Percent change in weight	(4) WTP for hp change (2005 $)	(5) WTP for mpg change (2005 $)
U.S. cars	0.95	-15.86	-15.09	-389	725
U.S. light trucks	1.50	-18.52	-12.70	-595	983
European cars	0.30	-18.71	-18.58	-211	1,487

Notes : The table reports effects of a standard that is 10 percent more stringent than the actual standard in the final time period of the data (2010-2012 for the United States and 2008-2010 for Europe). The counterfactual standards are compared to the baseline scenario using the actual values of the standards and the predicted values of the frontier, fuel economy, horsepower/torque, and weight. The first row shows results for U.S. cars, the second row for U.S. light trucks, and the third row for Europe. Each column reports the difference in the indicated outcome between the actual and counterfactual scenarios, across all vehicles in the data and weighting observations by sales or registrations. The frontier shift is the difference in the technology frontier between the two scenarios, estimated using the coefficients from equation (2). The percentage changes in horsepower and weight are computed using the change in direction estimated from equation (3) and the change in rate estimated from equation (2). 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 4. In column 5, the WTP for the fuel economy increase is calculated using a fuel price of $2.65 per gallon in the United States and $6.15 per gallon in Europe, a 10 percent discount rate, and the vehicle miles traveled and survival estimates in U.S. EPA (2012).

Appendix Figure 1. U.S. Technology Adoption by Company, Cars

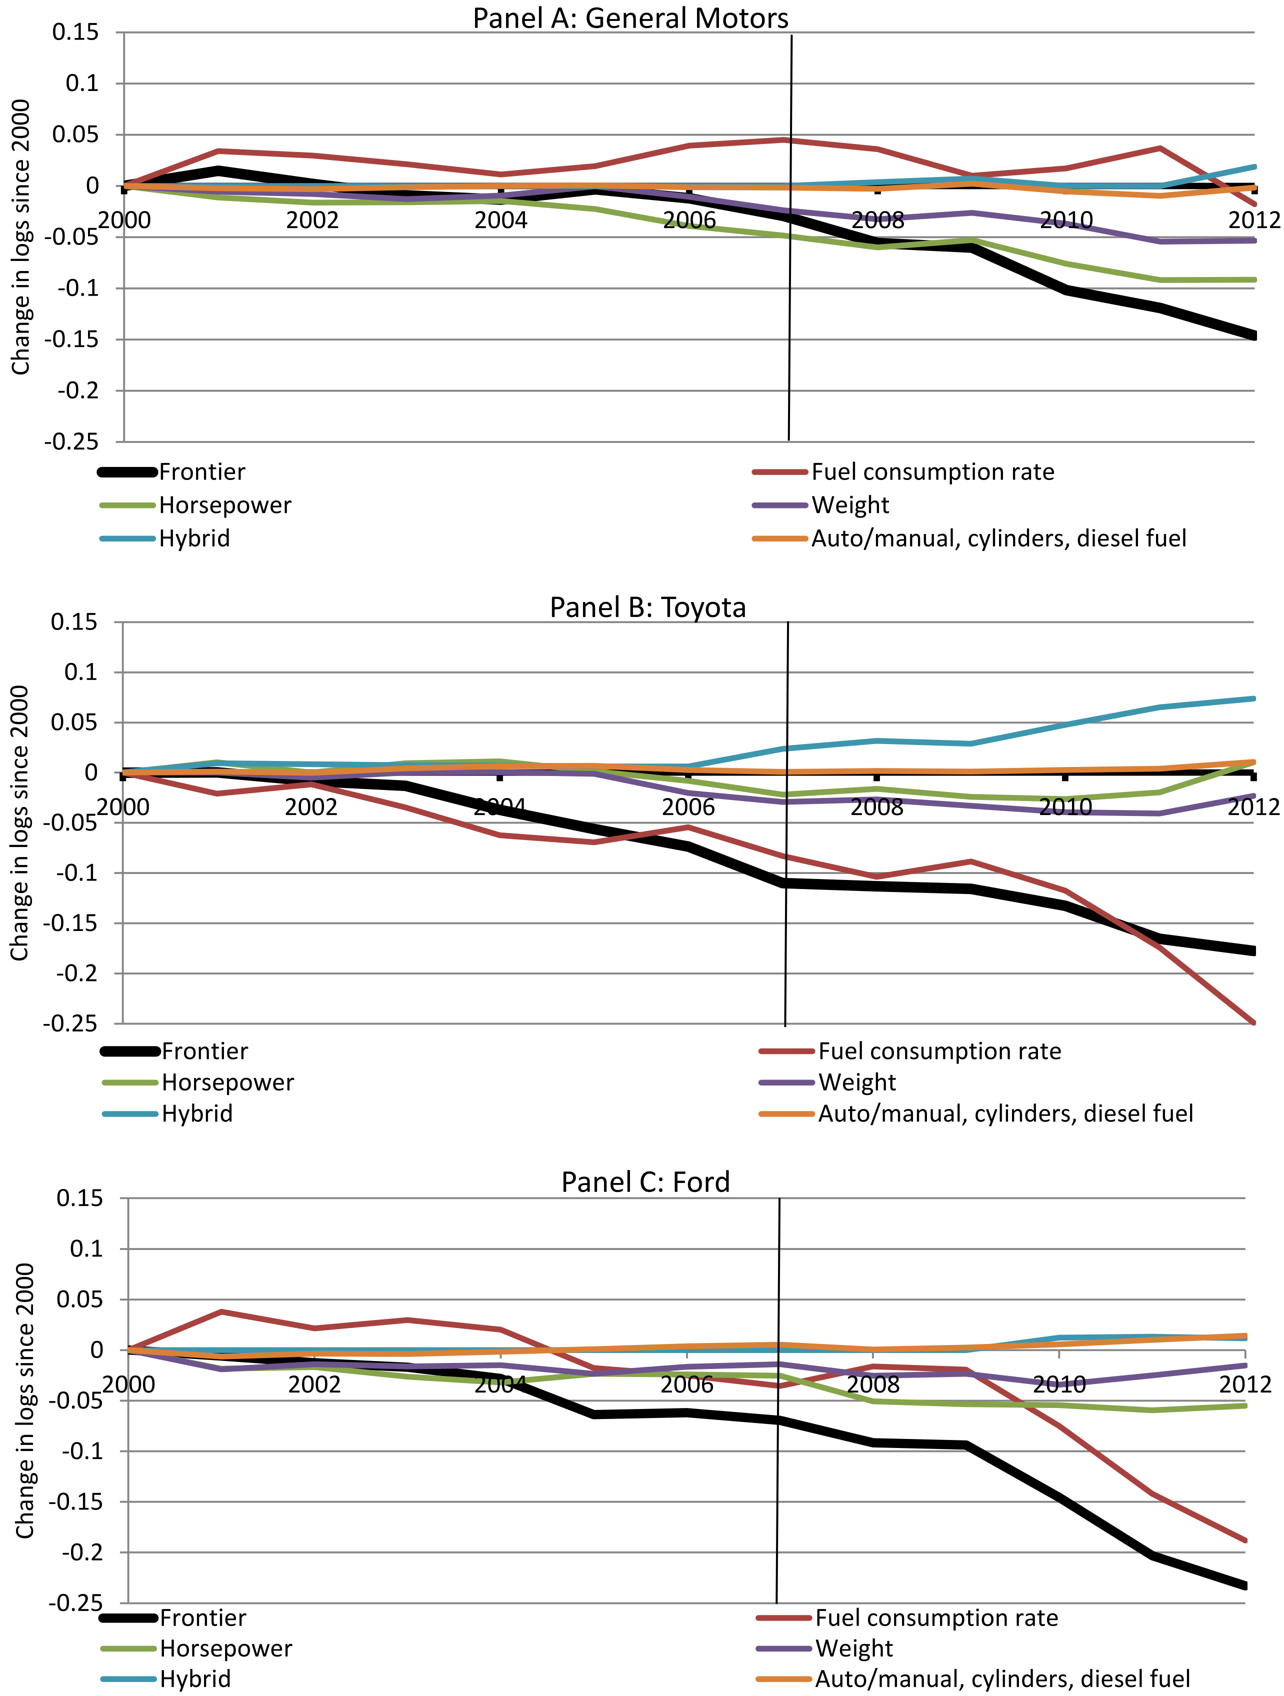

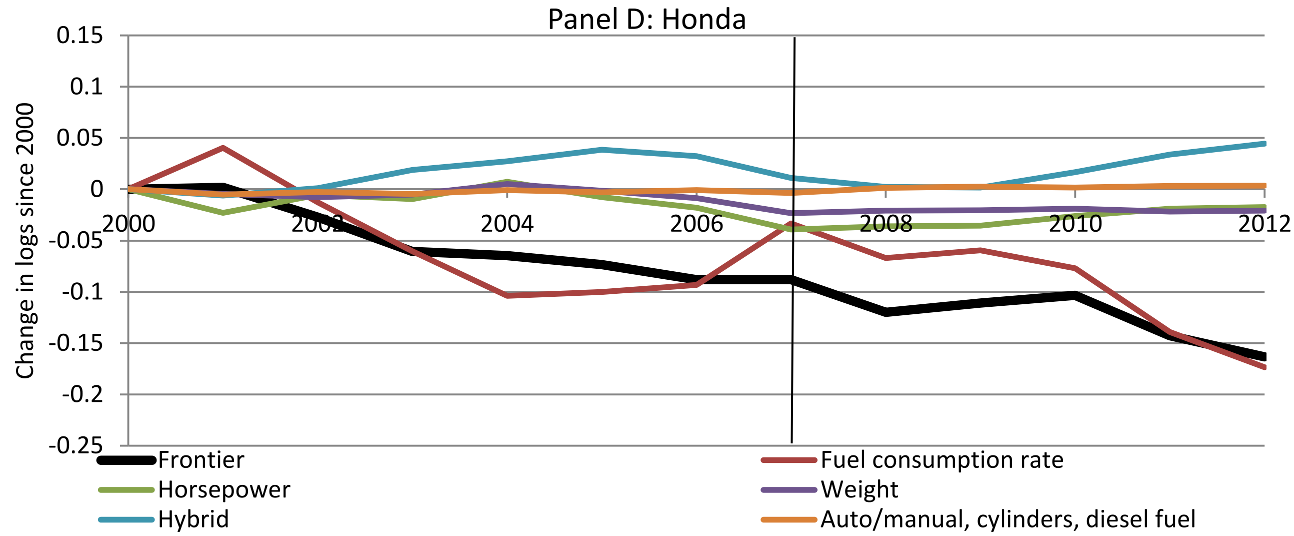
Panel D: Honda
Change in logs since 2000
0.15
0.1
0.05
0
-0.05
-0.1
-0.15
-0.2
-0.25
2000
2002
2004
2006
2008
2010
2012
Frontier
Horsepower
Hybrid
Fuel consumption rate
Weight
Auto/manual, cylinders, diesel fuel

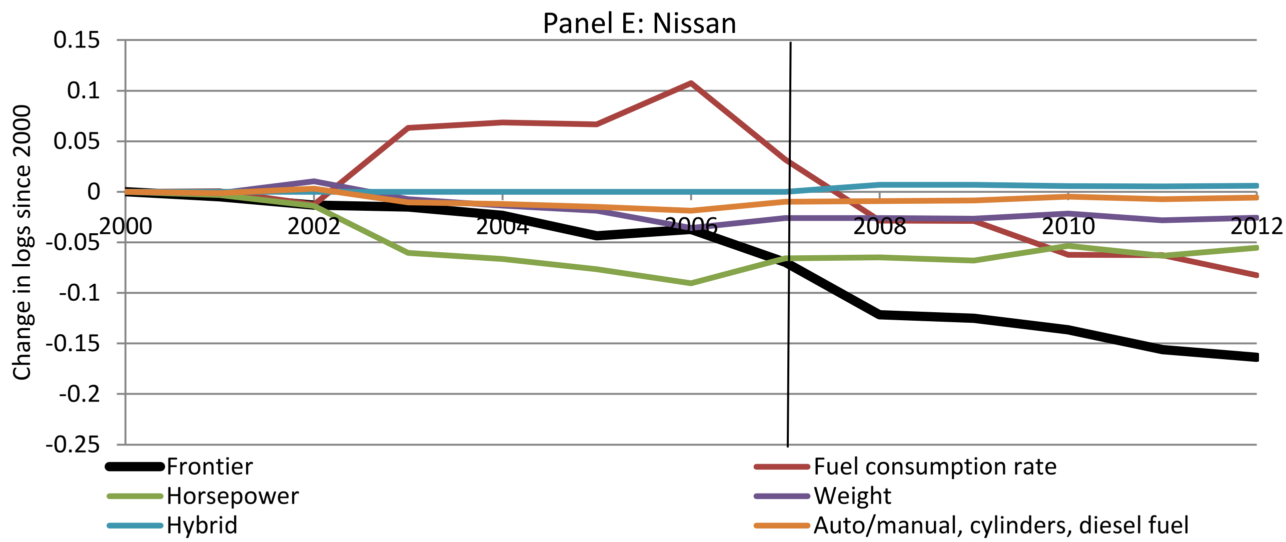
Panel E: Nissan
Change in logs since 2000
0.15
0.1
0.05
0
-0.05
-0.1
-0.15
-0.2
-0.25
2000
2002
2004
2006
2008
2010
2012
Frontier
Horsepower
Hybrid
Fuel consumption rate
Weight
Auto/manual, cylinders, diesel fuel

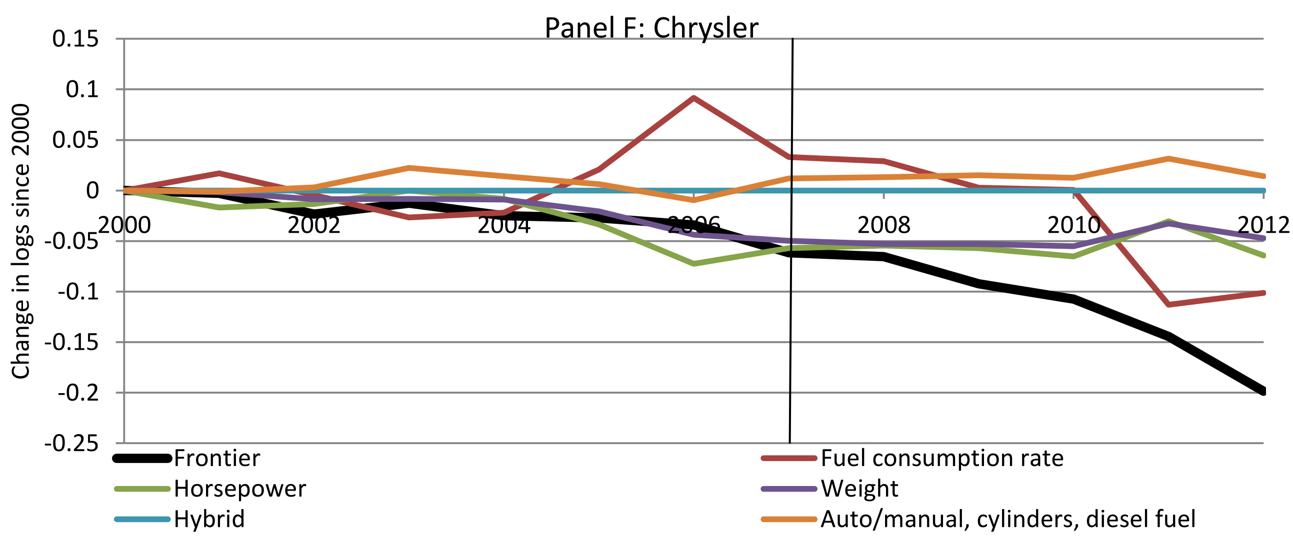
Panel F: Chrysler
Change in logs since 2000
0.15
0.1
0.05
0
-0.05
-0.1
-0.15
-0.2
-0.25
2000
2002
2004
2006
2008
2010
2012
Frontier
Horsepower
Hybrid
Fuel consumption rate
Weight
Auto/manual, cylinders, diesel fuel

Appendix Figure 2. U.S. Technology Adoption by Company, Trucks

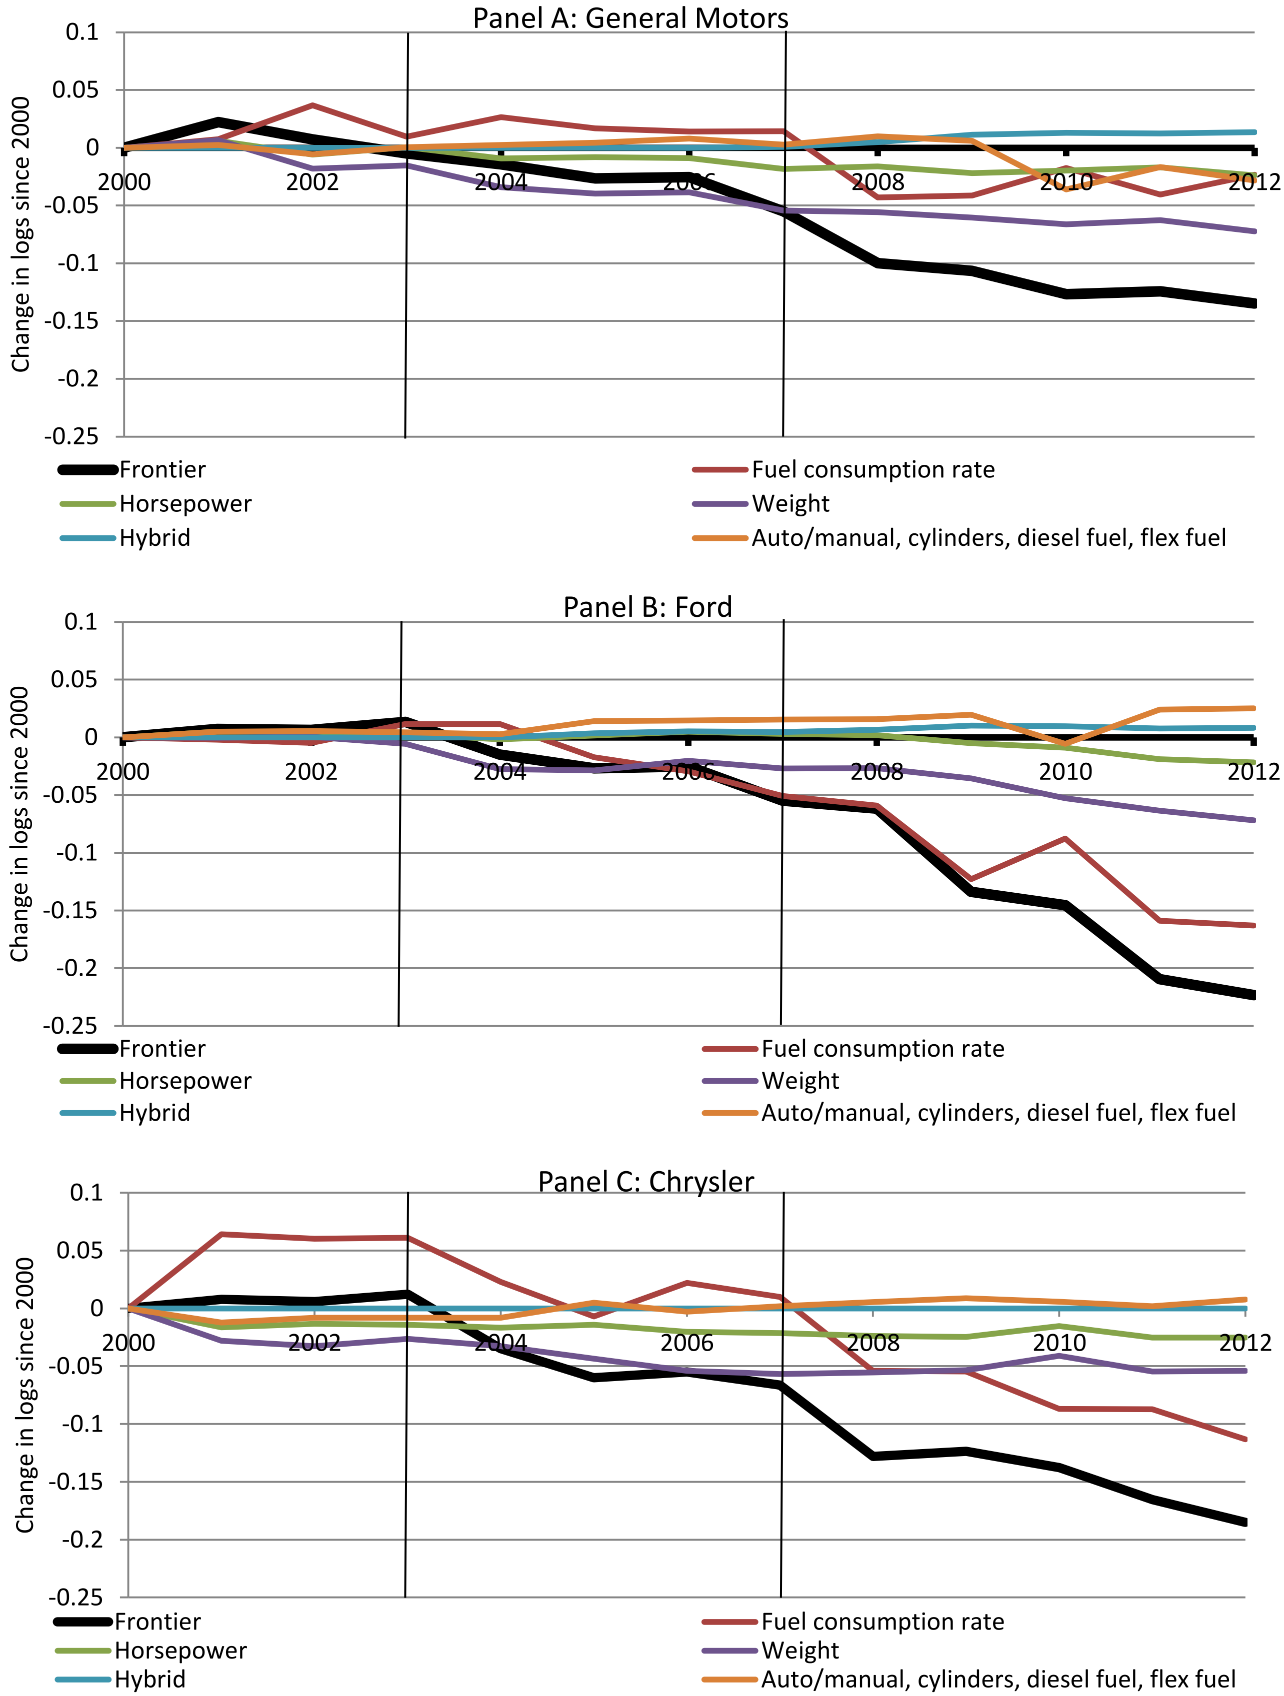

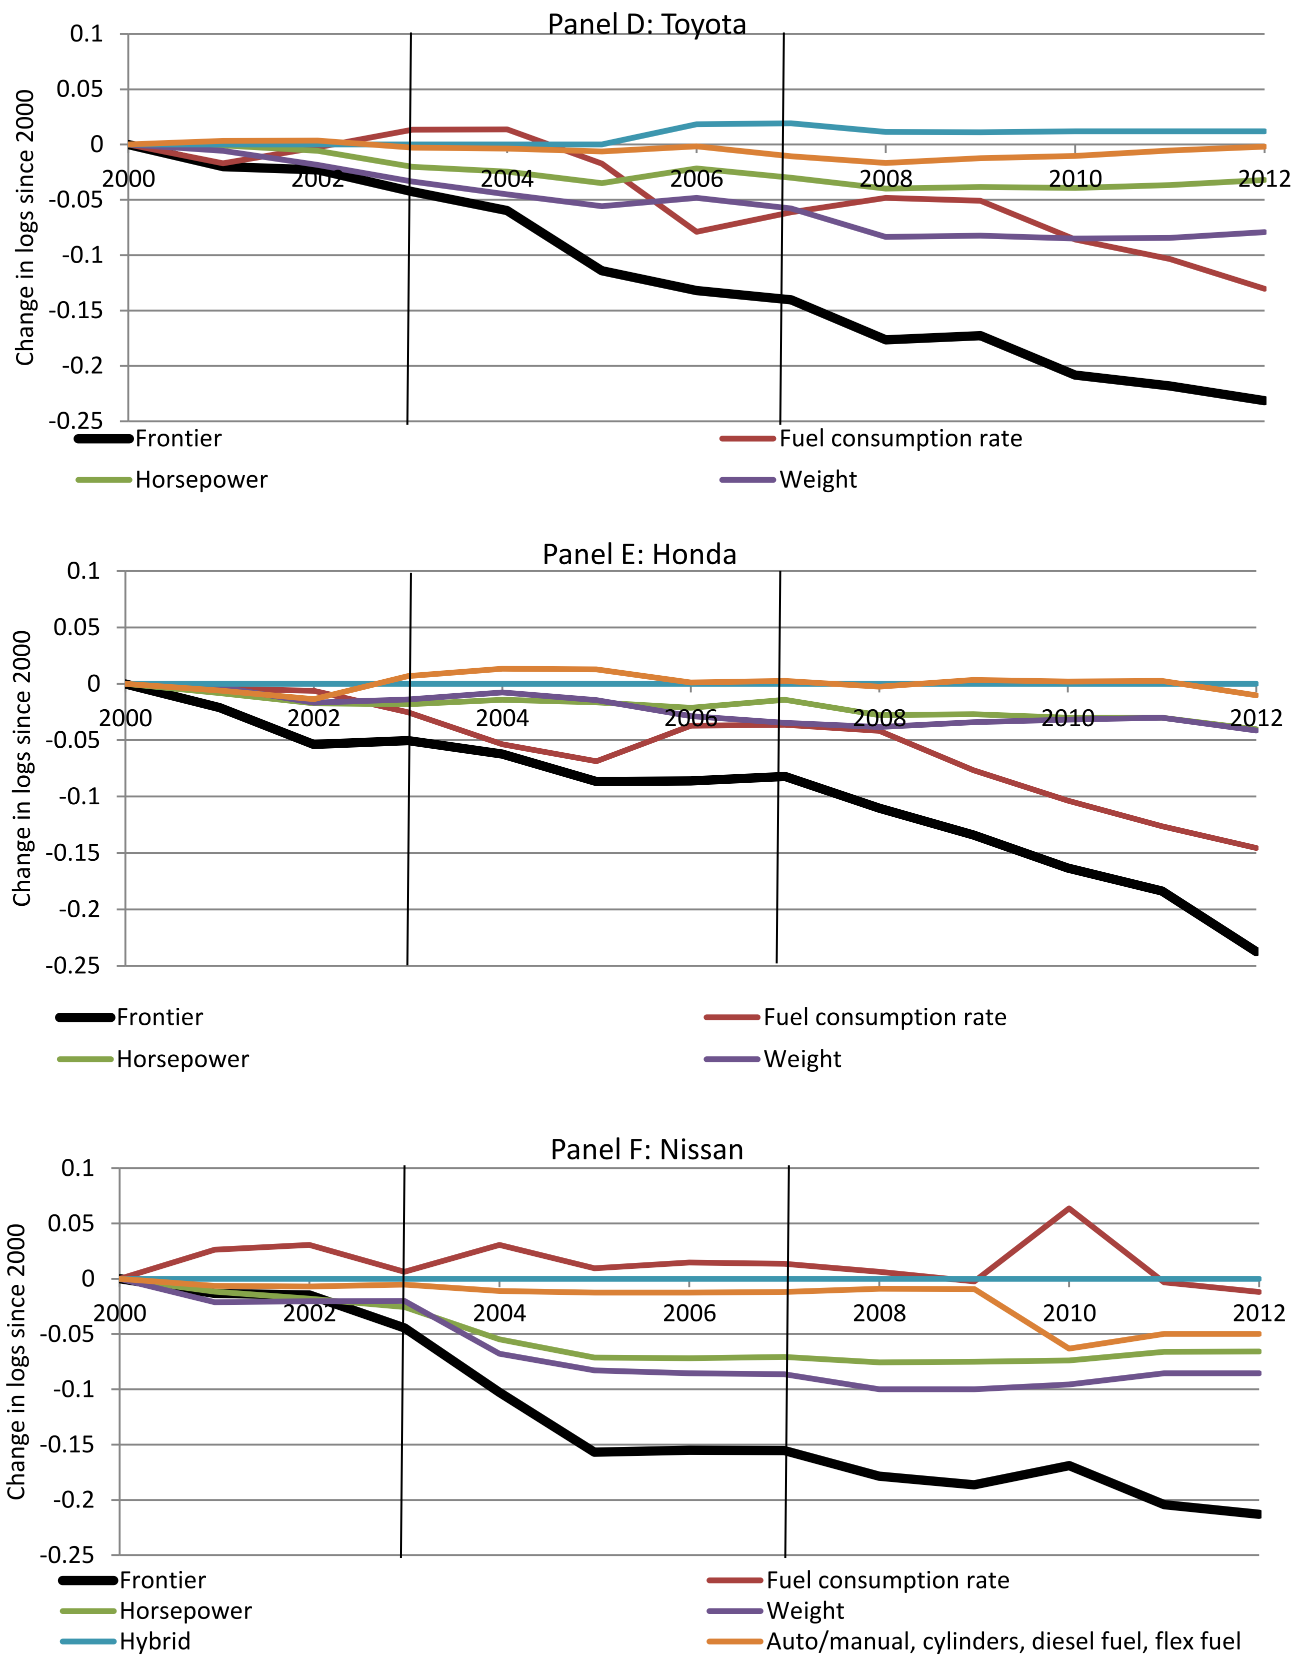
Panel D: Toyota
Change in logs since 2000
0.1
0.05
0
-0.05
-0.1
-0.15
-0.2
-0.25
2000
2002
2004
2006
2008
2010
2012
Frontier
Fuel consumption rate
Horsepower
Weight
Panel E: Honda
Change in logs since 2000
Frontier
Fuel consumption rate
Horsepower
Weight
Panel F: Nissan
Change in logs since 2000
Frontier
Fuel consumption rate
Horsepower
Weight
Hybrid
Auto/manual, cylinders, diesel fuel, flex fuel

Appendix Figure 3. European Technology Adoption by Brand

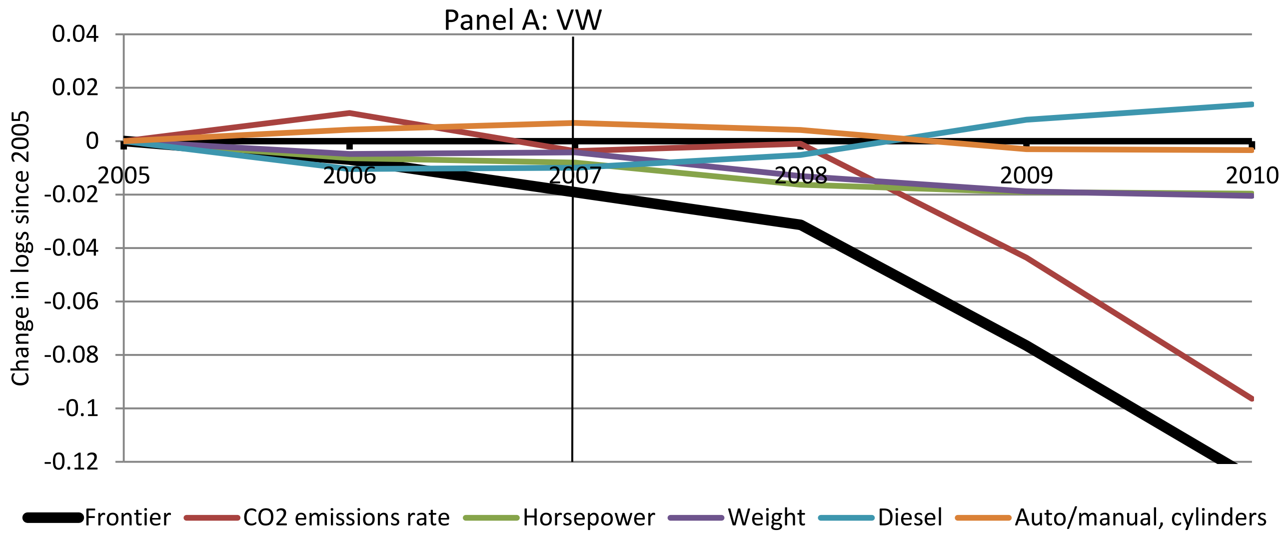

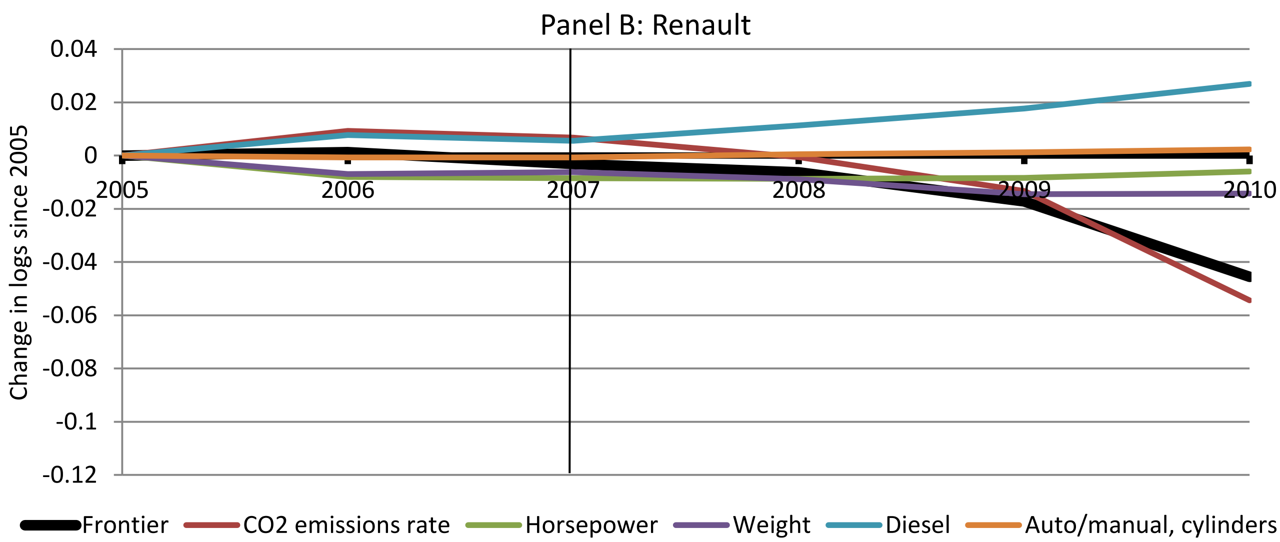

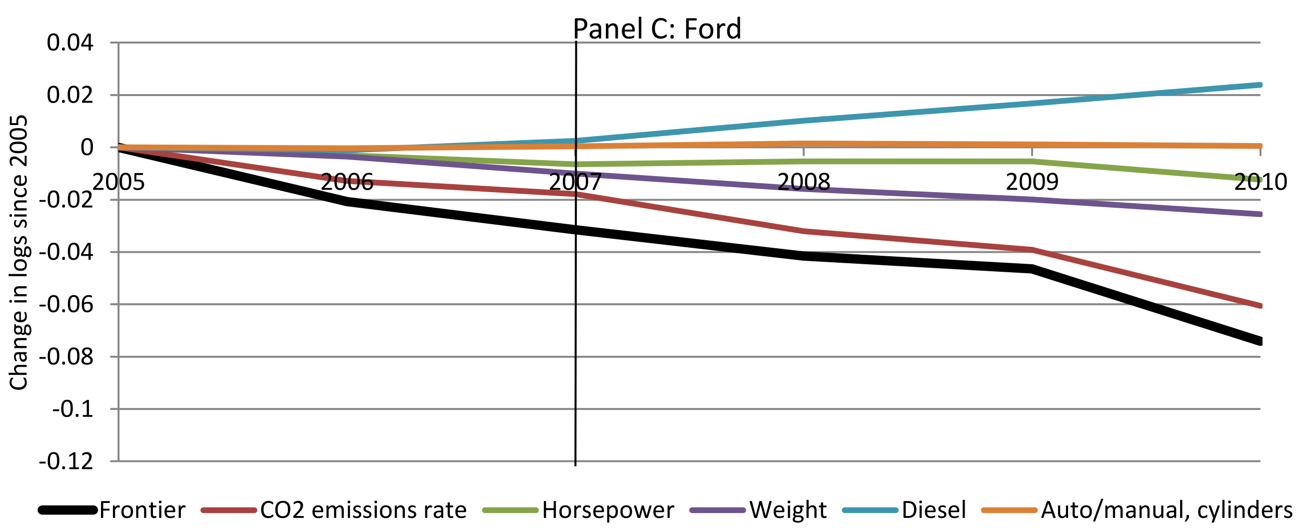

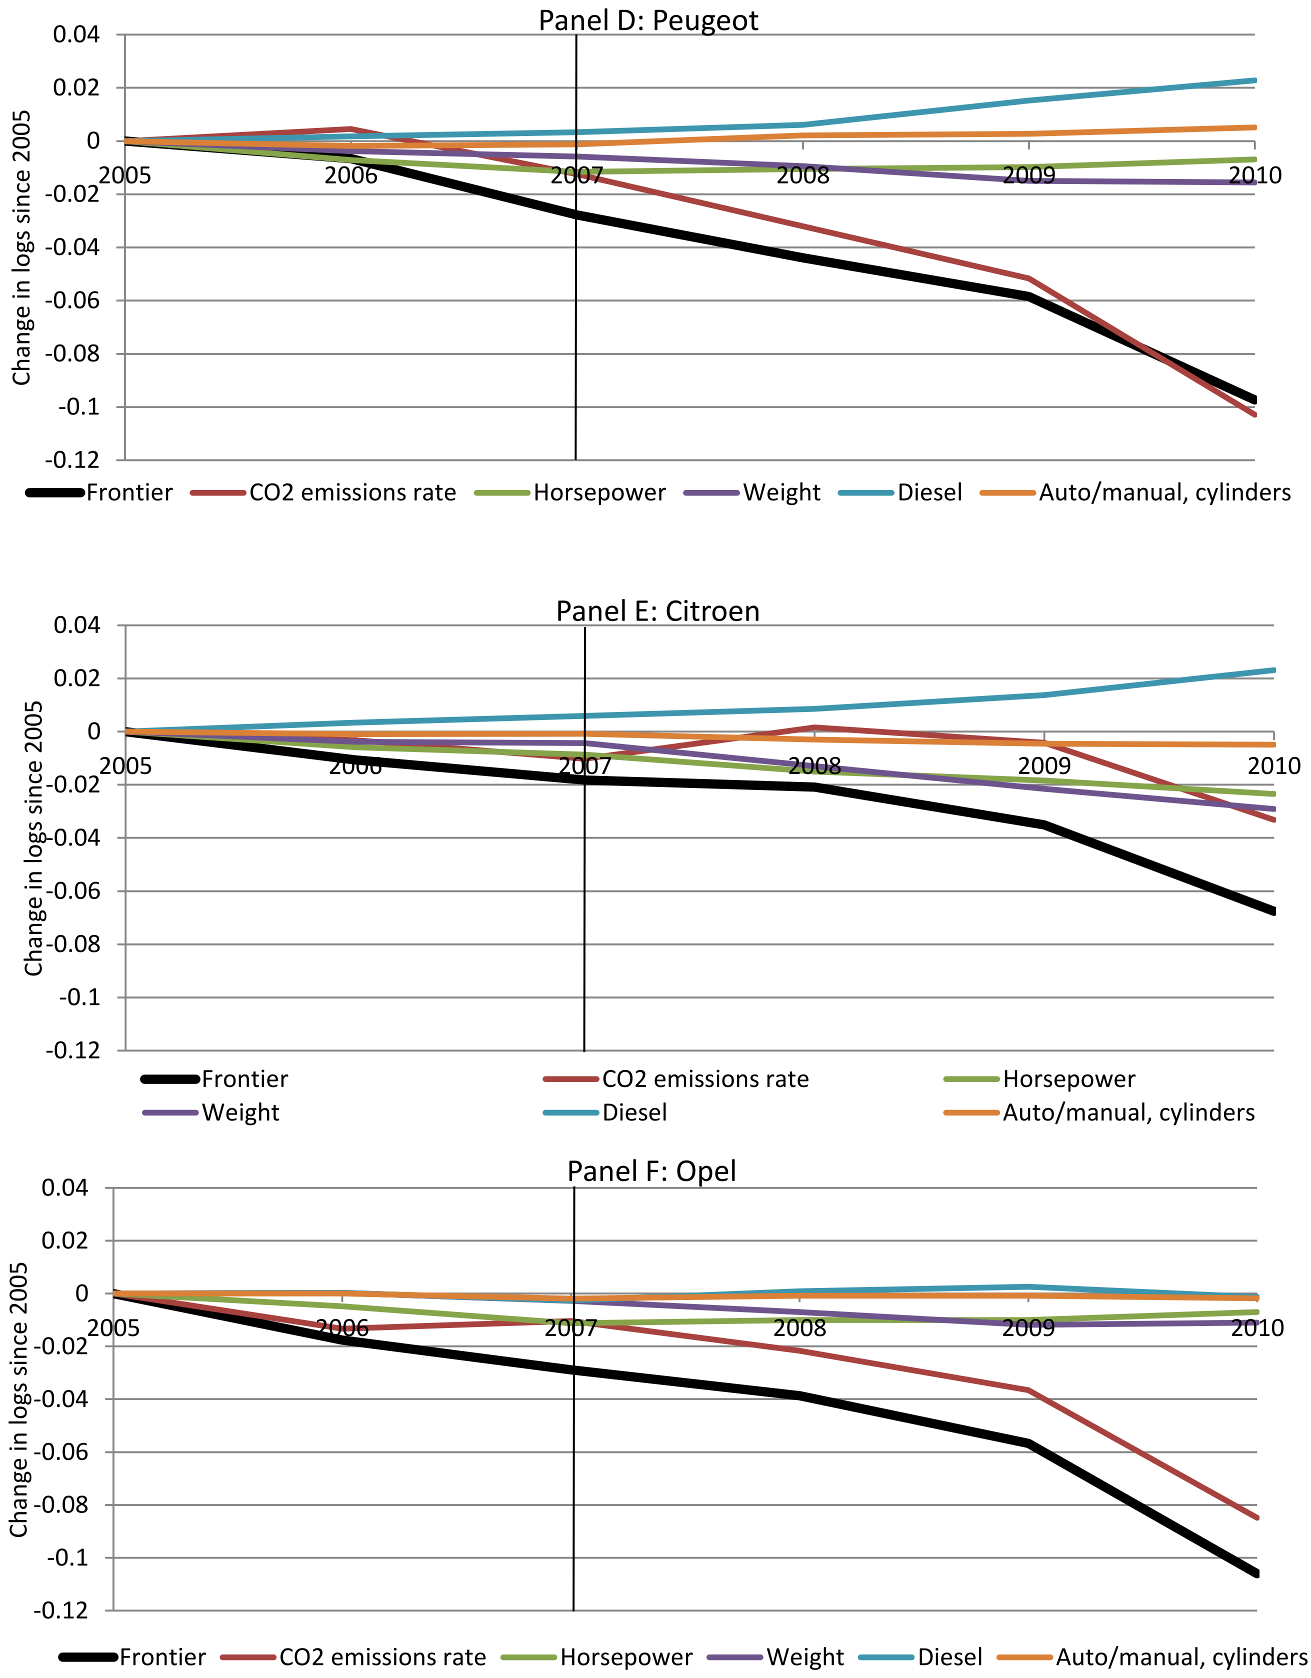

Panel D: Peugeot
Change in logs since 2005
0.04
0.02
0
-0.02
-0.04
-0.06
-0.08
-0.1
-0.12
2005
2006
2007
2008
2009
2010
Frontier
CO2 emissions rate
Horsepower
Weight
Diesel
Auto/manual, cylinders
Panel E: Citroen
Change in logs since 2005
0.04
0.02
0
-0.02
-0.04
-0.06
-0.08
-0.1
-0.12
2005
2006
2007
2008
2009
2010
Frontier
CO2 emissions rate
Horsepower
Weight
Diesel
Auto/manual, cylinders
Panel F: Opel
Change in logs since 2005
0.04
0.02
0
-0.02
-0.04
-0.06
-0.08
-0.1
-0.12
2005
2006
2007
2008
2009
2010
Frontier
CO2 emissions rate
Horsepower
Weight
Diesel
Auto/manual, cylinders

Appendix Table 1. Tradeoffs by Segment: United States

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Sample includes	Small cars	Medium cars	Large/luxury cars	Crossovers	Sport utility vehicles	Vans	Pickups
				Dependent variable: log fuel consumption rate			
Log horsepower or torque	0.245	0.188	0.200	0.165	0.154	0.183	0.122
	(0.028)	(0.028)	(0.016)	(0.022)	(0.026)	(0.051)	(0.021)
Log weight	0.654	0.186	0.275	0.571	0.587	0.358	0.399
	(0.107)	(0.060)	(0.053)	(0.105)	(0.035)	(0.051)	(0.058)
Diesel fuel	-0.375	-0.201	-0.309	-0.287	-0.228		
	(0.017)	(0.028)	(0.022)	(0.018)	(0.019)		
Hybrid	-0.350	-0.293	-0.105	-0.319	-0.286		-0.298
	(0.030)	(0.023)	(0.026)	(0.015)	(0.025)		(0.007)
Flex fuel				0.352	0.290	0.225	0.280
				(0.003)	(0.011)	(0.003)	(0.017)
Manual transmission	0.002	-0.006	0.012	-0.008	-0.003		0.004
	(0.010)	(0.006)	(0.007)	(0.008)	(0.005)		(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

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. Units are the same as in Table 1; torque is measured in Newton-meters.

Appendix Table 2. Tradeoffs by Segment: Europe

	(1)	(2)	(3)	(4)	(5)	(6)
Sample includes	Mini	Small	Lower medium	Medium	Upper medium	Large
	Dependent variable: log CO_2 emissions rate					
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

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 3 except that the sample is restricted to included observations from the market segment indicated at the bottom of the table. Units are the same as in Table 1.

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

	(1)	(2)	(3)	(4)
		Panel A: direction		
Dependent variable	Log (fuel cons rate / horsepower)	Log (fuel cons rate / weight)	Log (fuel cons rate / torque)	Log (fuel cons rate / weight)
Stringency X 2003–2006	0.047 (0.068)	0.032 (0.052)	-0.646 (0.120)	-0.320 (0.070)
Stringency X 2007–2009	0.003 (0.073)	0.052 (0.055)	-0.878 (0.134)	-0.441 (0.078)
Stringency X 2010–2012	-0.083 (0.077)	0.049 (0.066)	-0.714 (0.142)	-0.242 (0.092)
Number of observations	6,856	6,856	11,966	11,966
R^2	0.781	0.615	0.616	0.622
		Panel B: rate		
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 4 except that the independent variables include only the reported variables, model-year fixed effects, and model fixed effects. Units are the same as in Table 1; torque is measured in Newton-meters.

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

	(1)	(2)
	Panel A: direction	
Dependent variable	Log (emissions rate / horsepower)	Log (emissions rate / weight)
Stringency X post 2007	-0.076	-0.046
	(0.007)	(0.006)
Number of observations	275,675	275,675
R^2	0.764	0.585
	Panel B: rate	
Stringency X post 2007	-0.028	-0.022
	(0.003)	(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 5 except that the independent variables include the reported variables, model-year fixed effects, and model trim fixed effects. Units are the same as in Table 1.

Appendix Table 5. Effect of U.S. Standards on Rate of Technology Adoption, Including Company or Market Segment Time Trends

	(1)	(2)	(3)	(4)
		Panel A: rate, with company time trends		
Stringency X 2003–2006	0.074	0.066	-0.218	-0.213
	(0.066)	(0.067)	(0.075)	(0.077)
Stringency X 2007–2009	0.098	0.068	-0.263	-0.220
	(0.118)	(0.117)	(0.114)	(0.117)
Stringency X 2010–2012	0.102	0.091	-0.151	-0.079
	(0.162)	(0.161)	(0.146)	(0.149)
Number of observations	1,749	1,749	1,425	1,425
R^2	0.786	0.784	0.852	0.851
		Panel B: rate, with market segment time trends		
Stringency X 2003–2006	-0.006	-0.017	-0.261	-0.267
	(0.038)	(0.039)	(0.064)	(0.065)
Stringency X 2007–2009	-0.057	-0.086	-0.283	-0.268
	(0.046)	(0.046)	(0.064)	(0.065)
Stringency X 2010–2012	-0.113	-0.121	-0.154	-0.132
	(0.050)	(0.050)	(0.065)	(0.067)
Number of observations	1,749	1,749	1,425	1,425
R^2	0.775	0.763	0.840	0.837
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 Panel B of Table 4 except that Panel A includes the interactions of a linear time trend with a set of company fixed effects, and Panel B includes the interactions of a linear time trend with a set of market segment fixed effects.

Appendix Table 6. Effect of Standards on Rate of Technology Adoption, Allowing For Time Trends in First-Stage Estimation

	(1)	(2)	(3)
Stringency X 2003–2006	0.020	-0.259	
	(0.039)	(0.065)	
Stringency X 2007–2009	-0.022	-0.305	
	(0.046)	(0.066)	
Stringency X 2010–2012	-0.095	-0.166	
	(0.050)	(0.069)	
Stringency X post 2007			-0.026
			(0.004)
Number of observations	1,749	1,425	1,425
R^2	0.945	0.822	0.838
Sample includes	U.S. cars	U.S. light trucks	European cars

Notes : The table reports coefficient estimates, with standard errors in parentheses, clustered by redesign, model, and model-year. Regressions in columns 1 and 2 are the same as in columns 1 and 3 of Panel B of Table 4 except that equation (1) includes interactions of weight and horsepower with linear time trends. Regressions in column 3 is the same as in column 1 in Panel B of Table 5 except that equation (1) includes interactions of

Appendix Table 7. Direction Results by Market Segment: United States

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Sample includes	Small cars	Medium cars	Large/luxury cars	Crossovers	Sport utility vehicles	Vans	Pickups
			Panel A: Dependent variable is log (fuel cons rate / horsepower)				
Stringency X 2003–2006	-0.280	0.327	-0.031	-0.521	-0.826	-0.121	-1.505
	(0.217)	(0.223)	(0.070)	(0.289)	(0.142)	(0.287)	(0.493)
Stringency X 2007–2009	-0.552	0.423	-0.131	-0.648	-0.763	-0.553	-1.862
	(0.222)	(0.216)	(0.087)	(0.287)	(0.184)	(0.262)	(0.442)
Stringency X 2010–2012	-0.811	-0.251	-0.082	-0.627	-0.617	-0.365	-0.772
	(0.242)	(0.277)	(0.085)	(0.284)	(0.187)	(0.288)	(0.584)
Number of observations	1,798	2,188	2,870	2,416	2,826	1,105	5,861
R^2	0.655	0.661	0.698	0.605	0.741	0.560	0.563
			Panel B: Dependent variable is log (fuel cons rate / weight)				
Stringency X 2003–2006	0.578	-0.018	-0.064	-0.618	-0.133	-0.168	-0.512
	(0.300)	(0.179)	(0.050)	(0.188)	(0.095)	(0.161)	(0.318)
Stringency X 2007–2009	-0.131	0.231	-0.072	-0.773	0.130	-0.519	-0.589
	(0.300)	(0.174)	(0.064)	(0.188)	(0.123)	(0.156)	(0.334)
Stringency X 2010–2012	-0.134	0.096	0.005	-0.711	0.429	-0.545	0.446
	(0.311)	(0.231)	(0.067)	(0.197)	(0.137)	(0.189)	(0.457)
Number of observations	1,798	2,188	2,870	2,416	2,826	1,105	5,861
R^2	0.518	0.687	0.642	0.514	0.727	0.766	0.558

Notes : The table reports coefficient estimates, with standard errors in parentheses, clustered by model and model-year. The dependent variable in Panel A is the log of the ratio of the fuel consumption rate to horsepower and the dependent variable in Panel B is the log of the ratio of the fuel consumption rate to weight. Regressions are the same as in Table 4 except that each sample includes the market segment indicated at the top of the table.

Appendix Table 8. Direction Results by Market Segment: Europe

	(1)	(2)	(3)	(4)	(5)	(6)
Sample includes	Mini	Small	Lower medium	Medium	Upper medium	Large
	Panel A: Dependent variable is log (fuel cons rate / horsepower)					
Stringency X post 2007	0.146 (0.060)	0.148 (0.028)	0.047 (0.021)	-0.109 (0.015)	-0.021 (0.012)	-0.057 (0.023)
Number of observations	8,263	47,425	91,430	90,206	35,170	3,882
R^2	0.666	0.526	0.579	0.735	0.833	0.902
	Panel B: Dependent variable is log (fuel cons rate / weight)					
Stringency X post 2007	-0.041 (0.045)	0.022 (0.021)	0.016 (0.020)	-0.058 (0.013)	0.017 (0.010)	0.012 (0.023)
Number of observations	8,263	47,425	91,430	90,206	35,170	3,882
R^2	0.617	0.410	0.460	0.599	0.775	0.811

Notes : The table reports coefficient estimates, with standard errors in parentheses, clustered by model and model-year. The dependent variable in Panel A is the log of the ratio of the fuel consumption rate to horsepower and the dependent variable in Panel B is the log of the ratio of the fuel consumption rate to weight. Regressions are the same as in Table 5 except that each sample includes the market segment indicated at the top of the table.