



MIT Center for
Energy and Environmental
Policy Research

**Using Vehicle Taxes to Reduce
Carbon Dioxide Emissions Rates of
New Passenger Vehicles: Evidence from
France, Germany, and Sweden**

Thomas Klier and Joshua Linn

August 2012

CEEPR WP 2012-011

Using Vehicle Taxes To Reduce Carbon Dioxide Emissions Rates of New Passenger Vehicles: Evidence from France, Germany, and Sweden¹

Thomas Klier

Federal Reserve Bank of Chicago

Joshua Linn

Resources for the Future

July 2012

Abstract

France, Germany, and Sweden link vehicle taxes to the carbon dioxide (CO₂) emissions rates of passenger vehicles. Based on new vehicle registration data from 2005–2010, a vehicle's tax is negatively correlated with its registrations. The effect is somewhat stronger in France than in Germany and Sweden. Taking advantage of the theoretical equivalence between an emissions rate standard and a CO₂-based emissions rate tax, we estimate the effect on manufacturers' profits of reducing emissions rates. For France, a decrease of 5 grams of CO₂ per kilometer reduces profits by 24 euros per vehicle. We find considerable heterogeneity across manufactures and countries.

JEL Codes: L62, Q54

Keywords: feebate, fuel economy standards, emissions rate standards

¹ We thank participants at the summer conference of the Association of Environmental and Resource Economists, the Bergen Economics of Energy and Environment Research Conference, and University of Gothenburg seminars. We are grateful to the Massachusetts Institute of Technology Center for Energy and Environmental Policy Research, the Center for Climate and Electricity Policy at Resources for the Future, and the Swedish Energy Agency for supporting the research. Jeffrey Ferris and Taft Foster provided excellent research assistance.

1. Introduction

To reduce carbon dioxide (CO₂) emissions from passenger vehicles, many countries have recently adopted policies to reduce the emissions rates of new vehicles. Often, these policies set standards for CO₂ emissions rates or for fuel economy; manufacturers must meet the standards or pay fines for noncompliance. For example, the United States is increasing its Corporate Average Fuel Economy (CAFE) standards by about 40 percent from 2012 to 2016, with further increases expected to occur through 2025. The European Union sets greenhouse gas emissions rate standards and has mandated a decrease in emissions rates of about 20 percent from 2009 to 2015; as in the United States, further decreases are expected through 2020. In addition to the increasing stringency of the standards, there has been a recent trend toward using taxes on vehicle purchase and ownership to reduce CO₂ emissions rates. Historically, many European countries subsidized or taxed vehicle purchase or ownership depending on vehicle characteristics correlated with CO₂ emissions—for example, Germany taxed vehicle ownership largely according to engine size. Since the mid 2000s, many of these countries have reformed their tax systems to link taxes directly to the CO₂ emissions rates of the vehicles.

This paper examines recent tax changes in France, Germany, and Sweden and makes two contributions to the existing literature on CO₂ policy for new passenger vehicles. First, many vehicle industry analysts have advocated vehicle taxes over standards because of their cost-effectiveness, simplicity, and greater incentives for new technology. Furthermore, a major objective of the tax reforms has been to reduce emissions rates of new vehicles. However, compared to the extensive literature on fuel economy standards (e.g., Goldberg 1998 and Jacobsen 2010), very little empirical evidence supports the efficacy of taxes at reducing the average CO₂ emissions rates of new passenger vehicles. This paper compares major tax reforms in three countries regarding their effectiveness at reducing emissions rates.

Taxes can vary along two important dimensions: (a) the tax can be assessed at the time of purchase or in the form of an annual circulation (i.e., registration) tax and (b) the tax can be a linear function of the CO₂ emissions rate or can change with the emissions rate in discrete amounts. The different tax systems could

have different effects on vehicle purchases, and therefore on average CO₂ emissions rates. Consumers may respond more to an up-front purchase tax than to the present discounted value of expected circulation taxes, or the linear system may be less salient than the discrete system (Busse et al. 2006). The three countries—France, Germany, and Sweden—provide an excellent opportunity to examine the effects of vehicle taxes because the tax changes were large in magnitude, sudden, and administered differently. France taxes and subsidizes vehicle purchases, and the amount changes discretely with the vehicle's emissions rate. Germany and Sweden impose circulation taxes that increase linearly with the emissions rate. We compare the effect of taxes on registrations and average emissions rates across the three countries.²

The paper's second contribution is to use the empirical results from the first analysis to estimate the reduction in manufacturers' profits from tightening CO₂ emissions rate standards for new vehicles. According to standard economic theory, the cost of reducing CO₂ emissions is minimized when marginal abatement costs are equated across firms and sectors. Automobile manufacturers have asserted that costs to consumers and manufacturers (i.e., the reduced profits) of reducing emissions are much higher than in other sectors such as electric power, whereas regulators and many analysts have claimed that these costs are overstated. Furthermore, some manufacturers, particularly those specializing in larger vehicles, have claimed that they would be particularly adversely affected by tighter standards. If such heterogeneity across manufacturers exists, the cost-effectiveness of reducing emissions would be improved by allowing for credit trading under a standard or by introducing a feebate to equate marginal abatement costs across manufacturers (Anderson et al. 2011).

Recent research has attempted to estimate the effect of tightened standards on manufacturers' profits, estimating both the average effect and also considering whether the effect varies across manufacturers. In the academic literature (e.g., Goldberg 1998; Jacobsen 2010; Whitefoot et al. 2011; Klier and Linn 2012)

² For convenience, throughout the paper we describe the French system as a tax, although it is often referred to as a *bonus-malus* or *feebate* because some vehicles are taxed and others are subsidized.

and among regulatory agencies (e.g., U.S. Environmental Protection Agency 2010), the most common approach is to use a structural model of the vehicle market that captures how manufacturers adjust vehicle prices and emissions rates in response to the standards, and how consumers respond to those changes. Many assumptions about technology costs, manufacturer pricing behavior, and the structure of consumer demand are necessary to implement this approach. In contrast, we argue that we can use the theoretical equivalence between a linear CO₂ emissions rate tax and a CO₂ emissions rate standard, combined with observed market responses to the taxes in France, Germany, and Sweden, to estimate the reduction in manufacturers' profits from achieving a small decrease in emissions rates. Compared to the structural approach, this approach obviates the need for assumptions on technology costs and pricing behavior, and requires much milder assumptions on consumer behavior. Our paper is similar in spirit to Anderson and Sallee (2011), who use flex-fuel incentives in the CAFE program to estimate the cost to manufacturers of a marginal change in the emissions rate standard.

We next provide a more detailed account of the two analyses. First, we estimate the effect of taxes on the registrations of individual vehicles. Prior to 2008, France imposed vehicle purchase taxes largely on the basis of horsepower; but at the beginning of 2008, the purchase taxes changed dramatically. The tax increases in discrete steps from a subsidy of 1,000 euros for vehicles below 100 grams of CO₂ per kilometer (g CO₂/km, which corresponds to 68 miles per gallon [mpg] for a gasoline vehicle) to a tax of 2,600 euros for vehicles above 250 g CO₂/km (27 mpg).³ Germany and Sweden, on the other hand, have introduced annual circulation taxes that increase linearly with the emissions rate. Germany, which had previously taxed vehicles largely based on engine size, changed to the linear tax in July 2009. In October 2006, Sweden changed from a weight-based circulation tax to a linear CO₂-based tax. In both countries, the tax varies significantly across vehicles; for example, in both countries, the tax for a gasoline vehicle with 250 g CO₂/km is more than twice as high as the tax for a vehicle with 150 g CO₂/km (45 mpg).

³ Vehicles below 60 g CO₂/km receive a subsidy of 5,000 euros, but no such vehicles were in the market at the time of the tax reform.

Below, we provide evidence that the reforms caused large changes in the taxes imposed on individual vehicles and that the reforms were largely unexpected, at least no more than a few months in advance.

The empirical analysis separately estimates the short- and long-run effect of the reforms, where the short run is defined as the period of time in which vehicle characteristics are held fixed. Because of regular production cycles in the European markets, the short run corresponds to one calendar year (Klier and Linn 2011). We are interested in the short-run effect of taxes on equilibrium registrations; consequently, we estimate a reduced-form relationship between taxes and registrations rather than estimating a consumer demand function. During the period of interest, 2005–2010, several major demand- and supply-side changes occurred, including the adoption of CO₂ emissions rate standards and shifting market shares across vehicle classes. We take advantage of highly detailed data on vehicle characteristics and, by including vehicle fixed effects and other controls, we control for demand- and supply-side variables that may be correlated with taxes. We find that the tax change in France had a very large effect on vehicle registrations and explains nearly all of the observed reduction in the average emissions rate between 2007 and 2008. The results are similar to those of D’Haultfoeuille et al. (2010), who estimate the short-run effects of the French program.⁴ Compared to the French tax, the effects of the German and Swedish taxes are smaller and are less robust across regression models, particularly for Sweden.

We take advantage of the nonlinear CO₂ tax in France to identify whether firms reduced the emissions rates of individual vehicles in the long run. Firms regularly change emissions rates and other characteristics at the beginning of the year, and because the tax change in 2008 was unexpected, firms probably did not change the emissions rates of individual vehicles at the beginning of 2008. However, because the tax rates did not change between 2008 and 2009, firms may have changed emissions rates of individual vehicles in 2009 or 2010. Because of the discrete tax system, we test the hypothesis that firms reduced emissions rates more for vehicles that were initially above the cutoffs around which the tax

⁴ As discussed in more detail below, D’Haultfoeuille et al. (2010) estimate a demand model in contrast to our reduced-form approach. Furthermore, they do not examine the long-run effects of the policy or estimate the effect of reducing emissions rates on manufacturers’ profits.

changes, compared to vehicles that were initially below the cutoffs. That is, in the spirit of Sallee and Slemrod (forthcoming), we implement a simple difference-in-difference strategy and compare changes in emissions rates above and below the cutoffs after 2008, compared to changes in earlier years. We find very little evidence that the French tax affected the CO₂ emissions rates of individual vehicles.

In the second part of the analysis, we use the empirical results from the first part to estimate the change in manufacturers' profits from reducing average emissions rates by a small amount. A linear CO₂ emissions rate tax is theoretically equivalent to an emissions rate standard in the sense that the two policies can be designed to achieve the same emissions rate reduction at the same cost to vehicle manufacturers and consumers. Many analysts have advocated a tax over an emissions rate standard because it (a) is administratively simpler; (b) provides stronger incentives for new technology unless the standard is increased continually over time; and (c) automatically equalizes the marginal costs of reducing emissions rates across manufacturers compared to a standard, which would have to allow for credit banking and trading among firms to achieve this outcome (Anderson et al. 2011). We take advantage of the equivalence between a linear tax and an emissions rate standard and assume that (a) firms respond to the tax as they would to an emissions rate standard and (b) consumers respond to vehicle prices the same whether the price change occurs because of a tax change or because of emissions rate standards. The first assumption is supported by the theoretical equivalence of the two policies demonstrated in Section 6. D'Haultfoeuille et al. (2010) suggest that the second assumption may not have held precisely for the French tax, but their results suggest that the resulting bias is likely to be small.

We estimate the effect of reducing emissions rates on manufacturers' profits separately for France, Germany, and Sweden; within France, we estimate this effect separately for firms according to their initial emissions rates. For France, we estimate that reducing the average emissions rate by 5 g CO₂/km (about 3 percent) reduces manufacturers' profits by about 24 euros/vehicle, and that this estimate varies across manufacturers from about 10 to 50 euros/vehicle. For Germany and Sweden, the estimates are substantially higher. The heterogeneity across manufacturers and countries suggests that the cost to

vehicle producers and consumers of reducing emissions rates of new vehicles could be reduced by allowing credit trading under an emissions rate standard or by using a feebate. The French estimates are significantly lower than the penalties that will be imposed for noncompliance with the standards. Although the literature includes no comparable estimates for the European markets, our estimates are lower than structural estimates from the United States and are close to those of Anderson and Sallee (2011). The estimates are roughly similar to engineering estimates from Europe (e.g., TNO 2011), but the comparison is difficult because we consider only small changes in emissions rates, and the engineering estimates typically focus on manufacturers' costs rather than profits.

The remainder of the paper proceeds as follows. Section 2 describes the data for vehicle taxes, registrations, and characteristics in each country. Section 3 provides some background to vehicle taxes and emissions rate standards. Section 4 describes the empirical strategy for estimating the effect of the taxes on registrations. Section 5 reports the estimated effects of the taxes on registrations in the short run for each country, and in the long run for France. Section 6 provides the theoretical framework for using the results in Section 5 to estimate the costs to manufacturers of reducing emissions rates, and reports the estimates. Section 7 concludes.

2. Data

2.1 Vehicle Taxes

Vehicle tax data are collected in the annual tax guide of the European Automobile Manufacturers Association. For each country and year (2005–2010), the guide describes vehicle purchase and annual circulation taxes at the national level and, in some cases, regionally. The taxes include the value-added tax, fixed taxes that do not vary across vehicles, and taxes that depend on the characteristics of the vehicle or owner. For example, many countries have historically levied purchase taxes that depend on the vehicle's weight or engine size. Many countries also differentiate between vehicles that are purchased by

the owner and those that are purchased by an employer. In cases in which the country changed the tax structure, the guide provides the month in which the tax change occurred.

2.2 Vehicle Registrations and Characteristics

Vehicle registration data for France, Germany, and Sweden were obtained from R.L. Polk & Co. The data cover the years 2005–2010. We have monthly registrations by vehicle specification, where a *specification* is defined by the vehicle’s model name, trim, number of doors, engine size (cubic centimeters of displacement), power, transmission type (manual or automatic), and fuel type (gasoline or diesel fuel). Thus, the data distinguish versions of the same vehicle *model* that have different engines. For each specification, the data include fuel consumption (liters of fuel per 100 km) and the CO₂ emissions rate. The purchase tax and annual registration tax can be calculated for each specification and month using the vehicle characteristics in the data.

Although vehicles that use fuels other than gasoline or diesel receive substantial tax incentives, the analysis for France and Germany includes only gasoline and diesel fuel vehicles, and for Sweden it includes gasoline, diesel fuel, and flex-fuel vehicles (vehicles capable of using gasoline or ethanol). Flex-fuel vehicles have become increasingly popular in Sweden, at least partly because of their preferential tax treatment (Huse 2012), but their market share is very small in France and Germany. We exclude hybrids, plug-ins, and natural gas vehicles because they account for a very small share of the market in all three countries.

We have obtained monthly prices of gasoline and diesel fuel for 2005–2010. The prices, which include all taxes, are converted to 2005 euros for France and Germany and to 2005 kronor for Sweden using country-specific consumer price indexes. The fuel prices and price indexes are from Eurostat.

Figure 1 provides some summary information about the three markets studied in this paper. Panel A shows the log of quarterly registrations and demonstrates regular seasonal patterns in registrations. France and Germany are the countries with the two largest passenger vehicle markets in Europe; annual new

vehicle registrations average about 2 million in France and about 3 million in Germany. Sweden's market is much smaller, with about 250,000 new vehicle registrations per year.

Panels B and C of Figure 1 show other important dimensions of the passenger vehicle markets. Average CO₂ emissions rates decreased steadily in all three countries. Throughout the period, Germany had a higher average emissions rate than France, and the emissions rate in both countries followed roughly parallel trends. In Sweden, the average emissions rate was initially much higher than that of the other two countries, but it decreased more rapidly and was close to that of Germany by 2010.

In all three countries, diesel fuel vehicles typically have lower CO₂ emissions rates than closely related gasoline versions—roughly 20 percent lower, on average. Panel C shows the market share of diesel fuel vehicles. The market share is much higher in France than in the other countries throughout the period. The diesel market share in Sweden grew dramatically from about 15 percent in 2005 to almost 60 percent in 2010. The diesel market share in Germany was fairly stable during this period, although it decreased significantly in 2009 and then recovered.

Table 1 provides further insight about the factors underlying the differences in average emissions rates over time and across countries. The Polk data include a market segment with each vehicle specification, and the table shows the market shares for each segment and country (along with the top-selling model for each country and segment). The right side of the table shows the average emissions rates by country, year, and segment. Across the countries, the market shares and average emissions rates differ substantially; both factors contribute to the cross-country variation in emissions rates. However, it turns out that, because the powertrain mixes of a model vary across countries, emissions rates of particular models vary noticeably across countries and explain a large share of the overall cross-country differences. Focusing on models sold in all three countries (which account for roughly 95 percent of total registrations), differences in average emissions rates of each model explain nearly two-thirds of the overall difference in average emissions rates between France and Germany in 2005; differences in market shares of the models explain

the remainder. Comparing Sweden and Germany in 2005, differences in emissions rates of each model explain about half of the overall difference.

3. Overview of CO₂ Emissions Rate Standards, Taxes, and Retirement Programs

3.1 CO₂ Emissions Rate Standards

The regulation of CO₂ emissions rates from new vehicles changed during the period of analysis. During 1998 and 1999, the European Commission reached an agreement with the major manufacturers to voluntarily reduce the average emissions rates of new passenger vehicles. The agreement gave manufacturers 10 years to reduce the average emissions rate to 140 g CO₂/km. The manufacturers did not have to pay fines if they failed to meet the standards. Instead, the agreement provided that if the manufacturers did not meet the standards, mandatory new standards would be introduced that would be backed by fines. By 2005, it became clear that the manufacturers would not meet the standards, and in 2007 the European Commission created a new program in which the manufacturers would have to reduce emissions rates to 130 g CO₂/km by 2015. Each manufacturer must comply with the standard, although manufacturers can form cooperative agreements in which they are evaluated jointly. If a manufacturer exceeds the standard by more than 3 g CO₂/km, it has to pay a fine of 95 euros/g CO₂/km for each vehicle it sells.⁵ Furthermore, the standard in 2020 is expected to be 95 g CO₂/km. In short, these standards, first the voluntary and then the mandatory, probably contributed to the overall downward trends in emissions rates observed in Figure 1.

⁵ Manufacturers within 3 g CO₂/km of the standard do not have to pay a fine. However, the fine for a manufacturer more than 3 g CO₂/km above the standard includes a small penalty for the first 3 g CO₂/km, plus a larger penalty assessed in proportion to the amount by which the manufacturer exceeds the standard.

3.2 Vehicle Purchase and Circulation Taxes

Traditionally, European countries have taxed vehicle purchase or ownership much more extensively than most other countries, such as the United States. Often, the European taxes were correlated with fuel consumption or CO₂ emissions rates, but did not depend directly on these characteristics and instead depended on characteristics such as weight, horsepower, or engine size. Beginning around 2006, many countries shifted their tax systems to more directly target CO₂ emissions rates. The reforms were widely viewed as being complementary to the CO₂ emissions rate standards and were intended to encourage consumers to purchase vehicles that emit less CO₂.

Prior to 2008, France taxed vehicle purchases on the basis of horsepower and CO₂ emissions rates. Panel A in Figure 2 plots the tax against the emissions rate for all vehicle specifications sold in 2007, where each specification is a unique observation. The tax is positively correlated with the emissions rate.⁶

In January 2008, France introduced a program that offered subsidies to purchases of low-emissions vehicles and imposed taxes on purchases of high-emissions vehicles. These taxes and subsidies were added to the preexisting taxes. Under the new program, the subsidy and tax amounts changed discretely with the emissions rate. For example, purchasers of vehicles with emissions rates between 120 and 130 g CO₂/km received a subsidy of 200 euros, and purchasers of vehicles with emissions rates between 100 and 120 g CO₂/km received a subsidy of 700 euros. The taxes and subsidies were designed so that the program would be budget-neutral in expectation. However, in 2008, the program incurred a debt of 225 million euros, which was widely attributed to an unexpectedly large consumer response to the program (D'Haultfoeuille et al., 2010). To try to reduce the deficit, the French government made slight changes to the taxes and subsidies in 2010.

⁶ The tax rates vary across regions. Because this paper uses national registration data, we use the average tax rates across regions, which are computed as the midpoints of the minimum and maximum tax rates reported in the tax guides.

Panel B of Figure 2 shows a scatter plot of the 2008 tax or subsidy against the emissions rate.⁷ Compared to Panel A, the taxes in 2008 were more clearly associated with the emissions rate. However, although the CO₂ tax changes in discrete steps, after 2007, a portion of the total purchase tax continued to depend on the vehicle's power; consequently, the total tax may vary across vehicles that have the same emissions rate.

Panel C of Figure 2 summarizes the effect of the tax change on the total purchase tax of individual vehicle specifications. For each specification sold in 2007, the purchase tax is computed using the 2007 tax rates and separately using the 2008 tax rates. The figure plots the estimated density functions of the two distributions of taxes. The plot demonstrates the large change in the distributions of tax rates across the two years.

Whereas France has focused on vehicle purchase taxes, Germany and Sweden have not taxed purchases extensively, but instead impose annual circulation taxes. In July 2009, Germany changed from a system based on engine size to a tax that increases linearly with the CO₂ emissions rate. For all vehicles purchased before July 2009, circulation taxes increased linearly with engine capacity, and at a higher rate for diesel fuel vehicles compared to gasoline vehicles. For vehicles purchased in July 2009 or later, registration taxes were the sum of a linear function of engine capacity and a linear function of the CO₂ emissions rate.⁸

Analogously to Figure 2, Figure 3 summarizes the tax changes for Germany. Panels A and B plot the circulation tax against the CO₂ emissions rates of vehicle specifications in the first half of 2009 (2009H1) and the second half of 2009 (2009H2). In both cases, taxes and emissions rates are positively correlated, but the correlation is much stronger for the second half of 2009 because of the linear emissions rate tax.

Panel C shows the distributions of taxes using the 2009H1 and 2009H2 tax rates and including all

⁷ France imposes a linear CO₂ emissions rate tax on the ownership of company cars. Because the tax did not change during the period of analysis, 2007 and 2008, these taxes are not discussed further.

⁸ Depending on the characteristics of the vehicle, Germany offered a one- or two-year exemption to the annual circulation tax for vehicles purchased in late 2008 or early 2009. The exemption probably had a small effect on the present discounted value of future circulation taxes, and therefore on new vehicle registrations.

specifications that have positive registrations in 2009H1. As for France, the changes in taxes are quite dramatic. However, it is important to note that vehicles registered in the first half of 2009 are subject to the new taxes starting in 2013. Furthermore, vehicle owners of vehicles purchased prior to July 2009 may opt into the new tax system if it is advantageous to do so, and as Panel C shows, for most vehicles the new circulation tax is less than the old tax. Therefore, the new tax may apply to some vehicles purchased before July 2009, but consumers only knew this after the announcement of the tax reform; therefore this feature probably had a small effect on new vehicle registrations prior to July 2009.⁹

In the fourth quarter of 2006, Sweden introduced a circulation tax that is a linear function of the vehicle's emissions rate. Figure 4 summarizes the information for Sweden analogous to that in the previous figures for France and Germany. Panel C shows large changes in tax rates between the third and fourth quarters of 2006. The tax rate for diesel fuel vehicles is higher than the tax rate for gasoline vehicles; the tax rate for alternative-fuel vehicles is lower than that for gasoline vehicles (only a few specifications used alternative fuels in the market in 2006).

A second tax change occurred in April 2007 when Sweden introduced a Green Car Rebate. The rebate was a flat subsidy (10,000 kronor) for vehicles that met specific criteria—emitting less than 120 g CO₂/km for diesel fuel and gasoline vehicles.¹⁰ The rebate was claimed for a large number of vehicles, as the program cost the government close to 400 million kronor in 2007 and 2008.

The vehicle tax changes were largely unanticipated, particularly in France and Sweden. Based on counts of articles in major daily newspapers, media coverage of the French tax system increased somewhat in the last quarter of 2007. But coverage was five times higher in the third quarter of 2008 than in the last quarter of 2007. In Germany, some discussion centered on a CO₂-based circulation tax as early as 2004, and several attempts to introduce such a tax over the following years failed. The tax change, introduced as

⁹ The tax change for existing vehicles may cause some households to sell their vehicles and purchase new ones. We discuss implications for the empirical analysis in Section 4.

¹⁰ The emissions rate cutoff is equivalent to about 56 mpg for a gasoline vehicle. The rebate also applied to flex-fuel vehicles, which had a much less stringent mpg cutoff. Huse (2012) argues that the distinction between flex-fuel and other vehicles led to a large increase in the market share of flex-fuel vehicles.

part of a large economic stimulus package, generated a considerable amount of media coverage in the second quarter of 2008 and the first quarter of 2009. Based on counts of articles in major daily newspapers, media coverage in the second quarter of 2008 was twice the overall average for 2008, and coverage in the first quarter of 2009 was even higher—three times more than the coverage in the second quarter of 2008. By comparison, coverage after the tax change was minimal. Finally, in Sweden, the tax change generated some media coverage and public discussion in mid-2006, but not earlier.

3.3 Vehicle Retirement Programs

France and Germany introduced vehicle retirement programs in 2009. To qualify for the French program, the trade-in had to be at least 10 years old. If the new vehicle had an emissions rate below 160 g CO₂/km, the consumer received a subsidy of 1,000 euros. The amount of the subsidy and emissions rate cutoff decreased in 2010. Importantly for the empirical analysis, France introduced its program about a year after the changes in purchase taxes.

The German program, which started at the beginning of 2009, offered 2,500 euros for anyone trading in a vehicle at least nine years old and purchasing a vehicle no more than one year old. The subsidy did not depend on the CO₂ emissions rate of the new vehicle, but the vehicle had to attain emissions rate standards for other pollutants. The initial program budget was 1.5 billion euros but the budget was subsequently expanded to 2.5 billion euros after the program proved unexpectedly popular.

4. Empirical Strategy for Estimating the Short-Run Effect of Taxes on Registrations

Our empirical objective is to estimate the effect of purchase or circulation taxes on the registrations of individual vehicles. Despite the recent tax reforms in Europe, very little empirical work has analyzed their effectiveness at reducing emissions rates. Consequently, we implement a reduced-form analysis of the effect of taxes on equilibrium new vehicle registrations. This approach allows us to control flexibly for

contemporaneous demand and supply shocks, although it does not enable a welfare analysis of the reforms.

We focus on the short run, which we define as the period of time over which the set of vehicles in the market and the characteristics of those vehicles is fixed. Because firms tend to redesign vehicles and introduce new versions at the start of the calendar year, the short run corresponds to one year; Klier and Linn (2011) document these patterns.

We estimate a simple reduced-form regression that links new registrations of vehicle j in quarter t , q_{jt} , to the vehicle tax, T_{jt}

$$\ln q_{jt} = \alpha T_{jt} + \varphi_j + \tau_t + \nu_{mt} + FC_{jt} \beta + \varepsilon_{jt} \quad (1)$$

where φ_j is a vehicle fixed effect, τ_t includes a set of year-quarter interactions, ν_{mt} includes a set of model-year-quarter fixed effects, FC_{jt} is the fuel costs of the vehicle, and α and β are coefficients to be estimated. Observations are defined by year, quarter, and vehicle, which is defined hereafter as a unique model and CO₂ emissions rate; the monthly specification data are aggregated to vehicle, year, and quarter.¹¹ The equation is estimated separately for France, Germany, and Sweden. For France, the tax is the purchase tax, and for Germany and Sweden, the tax is the annual circulation tax. In other words, we assume that circulation taxes remain constant over the life of the vehicle, in which case, the current tax is proportional to the present discounted value of future taxes. This assumption is consistent with the fact that each country introduced major tax reforms only once during the years 2000–2010.

To reduce the likelihood that other policies affect the results, the estimation samples are restricted to the time periods around the tax reforms. The sample for France includes 2007 and 2008, the sample for Germany includes 2009, and the sample for Sweden includes 2005, 2006, and the first quarter of 2007.

¹¹ This aggregation preserves nearly all of the tax variation in the data and reduces the likelihood that measurement error at the specification-month level affects the results. Table 3 reports estimates using the disaggregated data.

The French sample ends prior to the introduction of the retirement program, and the Swedish sample ends before the introduction of the Green Car Rebate, which reduces the possibility that the policies affect the estimates of α . Unfortunately, the German retirement program was introduced just prior to the tax change, which makes it more challenging to control for the program's effect. However, we find that the results are very similar if we take long differences—that is, we compare registrations in 2008 and 2010 because taxes changed significantly between these years, but the retirement program was not in effect during either year.¹²

The coefficient α is the effect of the tax on vehicle registrations. We first discuss the identification of the coefficient and then its interpretation.

Equation (1) includes vehicle fixed effects, time dummies, and model-quarter fixed effects, all of which help identify α and control for demand and supply shocks. Because the vehicle fixed effects, φ_j , are defined at the level of the model by CO₂ emissions rate, the tax coefficient is identified by changes in the tax rates combined with cross-sectional variation in emissions rates. The fixed effects control for all characteristics of the vehicle that do not vary over time.¹³

We control for model-level demand and supply shocks by including model-year-quarter interactions. Therefore, the tax coefficient is identified by within-model variation in taxes that arise because of within-model cross-sectional variation in CO₂ emissions rates interacted with the change in the tax rate. In other words, a given change in the tax rate has a larger effect on T_{jt} for a version of a particular model with a high CO₂ emissions rate compared to a version with a low emissions rate. The inclusion of the model-

¹² This also addresses the potential concern that because the tax reform applied to some vehicles purchased in the first half of 2009, the reform may have affected new vehicle registrations prior to 2009.

¹³ In practice, characteristics such as horsepower may change over time for a given vehicle. These changes tend to be very small in magnitude and therefore should not greatly affect the results. Below we report estimates of equation (1) at the specification level and including specification fixed effects, which controls for such changes.

year-quarter interactions also helps control for the effects of other policies, such as the CO₂ emissions standards and the vehicle retirement programs discussed above.¹⁴

We control for the vehicle's fuel costs, which may be correlated with its taxes. Because the postreform vehicle taxes depend on CO₂ emissions rates, and because CO₂ emissions rates are proportional to fuel consumption rates (which are inversely related to fuel economy), taxes are likely to be correlated with fuel costs. Under the assumption that fuel prices follow a random walk (which is consistent with the data), the lifetime expected fuel costs of the vehicle are proportional to the vehicle's fuel consumption multiplied by the price of fuel. Therefore, equation (1) includes the interaction of the fuel price with the vehicle's fuel consumption, along with the main effects of the variables. Theory predicts that the coefficient should be negative, as higher fuel costs reduce the consumer's disposable income, although Klier and Linn (2011) find the magnitude of the coefficient to be small using similar data.

In summary, we identify α using within-model and within-specification changes in the tax caused by changes in the tax rates over time. Three features of equation (1) address potential concerns about omitted variables bias—that is, the possibility that changes in vehicle characteristics or consumer demand for the characteristics are correlated with changes in the vehicle taxes. First, the sample is restricted to observations close in time to the tax reforms. Second, the vehicle fixed effects control for time-invariant vehicle characteristics and demand for the vehicle. The fixed effects also control for new technology introduced in response to the CO₂ standards discussed in Section 3.1. Third, model-year-quarter intercepts control for demand and supply shocks at the model level.

Alternatively, we could estimate the effect of taxes on registrations using a structural model of consumer demand. There are two main advantages to using equation (1). First, it is not necessary to make assumptions about consumer substitution patterns. Second, we do not need to control for vehicle prices,

¹⁴ The model-year-quarter intercepts absorb a substantial amount of tax variation. A regression of the first-differenced vehicle tax on model-year-quarter intercepts and time fixed effects has a mean squared error roughly half of that from a regression of the first-differenced vehicle tax on time fixed effects. However, as we show below, the remaining variation is sufficient to identify the coefficient on the tax variable.

and it is difficult to control for vehicle prices in a setting such as this, in which vehicle characteristics are changing and are unlikely to be exogenous (Klier and Linn 2012); that is, it is difficult to find plausible instruments for vehicle prices.

On the other hand, the main disadvantage of equation (1) is that we need to be cautious in interpreting the results, particularly in using the estimates to simulate equilibria under counterfactual tax scenarios, as we will do in the following two sections. We minimize this downside in three ways. First, we focus on scenarios that are similar to the tax rates observed in the data. Second, we demonstrate the robustness of the results to structural models, such as a nested logit model, that have been used in the recent literature on European vehicle markets (e.g., Zachariadis 2011). Third, we show that the results are similar if we allow the tax coefficient to vary across vehicles or control for competitive effects of tax changes for close substitutes.

Turning to the interpretation of α , the coefficient represents the reduced-form effect of taxes on vehicle registrations. We can illustrate this interpretation using a simple heuristic example of a tax increase for a particular vehicle. The tax increase causes an inward shift of the vehicle demand curve along the short-run supply curve, and the equilibrium price of the vehicle is likely to decrease. The coefficient α represents the net effect of the tax and price changes on equilibrium registrations, which is expected to be negative. Note that, because of the model-year-quarter fixed effects, a negative coefficient does not imply that increasing emissions tax rates for all vehicles causes registrations of all vehicles to decrease. It simply means that, for a particular model, the registrations of vehicles with high CO₂ emissions rates decrease compared to registrations of vehicles with low CO₂ emissions rates. A negative estimate of α allows for the possibility that an increase in the tax rate increases the registrations of some vehicles. Because equation (1) includes vehicle fixed effects, we interpret this as the short-run effect of a tax change on emissions rates. The next section reports these results and also presents a long-run analysis.

5 Effects of Vehicle Taxes on Registrations

5.1 Baseline Estimates for the Short Run

Table 2 reports the estimated short-run effects of vehicle taxes, where each column shows a separate country. Panel A reports the estimated coefficient on vehicle taxes in equation (1). For each country, variables are in first differences to account for the persistence of quarterly vehicle registrations and taxes. All regressions include a set of model-year-quarter interactions to control for omitted demand and supply shocks at the vehicle model level. Standard errors are clustered by market segment-quarter to allow for correlation in the error term within segments. The samples contain observations immediately before and after the tax changes in each country. Column 1, for France, includes the years 2007–2008; Germany, in column 2, includes 2009; and Sweden, in column 3, includes 2005, 2006, and the first quarter of 2007. We refer to these as the baseline regressions.

The estimated coefficient for France is -0.561 with a standard error of 0.113 , which is significant at the 1 percent level. The estimates for Germany and Sweden are also statistically significant at the 1 percent level. Because France uses a purchase tax and Germany and Sweden use annual circulation taxes, and because the Swedish tax is denominated in kronor rather than euros, the point estimates in Panel A are not directly comparable across countries. Therefore, Panel B reports the estimated elasticity of vehicle registrations to vehicle taxes, where the elasticity is evaluated at the corresponding sample means. The elasticity for France is the largest, followed by Germany and Sweden.

We provide context for the magnitude of the estimates in two ways. First, the point estimate suggests that a tax increase of 1,000 euros (measured in 2005 euros) reduces log registrations by 0.56 in France. Assuming a 10 percent discount rate and a 12-year vehicle lifetime, the corresponding decreases in log registrations for Germany and Sweden are 0.19 and 0.07. By comparison, Klier and Linn (2011) estimate the effect of fuel costs on registrations in Western Europe from 2002 to 2007; the results suggest that a

1,000-euro increase in the present discounted value of fuel costs reduces log registrations by 0.09.¹⁵ Thus, particularly for France and Germany, the tax changes had a larger effect on registrations than a commensurate change in fuel prices during roughly the same time period.¹⁶

Second, we simulate the effect of the actual tax changes on CO₂ emissions rates. Starting from the observed registrations prior to the tax change in each country, we simulate the change in registrations from changing taxes to the postreform levels. Before discussing the results, we provide a note of caution. Equation (1) describes the reduced-form relationship between taxes and vehicle registrations. We do not explicitly model the substitution that would arise from a particular tax change; instead, we estimate the average effect of taxes on registrations. A hypothetical tax change could cause substitution patterns that differ from those observed in the data. That said, the results are very similar if we explicitly model substitution, e.g., using a nested logit model, which is reported in the next section.

Panel C of Table 2 reports the results of the simulations. The first row shows the observed registration-weighted average CO₂ emissions rate prior to the tax changes—2007 in France, the first half of 2009 in Germany, and the first three quarters of 2006 in Sweden. The second row shows the change in emissions rates from changing the taxes to the postreform levels. Standard errors, in parentheses, are estimated using the delta method, in which we assume that emissions rates and observed registrations are measured without error. The tax reduces the average emissions rate by 7.95 g CO₂/km in France, 1.67 g CO₂/km in Germany, and 0.57 g CO₂/km in Sweden. All estimates are statistically significant at the 1 percent level. For comparison, in France, the average emissions rates actually declined by 8.70 g CO₂/km between 2007 and 2008. In Germany, emissions rates decreased by 1.55 g CO₂/km between the first and second half of 2009, and in Sweden, emissions rates decreased by 1.54 g CO₂/km between 2005 and 2007. Thus, in all

¹⁵ The calculation assumes that the average vehicle is driven 15,000 km per year for 12 years and that consumers discount future fuel costs at 10 percent.

¹⁶ It is tempting to compare the tax elasticities with own-price elasticities estimated in the literature, but the elasticity reported in this paper includes the effect of the tax on vehicle prices. We would expect taxes to positively affect equilibrium prices, in which case, the elasticity estimated in Table 2 would be smaller in magnitude than the elasticity of registrations to taxes, holding vehicle prices fixed.

three countries, the tax changes explain large fractions of the actual changes in average emissions rates during the periods surrounding the tax changes, although less so for Sweden.

5.2 Addressing Functional Form Assumptions and Potential Omitted Variables

We first assess whether the results are robust to alternative functional form assumptions. We then discuss alternative approaches to addressing omitted variables. Overall, we find that the empirical estimates are quite robust for France, slightly less so for Germany, and less so for Sweden.

Table 3 shows several regression models that use different functional form assumptions or different levels of aggregation, with each country shown in a separate panel. The first two rows in each panel show the estimated coefficient on the tax variable with standard errors in parentheses, and the bottom two rows show the elasticity of registrations to taxes with standard errors in parentheses.

The first column shows the results from estimating equation (1) in levels rather than first differences. As noted above, we prefer the first-differenced regression because the residuals from the levels regressions suggest substantial serial correlation, which is not the case with the first-differenced regression.

Nevertheless, column 1 shows that the estimates are statistically significant if we estimate equation (1) in levels, although for France and Germany the point estimate is smaller than in Table 2, and the estimate is somewhat larger for Sweden.

Because the dependent variable in equation (1) is the log of registrations, the estimation does not include vehicles with zero registrations. Dropping these observations could bias the estimates, and column 2 shows the results if the dependent variable is the share of registrations in total quarterly registrations. The point estimates are somewhat smaller in magnitude for all three countries and are not statistically significant for France and Germany. Because the R-squared is so much smaller in column 2 than in Table 2, we prefer the latter regression.

If taxes have a proportional effect on registrations at different levels of the taxes, it would be more appropriate to use the log of the tax rather than the tax in equation (1). However, for some vehicles the French tax is negative, and using the log of the tax would cause those observations to be dropped from the regression. Consequently, we use the tax in equation (1), but column 3 shows the results if we use the log of the tax instead. For France and Germany, the estimated elasticity is similar to that in Table 2, although the estimate for Sweden is larger.

We next show that the results are robust to allowing for different substitution patterns than assumed in equation (1). Equation (1) imposes the assumption that α is the same for all vehicles. Allowing the coefficient to vary across vehicles may better capture substitution patterns than would imposing the restriction that the coefficient is the same. The results (not shown) are similar if we allow the coefficient to vary by market segment or by tax level (i.e., computing quartiles of vehicle taxes by year and estimating a separate coefficient by quartile); Sweden is an exception, where the estimated effect of the tax on registrations is considerably smaller if we allow the coefficient to vary by tax quartile.

Earlier, when we simulated a hypothetical linear CO₂ purchase tax, we noted the caveat that equation (1) does not explicitly model substitution across vehicles. In column 4, we use a nested logit model to estimate the effect of taxes on registrations. The nests correspond to the market segments in Table 1. We allow consumer utility to depend on the price, tax, engine size, weight, height, length, number of cylinders, and horsepower. The price of the vehicle and the tax enter the equation in logs. Similarly to D'Haultfoeuille et al. (2010) and Zachariadis (2011), who model the French and German markets, respectively, we instrument for the price and within-class share of registrations using the average of the same characteristics of the vehicles sold under other brands in the same segment, and the average of vehicles sold under the same brand in other segments. We use observations from 2005–2010 to increase the amount of variation in the vehicle prices and characteristics, thereby easing the identification of the coefficients on these variables. Because the nested logit model is estimated in levels rather than first differences, the results should be compared with column 1 of Table 3. Focusing on the elasticities, the

estimates in columns 1 and 4 for France are similar. The estimate is also similar to D'Haultfoeuille et al. (2010), who use similar data but a different sample period. The nested logit estimate for Germany is quite similar to the estimate in column 1, although the estimate for Sweden is considerably smaller.

Besides estimating a nested logit equation, it is possible to model substitution more explicitly by controlling for taxes of closely related vehicles in equation (1). Adding taxes of closely related vehicles may better capture substitution because the coefficient on the tax variable is identified by differences between the vehicle's tax and the tax of similar vehicles, whereas the coefficient in equation (1) is identified by within-model tax differences. The results are similar to the baseline if we control for the average tax of other vehicles in the same segment or other vehicles in the same tax quartile. These results are not reported but are available upon request.

Next, we discuss aggregation. In Table 2, a unique observation is a model by CO₂ emissions rate by year and quarter. As noted above, this level of aggregation is natural to use as the baseline because it preserves nearly all of the tax variation and reduces measurement error that may be present at lower levels of aggregation. This regression model may introduce some aggregation bias, however, because fuel costs, which are included in equation (1), vary at the specification level. Therefore, for comparison with Table 2, column 5 reports estimates of equation (1) at the specification level rather than model by CO₂ emissions rate. The estimates vary somewhat from Table 2, but remain statistically significant.¹⁷

Aggregation pertains not just to the definition of the vehicle, but also to the time dimension. We use quarterly observations because taxes changed in Germany and Sweden at the beginning of the third and fourth quarters. Vehicle purchases may exhibit seasonality that, by chance, is correlated with the tax

¹⁷ In addition to disaggregating observations as in column 5, we have estimated equation (1) by aggregating observations to the model-fuel type-year-quarter level or the brand-segment-fuel type-year-quarter level (in the two cases, we replace model-year-quarter interactions with brand-fuel type-year-quarter and segment-fuel type-year-quarter interactions). For France, the estimates are statistically significant in each case and are similar to Table 2, but for Germany and Sweden the estimates are much smaller and are not statistically significant. A possible explanation for the difference is that the German and Swedish linear taxes may cause consumers to substitute across versions of the same model, and such substitution is masked by aggregating to the model or brand level, whereas in France the discrete tax changes cause substantial cross-model substitution. The difference across countries demonstrates the importance of using such disaggregated data in the baseline regressions.

changes. For example, consumers who purchase diesel versions of a particular model may be more likely to purchase at some times of year than at others. To address this possibility, we aggregate observations to the annual frequency for France and Sweden and to the half-year for Germany (recall that the baseline regression for Germany includes observations only for 2009). Column 6 reports these results and, as in column 5, although the point estimates differ from the baseline, the estimates remain statistically significant at the 5 percent level.

Regarding the possibility of omitted variables bias, recall that we take several approaches to account for unmeasured demand and supply shocks: restricting the sample to the time period surrounding the tax change, including model-CO₂ fixed effects to control for time-invariant vehicle characteristics, and including model-year-quarter fixed effects to control for demand and supply shocks at the model level.

Table 4 reports a number of regression models that control for particular types of variables omitted from equation (1). Although equation (1) controls for demand shocks at the model level, there may be within-model demand or supply shocks. In column 1, we add to the baseline regression model interactions of vehicle characteristics (the same characteristics as in the nested logit model) with year and quarter fixed effects. The estimate for France is similar to the baseline, although the estimates for Germany and Sweden are larger in magnitude, which suggests that the tax changes in these countries may be correlated with changes in consumer preferences for vehicle characteristics.

Column 3 controls for the French and German vehicle retirement programs and for the Swedish Green Car Rebate. In the baseline regression model, the sample for France ends before the retirement program begins, but it is possible that the anticipation of the program affected registrations during the sample period. Column 3 adds a variable measuring the subsidy offered under the retirement program and includes observations from 2005 to 2010. Column 3 should be compared to column 2, which also includes the years 2005–2010, but does not include the retirement variable. The coefficient on the retirement subsidy (not reported) is positive, as expected, but adding the variable does not affect the

estimate of the tax coefficient. The same is true if we control for the Swedish Green Car Rebate. Because the retirement subsidy in Germany did not depend on the characteristics of the new vehicle, it is difficult to control directly for the effect of the program. However, we can compare the results if we use observations in 2008 and 2010 rather than 2009, thereby excluding the period of time in which the program was in effect. The estimates are close to the baseline (not reported).

The previous discussion about temporal aggregation noted the possibility that seasonality in purchase patterns could bias the estimates. As an alternative to aggregating observations over time, column 4 adds interactions of model fixed effects with quarter-of-year fixed effects. The results are somewhat smaller than the baseline, but the coefficient remains statistically significant for France. Estimates for Germany and Sweden are smaller than the baseline and are not statistically significant.

Another approach to addressing seasonality—and omitted variables more generally—is to conduct “placebo” tests. Suppose that the tax changes occurred during the same calendar month but in a different year; for example, if the tax change actually occurred in January 2008 we might suppose, for this test, that it occurred in January 2007. The fictional tax should not affect registrations in those years, and we would expect the tax coefficient in equation (1) to be small and not statistically significant. If seasonality or other omitted variables (such as the CO₂ emissions rate standards) bias the results, however, the estimate of the tax coefficient could be large or statistically significant. We perform this test for France in 2006–2007, Germany in 2008, and Sweden in 2005. For France and Germany, the coefficient on the placebo variable is small and is not statistically significant; the results reject at the 5 percent level the hypothesis that the elasticity is as large in magnitude as -0.2 . For Sweden, however, the coefficient on the placebo variable is statistically significant, although it is about half of that in Table 2; this further suggests that seasonality explains at least some of the estimated effect of the Swedish tax.

By the end of 2005, the European Commission required that all manufacturers provide labels for new passenger vehicles that describe the vehicle’s CO₂ emissions rate and fuel consumption. The labels

include the CO₂ emissions rate and fuel consumption and provide a color-coded scheme that compares the vehicle's emissions rate with that of other vehicles; the vehicle is assigned a color based on its emissions rate. The objective of the labeling requirement is to provide more information about fuel costs and CO₂ emissions to consumers than they had previously. The labels may have affected consumer preferences and demand for individual models (D'Haultfoeuille et al. 2011). Because the labels were introduced several years before the tax reforms in France and Germany, the labeling requirement is unlikely to significantly bias the estimates for those countries. However, the requirement did occur just prior to the tax reform in Sweden. If the labeling requirement did bias the results, we would expect the bias to be larger in magnitude close to the cutoffs for the color codes. Column 5 shows that the results are similar for all three countries if we omit observations within 5 g CO₂/km of the cutoffs.

The model-year-quarter fixed effects are included to control for demand and supply shocks that affect model registrations proportionately. If the results differ considerably when the fixed effects are omitted, we would be concerned that such shocks are correlated with the tax changes and that the fixed effects may not completely control for them. Column 6 omits the model-year-quarter fixed effects and shows that the estimated coefficients are somewhat larger than, although similar to, the baseline estimates in Table 2.

To summarize the results from Tables 3 and 4, the estimated effect of taxes on registrations varies somewhat across regression models. For France, taxes consistently have a large effect on registrations. The magnitudes imply that the feebate causes all or most of the decrease in the CO₂ emissions rate between 2007 and 2008. Compared to France, the results for Germany and Sweden vary somewhat more across regressions, but nearly all of the regressions indicate that the taxes have had a statistically significant effect on registrations. The estimates for Sweden are smaller than those for Germany and are perhaps less robust.

5.3 Why Do the Results Differ across Countries?

Tables 2–4 showed that taxes affected vehicle registrations for all three countries. The magnitude of the effect is larger for France than for Germany and Sweden, and the French results appear to be more robust. These differences merit some discussion.

There are two fundamental differences between the French tax and the German and Swedish taxes: the French tax is a purchase tax as opposed to an annual circulation tax; and the French tax is nonlinear. A recent body of work has concluded that consumers in vehicle markets and other markets are more responsive to tax and price changes that are more salient (e.g., Busse et al. 2006; Chetty et al. 2009; Finkelstein 2009). In this case, the nonlinearity of the French tax may increase its salience, and therefore increase the effect of a proportional change in the tax on a vehicle's registrations.

A second possible explanation is that consumers respond more to purchase taxes than to annual circulation taxes, possibly because of uncertainty over future circulation taxes.¹⁸ And yet another possibility is that consumer preferences simply differ across countries. These possibilities are not mutually exclusive and, in general, it is difficult to distinguish among them empirically. However, we can gain some insight by examining the tax changes in France in 2010. Partly because of the unexpectedly high fiscal cost of the policy, France slightly decreased the subsidies and increased the taxes on particular vehicles. If we estimate equation (1) using observations from 2009–2010, the magnitude of the tax coefficient is much smaller; the implied elasticity is -0.17 with a standard error of 0.12 , which is not statistically significant. Because the 2010 tax changes were less dramatic than the 2008 changes, the smaller estimates for the 2009–2010 sample suggest that salience plays an important role in explaining the differences across countries. We leave further investigation of this question to future research.

¹⁸ In principle, discount rates could cause consumer responses to differ between purchase and circulation taxes. However, if the circulation tax is expected to be constant, the discount rate should not affect the estimated elasticity of registrations to taxes.

5.4 Long-Run Effects of Vehicle Taxes in France

We investigate whether firms changed the emissions rates of vehicles they supplied to the market, which we refer to as the long-run (or strategic) response to the tax. We identify the long-run response by taking advantage of the timing of the tax change and the discrete changes in taxes across vehicles. Because the tax change was announced in late 2007, it is unlikely that manufacturers changed emissions rates in 2008 in response to the tax. However, manufacturers may have begun responding by 2009, which suggests a comparison of emissions rates of individual vehicles after 2008 with emissions rates before 2008.

As Figure 1 shows, CO₂ emissions rate standards and other factors resulted in a downward trend in emissions rates over time. We use a simple difference-in-difference approach to control for this trend by focusing on the cutoffs around which the tax varies discretely after 2007. That is, we compare the change in emissions rates for vehicles that were initially just above the cutoff with the emissions rates of vehicles that were initially just below the cutoff. The hypothesis is that the emissions rate of vehicles above the cutoff should decrease more than that of other vehicles after 2008.

We provide both graphical and econometric evidence of a very small long-run response from 2008 to 2010. We focus on specifications within 2 g CO₂/km of the cutoffs (we exclude the cutoff for 100 g CO₂/km, around which we find very few specifications). To motivate our main hypothesis in the graphical exercise, Table 5 provides some summary statistics for the vehicles around the cutoffs. The first column shows the share of total registrations in 2007 accounted for by vehicles within 2 g CO₂/km of each cutoff. The second column shows the change in the tax or subsidy for vehicles above the cutoff compared to vehicles below the cutoff. We expect to see a larger manufacturer response around cutoffs with a larger tax or subsidy change (because of the larger benefit of reducing emissions rates) or around cutoffs with a large fraction of registrations (to the extent that fixed costs are associated with reducing the emissions rate of particular vehicles). Therefore, based on the statistics in the table, we expect a larger response around the cutoffs of 120 and 200 g CO₂/km.

To investigate this hypothesis, for each cutoff and year, we compute the ratio of the number of specifications above the cutoff to the number of specifications below the cutoff (in both cases, only counting specifications with positive registrations). Because of the downward trend in emissions rates over time, we expect to observe positive trends prior to 2008. If the tax change in 2008 caused manufacturers to reduce emissions rates of individual vehicles, the trend should become steeper after 2008. For each cutoff over time, Figure 5 plots the ratio of the number of specifications above the cutoff to the number of specifications below the cutoff. We observe some evidence that manufacturers reduced emissions rates around a few of the cutoffs—particularly, around 200 g CO₂/km and, to a lesser extent, around 160 g CO₂/km—but not for other cutoffs; overall, we do not find strong graphical evidence of a large response.

To implement the difference-in-difference estimation strategy, we focus on changes in emissions rates across years for specifications that are within 2 g CO₂/km of the cutoffs in the initial year. We estimate the following equation

$$\Delta \ln e_{iy} = \delta_0 POS_i + \delta_1 POS_i * POST_y + \tau_y + \varepsilon_{iy} \quad (2)$$

where the dependent variable is the change in the log emissions rate of specification i between years $y - 1$ and y , POS_i is a dummy variable equal to one for specifications with emissions rates above the cutoff in year $y - 1$, $POST_y$ is a dummy variable equal to one in years 2009 and 2010, $POS_i * POST_y$ is the interaction of the two variables, τ_y includes a set of year fixed effects, and δ_0 and δ_1 are coefficients. Note that equation (2) is estimated at the specification level as opposed to equation (1), which is estimated at the vehicle (model by emissions rate) level. If manufacturers reduce the emissions rates of specifications just above the cutoff, the coefficient δ_1 would be negative.

If there are fixed costs to reducing emissions rates, we would expect to observe a larger long-run response for top-selling specifications. Consequently, all observations are weighted by the specification's registrations in year $y - 1$ (the qualitative conclusions are the same if regressions are unweighted).

Table 6 reports the estimation results. For comparison with the difference-in-difference results, column 1 reports a simple difference estimate that includes observations in 2009 and compares the change in emissions rates of vehicles that are above and below the cutoff in 2008. The point estimate on the POS_i dummy variable indicates that emissions rates of specifications above the cutoffs decreased less than those of specifications below the cutoffs.

Column 2 reports the estimates of equation (2) and shows that the emissions rates decrease more for specifications that were initially above the cutoff after 2008 compared to specifications that were above the cutoff in earlier years. However, this effect is less than 1 percent and is statistically significant at about the 10 percent level. Therefore, we conclude that very few specifications moved from above to below the cutoffs.¹⁹

The long-run response could be greater near cutoffs around which the tax changes by a large amount. For example, the tax for specifications above 160 g CO₂/km is 200 euros, whereas there is no tax for specifications between 130 and 160 g CO₂/km; the difference in the tax from 128 to 132 g CO₂/km is 200 euros. This difference varies across the cutoffs, and we would expect to observe a larger decrease in the emissions rates in 2009 around cutoffs that have a large difference, such as around 200 g CO₂/km, where the difference is 850 euros. In column 3, we add to equation (2) the interaction of this difference variable with the $POS_i * POST_y$ variable (as well as all other main effects and interactions). We expect the coefficient on the triple interaction to be negative. The coefficient is negative, but the magnitude is small (note that the tax difference is divided by 1,000 so that the coefficients are readable); similarly to column

¹⁹ We have also estimated a model that estimates the probability that a specification that is above a cutoff in a particular year is below the cutoff in the following year. The estimates suggest that the probability increased after 2008, but the estimate is small and is not statistically significant.

2, the coefficients imply that the emissions rate decreases by about 0.5 percent for vehicles that are initially above the cutoff.

The remaining columns report other regression models that account for several possible reasons why we do not detect a large long-run effect of the tax in columns 2 or 3. There is probably a substantial fixed cost to reducing emissions rates of particular specifications. Consider a manufacturer that sells most of its vehicles outside of France. In that case, the manufacturer may not choose to incur the fixed cost of reducing emissions rates if the benefit of doing so would be realized only for the small number of vehicles that sell in France. In fact, of the top 10 selling brands in France, except for the French brands (Citroen, Peugeot, and Renault), at least about 90 percent of European sales occur outside of France. In contrast, French sales account for about 30 percent of European sales for the French brands. Therefore, column 4 includes only the three French brands, but we do not find evidence of a long-run response.

We may not detect a significant long-run response because we are able to perform the analysis for only two years after the tax change. In the European market, many vehicles are redesigned over roughly five- to seven-year design cycles. Therefore, manufacturers intending to redesign their vehicles in response to the tax may not have done so by 2010. However, many smaller modifications to the power train can be made within one to two years (Klier and Linn 2012), and these changes could reduce emissions rates by enough to move specifications within 2 g CO₂/km of a threshold across the threshold. Our analysis should detect such responses. Nevertheless, the cost of performing such modifications may vary over the design cycle. To allow for this possibility, in column 5, we add to column 2 the interaction of the

$POS_i * POST_y$ variable with a variable equal to one if a new version of the model began production in 2009 or 2010.²⁰ Based on the triple interaction term in column 5, emissions rates decreased slightly more for such vehicles, but the difference is small and is not statistically significant.

²⁰ We use production data from CSM to identify the redesign year for each model in the registration data. The production data include the year in which a new version of the model began production. We assume that specifications of the same model follow the same production cycle. For vehicles imported to Europe we rely on data

In columns 6–8, we repeat the estimation models in columns 2, 4, and 5 at the model level rather than the specification level (where we use the average emissions rate across specifications of the same model to construct the dependent variable). The aggregation accounts for the fact that the specification-level estimation includes only specifications that are sold in consecutive years. Because specifications sold in consecutive years account for about 85 percent of registrations, the regressions in columns 1–5 do not include a sizeable fraction of registrations. The conclusion for the model-level analysis is similar to the other analysis, although we do find some evidence in column 7 that the French brands responded to the tax. The sample size is very small, however, and given that we did not find a response for the specification-level analysis for the French brands in column 4, we treat this evidence as suggestive.

Besides those reported in Table 6, we have tried other regression models with similar results, including (a) omitting observations around the cutoffs for 160 and 165 g CO₂/km because the proximity of the cutoffs may make it difficult to detect a response and (b) restricting the sample to observations within 1 g CO₂/km of the cutoffs. Thus, we find very little evidence that firms reduced the emissions rates of individual specifications in response to the tax. The concluding section discusses some implications of this result.

6 Effect of Reducing CO₂ Emissions Rates on Manufacturers' Profits

This section uses the empirical estimates of equation (1) to estimate the effect of a CO₂ emissions rate standard on manufacturers' profits. We first describe the framework for estimating the short-run change in profits, and then present the estimates. Because we did not find evidence that the French tax affected emissions rates of individual vehicles, we do not perform a long-run analysis.

for production in the nearest region. Except for a few low-selling models, we are able to match the models in the Polk data to the CSM data. The matched sample includes almost 99 percent of total registrations in the Polk data.

6.1 Framework for Estimating Manufacturer Profits

We discuss the main assumptions under which we can use the effect of CO₂ taxes on registrations to estimate the change in manufacturers' profits from reducing the average CO₂ emissions rates of their vehicles.

To simplify the exposition, we focus on a monopolist selling multiple vehicles. The demand for product j is given by $q(p_j, e_j)$, where p_j is the price and e_j is the emissions rate. Quantity demanded decreases with the price and emissions rate (because the emissions rate is proportional to fuel consumption and inversely proportional to fuel economy). We make two additional simplifying assumptions—that emissions rates are fixed and that cross-price elasticities of demand equal zero. The assumption of fixed emissions rates could be easily relaxed to allow for a long-run response to the tax. Relaxing the assumption on cross-price elasticities does not affect the main results. Vehicle production entails zero fixed costs and positive, but constant, marginal costs, c .

The government imposes an emissions rate tax. Consistent with the actual tax reforms in France, Germany, and Sweden, the tax is imposed on consumers. We first consider a purchase tax and return to circulation taxes at the end of the section. Consumers pay a tax of $\tau(e_j - \bar{e})$ for purchasing the vehicle, where we refer to τ as the tax rate and to \bar{e} as the pivot; for vehicles with emissions rates above the pivot, the vehicle is taxed, and for vehicles with emissions rates below the pivot, the vehicle is subsidized. p_j is the tax-exclusive price.

The monopolist chooses prices to maximize profits

$$\max_{p_j} \sum_j [p_j + \tau(e_j - \bar{e}) - c]q(p_j, e_j) \quad (3)$$

The first-order condition is

$$q(p_j, e_j) + [p_j + \tau(e_j - \bar{e}) - c] \frac{\partial q}{\partial p} = 0 \quad (4)$$

This is the standard first-order condition for a profit-maximizing monopolist.

To compare with the emissions rate tax, suppose that the government instead imposes an emissions rate standard, where the registrations' weighted average emissions rate must meet a standard, C :

$$\sum_j q_j e_j = C \sum_j q_j .$$

The monopolist maximizes profits subject to the standard

$$\max_{p_j} \sum_j (p_j - c) q(p_j, e_j) \quad (5)$$

$$\text{s.t. } \sum_j q_j e_j = C \sum_j q_j$$

The first-order condition is

$$q(p_j, e_j) + [p_j - c + \mu(e_j - C)] \frac{\partial q}{\partial p} = 0 \quad (6)$$

where μ is the multiplier on the constraint. By applying the envelope theorem, we can interpret μ as being

proportional to the marginal effect of the emissions rate standard on profits: $\frac{\partial \pi}{\partial C} = \mu \sum_j q_j$, where π is

the firm's profits. In other words, the marginal effect of the constraint on profits, per vehicle, equals μ .

Comparing the first-order conditions from the emissions rate tax and the emissions rate standard, equations (4) and (6), we observe the equivalence between the two policies. If the pivot, \bar{e} , is set equal to the emissions rate standard, C , the tax rate equals the multiplier on the constraint ($\tau = \mu$). Therefore, the tax rate is equal to the marginal effect of the emissions rate standard on profits per vehicle.

This equivalence suggests a simple approach to estimating the change in manufacturers' profits from reducing average emissions rates. We assume that all manufacturers are evaluated jointly for compliance and estimate the average change in profits across manufacturers; below we report separate results by brand. Considering the case of France, we begin by setting the pivot equal to the registration-weighted average emissions rate across all manufacturers in France in 2007. We then introduce a very small tax rate, τ_0 , and use equation (1) to calculate the counterfactual equilibrium registrations and change in the registration-weighted average emissions rate, Δe_0 . The change in profits per vehicle is computed as

$$\frac{\Delta \pi}{\sum_j q_j} \approx \tau_0 \Delta e_0 \quad (7)$$

The pivot is reset to the new average emissions rate, and the tax rate is increased by a small amount to τ_1 . The new equilibrium registrations and change in emissions rate is calculated. The procedure is repeated until the counterfactual emissions rate is 5 g CO₂/km below the 2007 equilibrium. This emissions rate change is chosen to represent a nonmarginal but small decrease. The sum over iterations of the estimated change in profits is computed to obtain the total change in profits per vehicle (throughout, we assume that the tax does not affect total registrations, which is consistent with a linear regression of aggregate quarterly registrations on the average tax).

Note that the analysis is the same if the tax is implemented as a circulation tax as opposed to a purchase tax. Therefore, we can use the empirical results for Germany and Sweden to estimate manufacturers' profits in those countries, but the interpretation is slightly different. With a purchase tax, equation (7) yields an estimate of the change in profits per vehicle. With a circulation tax, equation (7) yields the change in profits per vehicle per year of the vehicle's life; we can estimate the change in profits per vehicle by assuming a discount rate and vehicle lifetime.

Compared to nearly all other estimates in the literature (e.g., Goldberg 1998; Jacobsen 2010; and Klier and Linn 2012), this approach relies on few assumptions about consumer demand and producer behavior. However, we note two caveats to our approach. First, as with the simulations in Table 2, we assume that the estimated tax coefficient captures how registrations would respond to a hypothetical linear CO₂ emissions rate tax. The similarity between the baseline estimate and the nested logit results, as well as the overall robustness of the estimates, supports this assumption for France. On the other hand, this assumption is perhaps stronger for France because the empirical estimates are based on a discrete tax rather than a linear tax.

Second, we assume that consumers would respond similarly whether a price change occurred under a CO₂ emissions rate standard or under a linear CO₂-based tax. D'Haultfoeuille et al. (2010) find some evidence that consumers respond more to vehicle price changes under the French CO₂ tax than to price changes during other time periods (where price is defined as the tax-inclusive price). This suggests, though not conclusively, that consumers may respond less to prices under an emissions rate standard. To the extent that this is true, we would underestimate the reduction in manufacturers' profits from achieving the emissions rate standard. However, although that paper reports a statistically significant difference in the response of registrations to tax-induced price changes, the magnitude of the difference is relatively small; this suggests that the resulting bias may also be small.

6.2 Short-Run Change in Manufacturers' Profits

Table 7 reports the short-run change in profits from two separate simulations. The first column shows results assuming that α in equation (1) is the same for all vehicles. We also allow for the possibility that α , and therefore the change in profits, varies across brands. We assign each brand a quintile based on its average emissions rate in 2007. Panel A reports the total market share and average CO₂ emissions rate for each quintile. The table shows the top two brands based on 2007 registrations.

Column 1 of Panel B repeats the baseline estimate of α from Table 2. We reestimate equation (1) allowing a separate α_g for each quintile g . The coefficients, which are reported in columns 2–6, are all statistically significant at the 1 percent level, and the point estimates vary substantially across quintiles.

Panel C reports the simulation results. Starting from the observed 2007 registrations of each vehicle, in column 1 we use the baseline estimate of α in Table 2 to simulate the change in market shares from imposing a linear purchase tax of 0.01 euros/g CO₂/km. We then estimate the slope of the profit function using equation (7) and repeat the procedure, increasing the tax by 0.01 euros/g CO₂/km until the average emissions rate is 5 g CO₂/km lower than the initial level. A tax of 10.17 euros/g CO₂/km achieves this emissions rate reduction, at an average profit reduction of 24.1 euros/vehicle. Assuming that the vehicle is driven 15,000 km/year for 12 years and assuming no rebound effect, this estimate translates to a change in profits of about 26 euros/ton of avoided CO₂ emissions. The estimate of 24.1 euros/vehicle is smaller than the fine of 95 euros/g CO₂/km that manufacturers pay for exceeding the standard by more than 3 g CO₂/km. This suggests that, at least for small changes in emissions rates compared to the manufacturers' initial levels, most manufacturers should meet the standard without paying the fine.

The remaining columns show substantial variation across the categories of firms. Each column reports the results of a separate simulation. The tax is imposed on all vehicles in the data, and the tax is increased until the emissions rate of the corresponding quintile has decreased by 5 g CO₂/km. The tax required to achieve this emissions reduction varies by nearly a factor of four across quintiles, from 5.46 to 20.29 euros/g CO₂/km. Likewise, the profit reduction per vehicle varies by about a factor of four, from 13.3 to 48.4 euros/vehicle. We have also performed simulations based on estimates of equation (1) that allow the tax coefficient to vary by market segment, quartile of the tax distribution, or brand. The results, which are not reported but are available upon request, likewise suggest considerable heterogeneity in the change in profits across manufacturers; across all simulations and categories, the largest estimate is about 80 euros/vehicle.

Table 8 reports an analogous set of results for Germany and Sweden. Because the estimates in Section 5 for these countries are less robust than the estimates for France, we emphasize the latter. Consequently, we only report simulations for the overall average profit reduction, and not separately by quintile.²¹ As noted above, because Germany and Sweden use circulation taxes, we interpret the estimates at the bottom of Panel B as the change in profits per year of the vehicle's lifetime. Consistent with the estimates in Table 2, the change in profits per vehicle is estimated to be much higher for Sweden than for Germany. Assuming a discount rate of 10 percent and a lifetime of 12 years, we can multiply the estimates in Panel B by about 7.5 to obtain the change in profits per vehicle. The change in profits per vehicle is higher in Germany and Sweden than in France.

7 Conclusions

It has become increasingly common in Europe and elsewhere to tax vehicle purchase and ownership according to CO₂ emissions rates. Very little research, however, has addressed the effectiveness of these policies at reducing emissions rates. In 2008, France began taxing and subsidizing purchase according to the vehicle's emissions rate, where the level of the tax or subsidy changed in discrete steps. In contrast, Germany and Sweden introduced CO₂-based circulation taxes that are linear in emissions rates—Germany in mid-2009 and Sweden in mid-2006. These changes provide an opportunity to examine the effectiveness of such policies and to estimate the change in manufacturers' profits from reducing emissions rates.

First, we take advantage of the rapid and large changes in taxes to distinguish between the short- and long-run effects of the policies, where the short run is defined as the period of time in which vehicle characteristics are fixed. For France, we find a large and statistically significant negative short-run effect of the taxes on vehicle registrations. The results are reasonably robust across alternative estimation strategies, and the baseline estimates imply an elasticity of vehicle registrations to vehicle taxes of about –

²¹ If we define quintiles here as we do for France, we do not precisely estimate the coefficient on the lowest quintile in either Germany or Sweden.

0.4. We also find evidence that taxes negatively affected registrations in Germany and Sweden, although the results are smaller and vary somewhat more across alternative estimation strategies compared to those for France, particularly in the case of Sweden.

Because the French tax and subsidy change around discrete cutoffs, we implement a simple difference-in-difference strategy to estimate manufacturers' long-run response to the tax, which includes changes in the emissions rates of individual specifications. We compare the change in emissions rates for vehicles that are initially just above the cutoffs against the change in emissions rates for vehicles below the cutoffs, and compare this difference before and after the tax change. We find little evidence of a long-run response.

The theoretical equivalence between a linear emissions rate tax and an emissions rate standard suggests that we can use the estimated effect of a tax on registrations to estimate the change in manufacturers' profits from a standard that imposes a small emissions rate reduction. Because the empirical results in the first stage were more robust for France than for Germany and Sweden, we focus on estimating the profit effects in France. A reduction in emissions rates of 5 g CO₂/km reduces profits by about 24 euros/vehicle. This estimate varies across brands from about 10 to 50 euros/vehicle, depending on the initial emissions rate of the brand. This variation suggests that a feebate or emissions rate standard with trading could be more cost effective than an emissions rate standard with limited trading, but we leave quantifying these benefits for future research.

Why do we find little evidence that manufacturers responded to the French tax by reducing the emissions rates of their vehicles in the long run? In the short-run model in Section 6, manufacturers reduce their average emissions rates by changing vehicle prices. The simulations suggested that the short-run change in profits from reducing emissions rates was relatively small. Evidence from the U.S. market indicates that reducing emissions rates by redesigning vehicles is less costly to manufacturers than adjusting prices in the short run (e.g., Whitefoot et al. 2011; Klier and Linn 2012). If the same were true in the French market, we would expect the French tax to cause manufacturers to reduce the emissions rates of their

vehicles in the long run. However, the short-run profit estimates in Section 6 do not include fixed costs, which manufacturers would not incur when adjusting vehicle prices in the short run (i.e., as in Table 7), but which they would incur when reducing vehicle emissions rates in the long run.²² In fact, some evidence suggests that fixed costs affect manufacturers' decisions, as we do observe a slightly larger long-run response among French brands than other brands. For most manufacturers, France accounts for a very small fraction of total sales in Europe. Consequently, it may not be profitable for these manufacturers to reduce the emissions rates of individual vehicles around the cutoffs if doing so incurs large fixed costs. The French brands, for which sales in France account for at least 30 percent of total European sales, may be more likely to recover the fixed costs and therefore more likely to adjust emissions rates in response to the tax.

The existence of large fixed costs would have important policy implications. If fixed costs do reduce manufacturers' long-run responses, the cost-effectiveness of emissions rate taxes or standards could be improved by harmonizing the policies across countries because manufacturers would have greater incentives to reduce the emissions rates of their vehicles.

Fixed costs to redesigning vehicles may explain the limited evidence of a long-run response to the French tax, but another possibility is that we are able to examine responses only for the first two years after the tax change. Because vehicle design cycles are typically about 5-7 years, manufacturers may not have had time to respond. Whether either of these explanations is the correct explanation remains a question for future research.

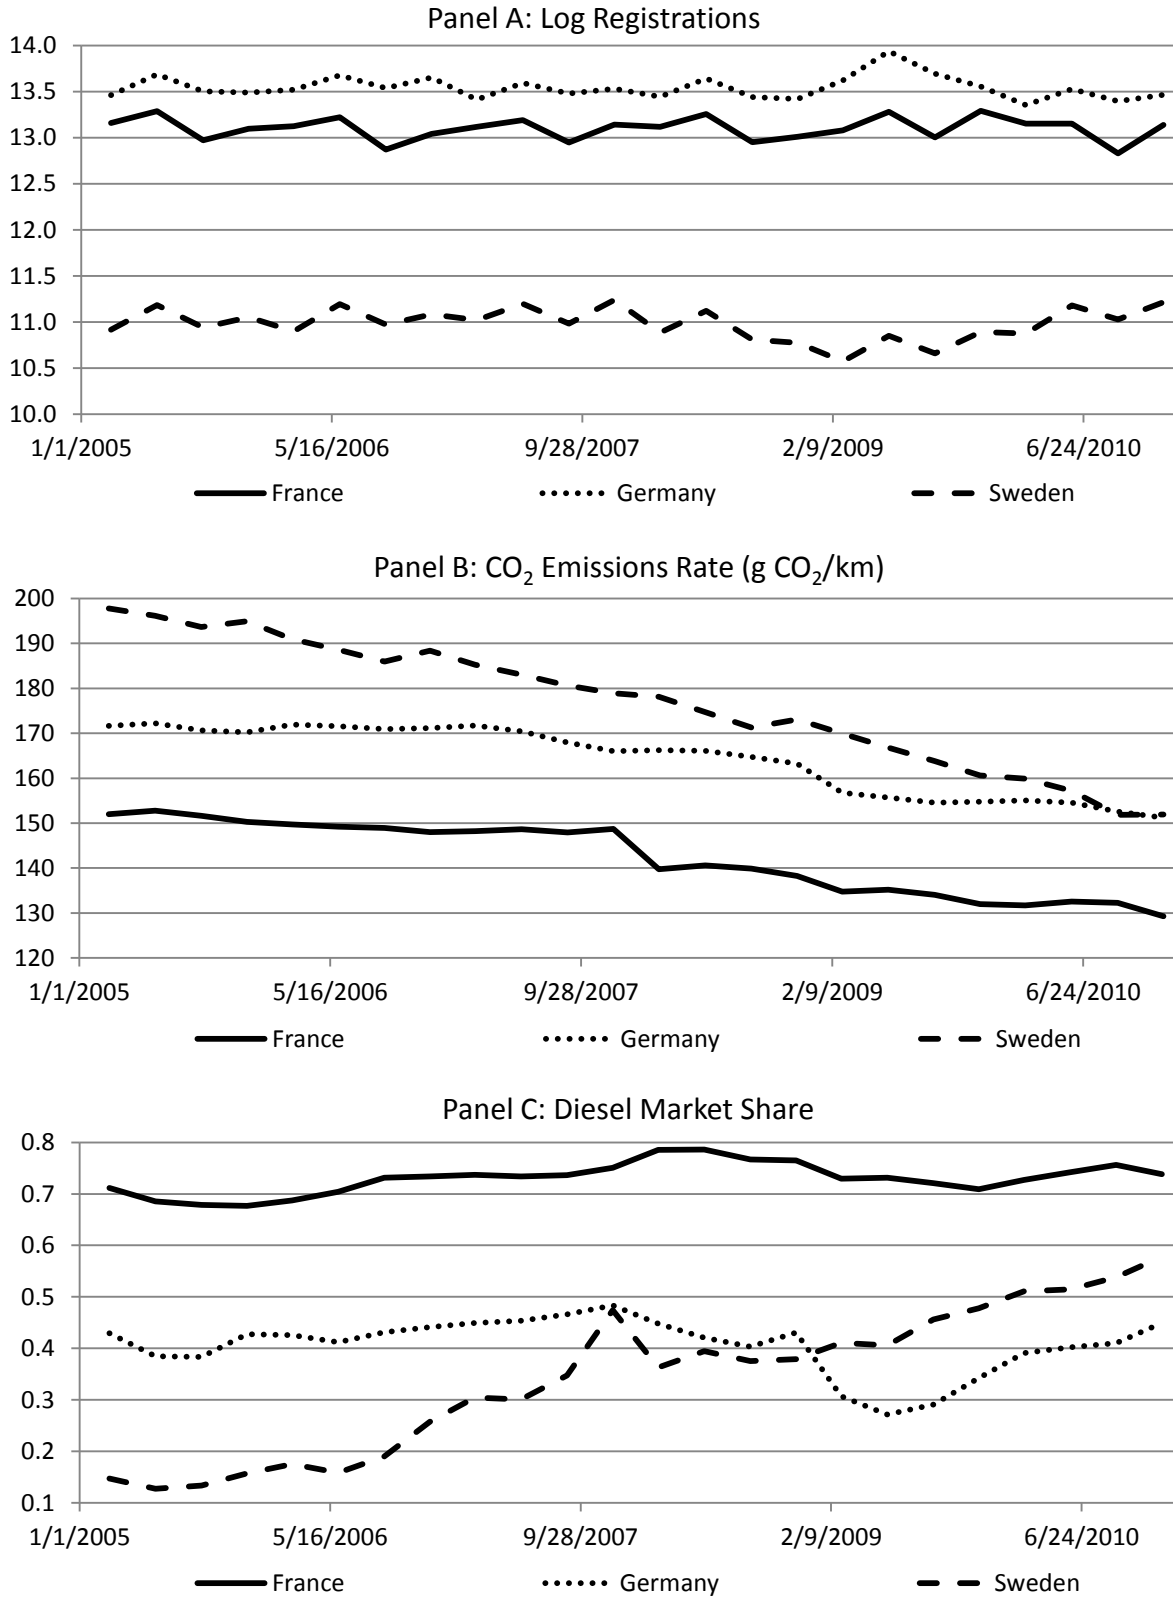
8 References

Anderson, S.T., I.W.H. Parry, J. Sallee, and C. Fischer. 2011. Fuel Economy Standards: Impacts, Efficiency, and Alternatives. *Review of Environmental Economics and Policy* 5(1): 89–108.

²² The studies of the U.S. market do not include fixed costs and are only suggestive that the long-run costs, including fixed costs, are smaller than the short-run costs.

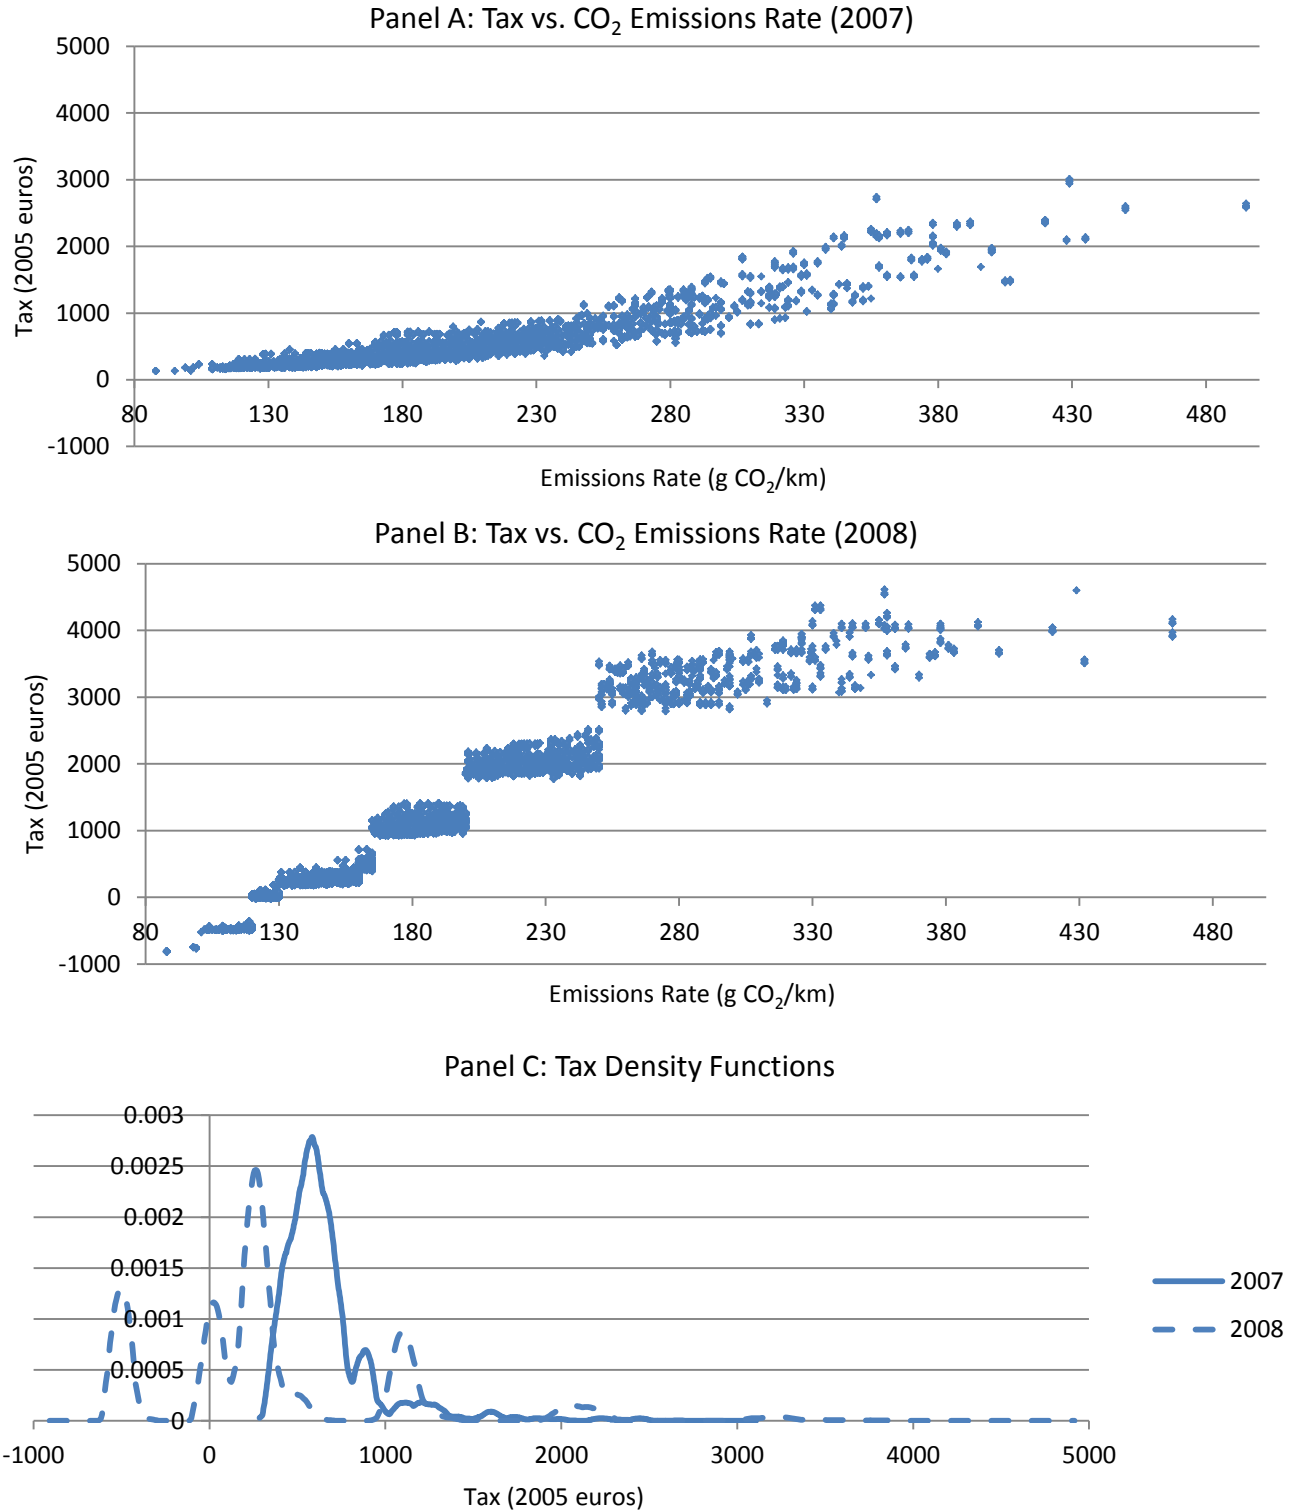
- Anderson, S.T., and J. Sallee. 2011. Using Loopholes To Reveal the Marginal Cost of Regulation: The Case of Fuel Economy Standards. *American Economic Review* 101(4): 1375–1409.
- Busse, M., J. Silva-Risso, and F. Zettelmeyer. 2006. \$1000 Cash-Back: The Pass-Through of Auto Manufacturing Promotions. *American Economic Review* 96(4): 1253–1270.
- Chetty, R., A. Looney, and K. Kroft. 2009. Saliency and Taxation: Theory and Evidence. *American Economic Review* 99(4): 1145–1177.
- D’Haultfoeuille, X., I. Durrmeyer, and P. Fevrier. 2010. What Did You Expect? Lessons from French Bonus/Malus. *Revue Économique* 62.
- D’Haultfoeuille, X., I. Durrmeyer, and P. Fevrier. 2011. Environmental Policies and Consumer Preferences: Lessons from the French Automobile Market. Working Paper.
- Finkelstein, A. 2009. E-Ztax: Tax Saliency and Tax Rates. *Quarterly Journal of Economics* 124(3): 969–1010.
- Goldberg, P. 1998. The Effects of the Corporate Average Fuel Efficiency Standards in the U.S. *The Journal of Industrial Economics* 46: 1–33.
- Huse, Cristian. 2012. Fast and Furious (and Dirty): How Asymmetric Regulation May Hinder Environmental Policy. Working Paper, Stockholm School of Economics.
- Jacobsen, M. 2010. Evaluating Fuel Economy Standards in a Model with Producer and Household Heterogeneity. Working Paper, University of California at San Diego.
- Klier, T., and J. Linn. 2011. Fuel Prices and New Vehicle Fuel Economy—Comparing the United States and Western Europe. RFF Discussion Paper 11-37.
- Klier, T., and J. Linn. 2012. New-Vehicle Characteristics and the Cost of the Corporate Average Fuel Economy Standard. *RAND Journal of Economics* 43(1): 186–213.
- Sallee, J. M. and J. Slemrod. Forthcoming. Car Notches: Strategic Automaker Responses to Fuel Economy Policy. *Journal of Public Economics*.
- TNO. 2011. *Support for the Revisions of Regulation (EC) No 443/2009 on CO₂ Emissions from Cars: Final Report*. Delft: the Netherlands: TNO.
- U.S. Environmental Protection Agency. 2010. *Final Rulemaking To Establish Light-Duty Vehicle Greenhouse Gas Emission Standards and Corporate Average Fuel Economy Standards: Regulatory Impact Analysis*. Washington, DC: U.S. Environmental Protection Agency.
- Whitefoot, K., M. Fowlie, and S.J. Skerlos. 2011. Product Design Response to Policy Intervention: Evaluating Fuel Economy Standards Using an Engineering Model of Product Design. Working Paper.
- Zachariadis, T. 2011. Environmental and Economic Effects of CO₂-Based Automobile Taxes in Germany. Working Paper.

Figure 1. Registrations, Diesel Market Share, and CO₂ Emissions Rates



Notes : The figure shows log quarterly registrations, registration-weighted quarterly average CO₂ emissions rates, and the share of registrations accounted for by diesel fuel vehicles, by country, for 2005–2010. All data from R.L. Polk & Co.

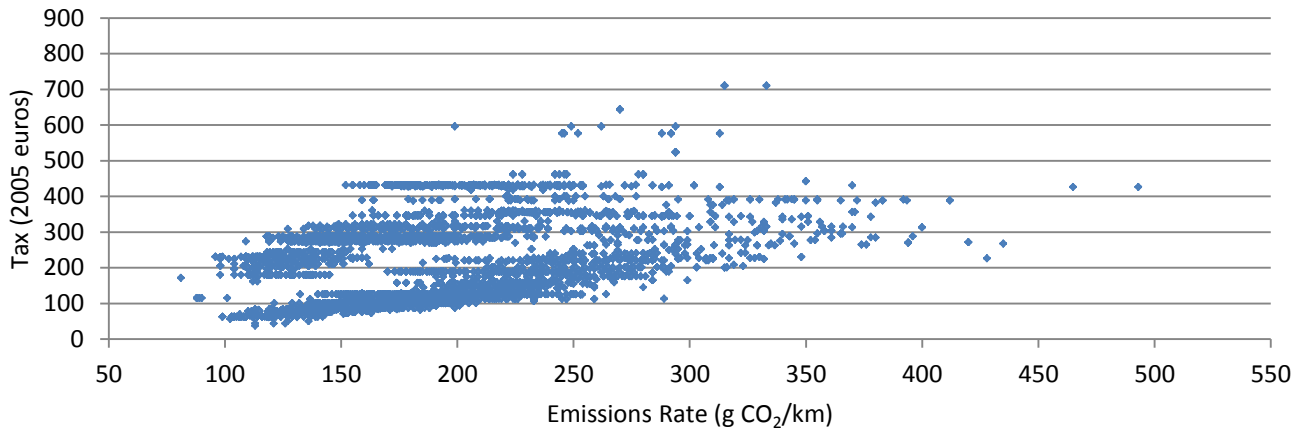
Figure 2. Purchase Taxes in France



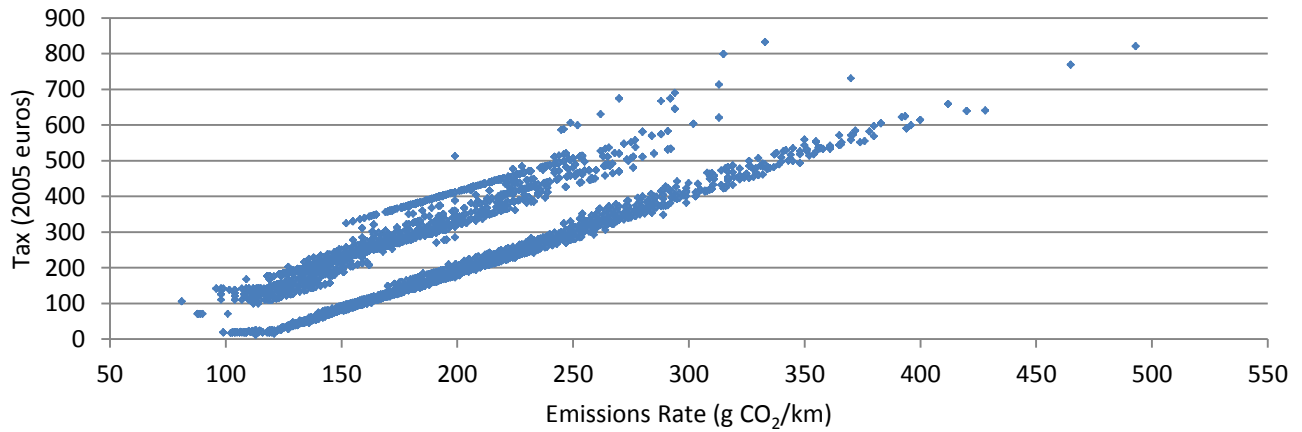
Notes : Panel A plots the purchase tax against the emissions rate for each vehicle specification with positive registrations in France in 2007. Panel B shows a similar plot using specifications in 2008 and the purchase taxes that applied in 2008. To construct Panel C, for each specification in 2007, two variables are constructed, using the 2007 and 2008 tax rates. Panel C plots the estimated density function of the two tax variables.

Figure 3. Annual Circulation Taxes in Germany

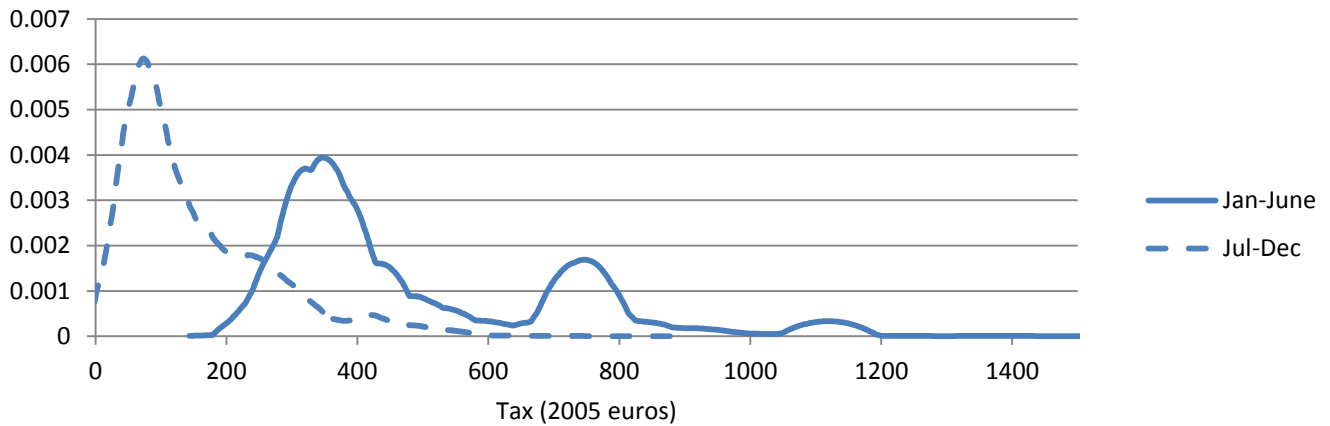
Panel A: Tax vs. CO₂ Emissions Rate (2009 Jan–Jun)



Panel B: Tax vs. CO₂ Emissions Rate (2009 Jul–Dec)



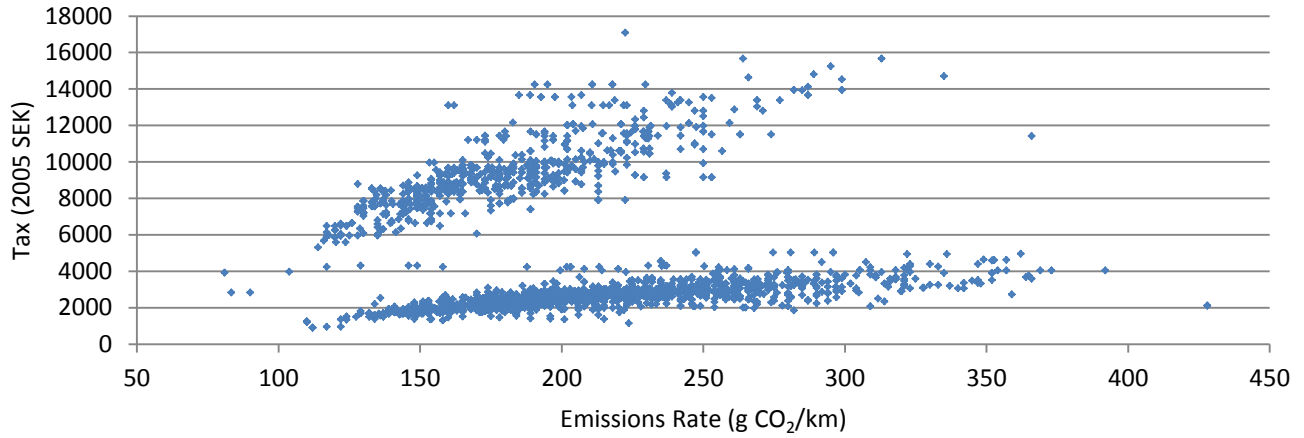
Panel C: Tax Density Functions (2009)



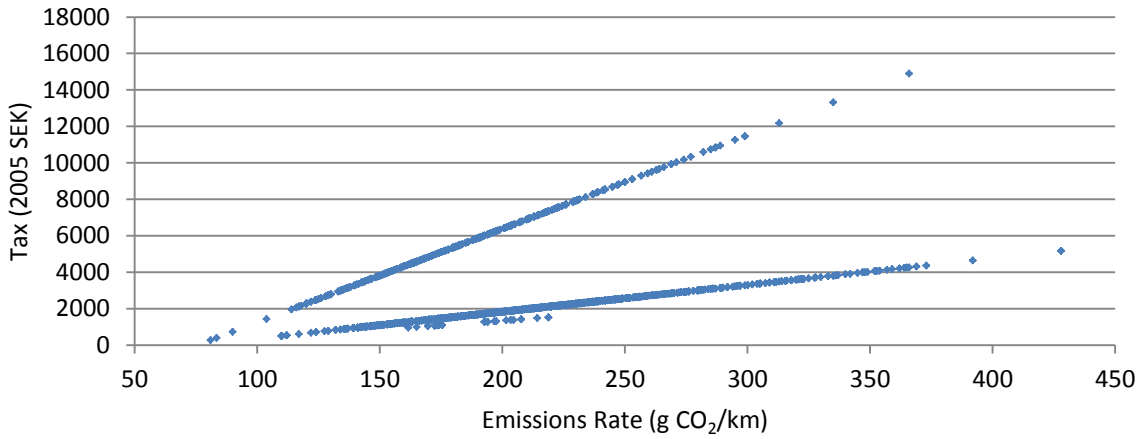
Notes: The figure is constructed similarly to Figure 2, using specifications that have positive registrations in the first or second half of 2009 in Germany.

Figure 4. Annual Circulation Taxes in Sweden

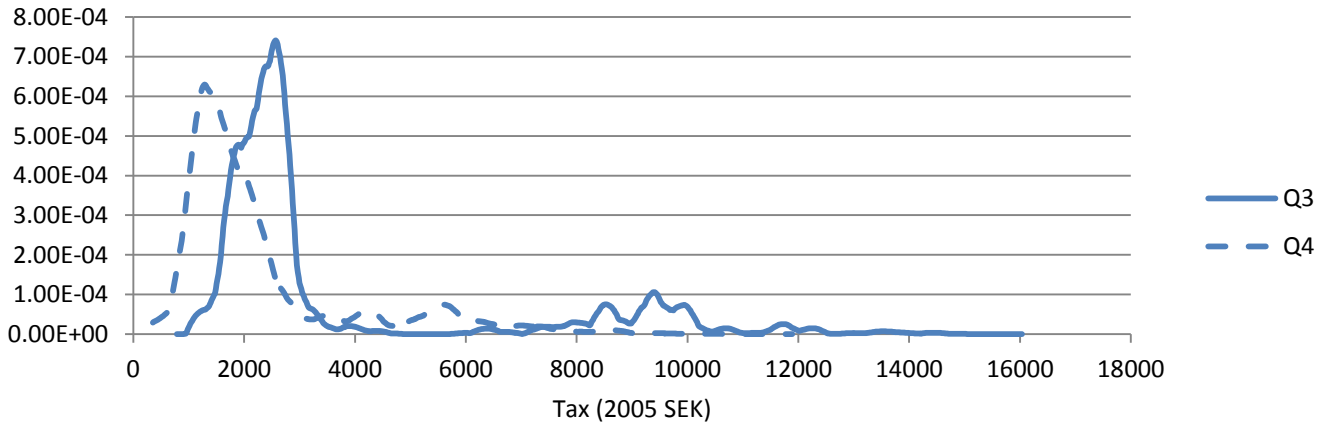
Panel A: Tax vs. CO₂ Emissions Rate (2006 Q3)



Panel B: Tax vs. CO₂ Emissions Rate (2006 Q4)

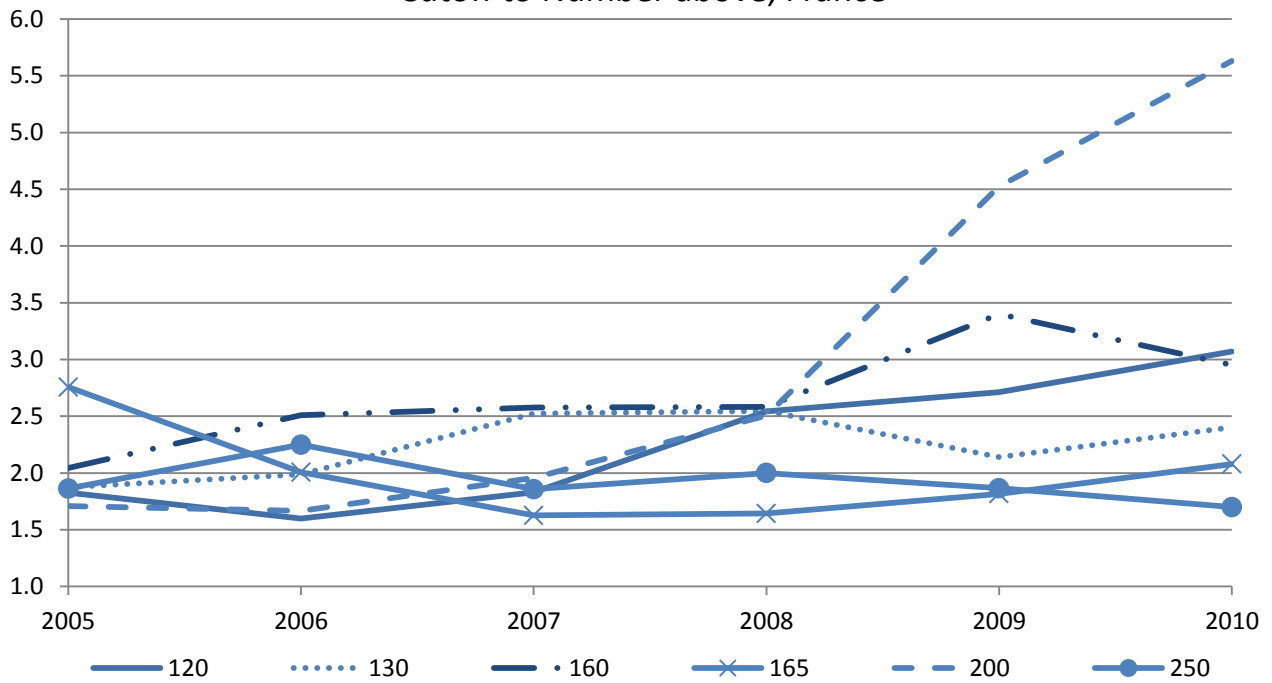


Panel C: Tax Density Functions (2006)



Notes : The figure is constructed similarly to Figure 2, using specifications that have positive registrations in the third or fourth quarter of 2006 in Sweden. SEK, Swedish kronor.

Figure 5. Ratio of Number of Specifications below Emissions Rate Cutoff to Number above, France



Notes : For each year and cutoff, the number of specifications in France below the cutoff and within 2 g CO₂/km of the cutoff is divided by the number of specifications in France above the cutoff and within 2 g CO₂/km. The figure plots the ratio for each year and cutoff.

Table 1.

Market Shares and Average CO₂ Emissions Rate by Market Segment

	<u>Panel A: France</u>			
	<u>Market share</u>		<u>Emissions rate (g CO₂/km)</u>	
	<u>2005</u>	<u>2010</u>	<u>2005</u>	<u>2010</u>
Mini (Renault Twingo)	0.067	0.130	131.9	115.6
Small (Renault Clio)	0.327	0.384	134.7	121.7
Lower Medium (Renault Scenic)	0.401	0.368	151.1	135.8
Medium (Peugeot 407)	0.161	0.097	174.3	159.3
Upper Medium (Peugeot 607)	0.042	0.019	232.2	198.5
Large (Nissan Patrol GR)	0.002	0.002	274.3	239.3
	<u>Panel B: Germany</u>			
	<u>Market share</u>		<u>Emissions rate (g CO₂/km)</u>	
	<u>2005</u>	<u>2010</u>	<u>2005</u>	<u>2010</u>
Mini (Smart Fortwo)	0.056	0.076	133.3	120.4
Small (VW Polo)	0.215	0.220	148.0	135.2
Lower Medium (VW Golf)	0.408	0.415	165.1	150.7
Medium (VW Passat)	0.219	0.201	187.5	171.5
Upper Medium (Audi A6)	0.095	0.080	227.5	193.2
Large (Audi A8)	0.008	0.009	262.1	234.1
	<u>Panel C: Sweden</u>			
	<u>Market share</u>		<u>Emissions rate (g CO₂/km)</u>	
	<u>2005</u>	<u>2010</u>	<u>2005</u>	<u>2010</u>
Mini (Toyota Aygo)	0.013	0.043	144.5	113.3
Small (Renault Clio)	0.135	0.150	154.4	132.3
Lower Medium (VW Golf)	0.309	0.327	177.0	149.6
Medium (Volvo V50)	0.321	0.299	207.6	164.9
Upper Medium (Volvo V70)	0.219	0.178	232.4	176.6
Large (Audi A8)	0.002	0.003	278.9	244.6

Notes : The table reports the market share and registration-weighted average CO₂ emissions rate for each market segment in 2005 and 2010. For each segment and country, the top-selling model is indicated in parentheses.

Table 2.

<u>Effect of Taxes on Vehicle Registrations in France, Germany, and Sweden</u>			
	(1)	(2)	(3)
	<u>France</u>	<u>Germany</u>	<u>Sweden</u>
<u>Panel A: Coefficient estimates</u>			
Tax coefficient	-0.561 (0.113)	-1.447 (0.334)	-0.060 (0.026)
R squared	0.36	0.32	0.37
Observations	13,221	10,717	8,942
<u>Panel B: Elasticity of registrations to taxes</u>			
Elasticity	-0.417 (0.084)	-0.322 (0.074)	-0.244 (0.106)
<u>Panel C: Effect of actual tax changes on emissions rates</u>			
Initial emissions rate (g CO ₂ /km)	148.4	156.2	188.5
Change in emissions rate (g CO ₂ /km)	-7.95 (1.26)	-1.67 (0.37)	-0.57 (0.26)

Notes : Panel A reports the estimate of the coefficient on vehicle taxes in equation (1). Observations are by vehicle (a unique model-fuel type-CO₂ emissions rate), quarter, and year, and all variables are in first differences. The dependent variable is the log of registrations. The tax is divided by 1,000 to improve readability. All regressions include model-quarter fixed effects, fuel costs, fuel prices, and fuel consumption. Standard errors are clustered by market segment and quarter. Panel B reports the elasticity of vehicle registrations to taxes using the coefficient estimate in Panel A and the corresponding sample mean. Panel C reports the change in the average CO₂ emissions rate (g CO₂/km) from under the assumption that the post-reform tax rates applied in 2007 in France, in the first half of 2009 in Germany, and in the first three quarters of 2006 in Sweden. The change in emissions rates is estimated using the coefficient from Panel A and the change in tax rate for each vehicle in the market prior to the reform. Standard errors are estimated by the delta method, assuming that actual registrations and CO₂ emissions rates are measured without error.

Table 3.

Functional Form Assumptions and Aggregation						
Specification	Equation (1) in levels	Dep var is share of regs	Log taxes	Nested logit	Observations by spec-quarter	Observations by year or half-year
	(1)	(2)	(3)	(4)	(5)	(6)
<u>Panel A: France</u>						
Tax coefficient	-0.364 (0.037)	-0.158 (0.117)	-0.407 (0.131)	-0.196 (0.020)	-0.451 (0.096)	-0.185 (0.092)
R squared	0.86	0.15	0.36	0.93	0.24	0.05
Obs	15,025	15,895	12,844	42,044	35,475	1,676
Elasticity	-0.279 (0.028)	-0.249 (0.184)	-0.407 (0.131)	-0.321 (0.043)	-0.267 (0.057)	-0.214 (0.107)
<u>Panel B: Germany</u>						
Tax coefficient	-0.843 (0.364)	-0.167 (0.176)	-0.306 (0.038)	-0.251 (0.046)	-1.792 (0.278)	-1.607 (0.672)
R squared	0.90	0.15	0.32	0.99	0.25	0.01
Obs	12,250	13,309	10,717	61,532	25,237	2,918
Elasticity	-0.189 (0.081)	-0.126 (0.133)	-0.306 (0.038)	-0.275 (0.057)	-0.375 (0.058)	-0.381 (0.160)
<u>Panel C: Sweden</u>						
Tax coefficient	-0.119 (0.025)	-0.043 (0.010)	-0.292 (0.164)	-0.164 (0.065)	-0.068 (0.026)	-0.374 (0.047)
R squared	0.80	0.12	0.37	0.99	0.31	0.05
Obs	12,774	13,120	8,942	30,602	15,772	1,856
Elasticity	-0.494 (0.105)	-0.329 (0.080)	-0.292 (0.164)	-0.176 (0.075)	-0.263 (0.102)	-1.120 (0.142)

Notes : Each panel reports a separate set of country-level specifications. The table reports the estimated coefficient on the tax variable and the elasticity of registrations to taxes, computed as in Table 2. Standard errors are in parentheses, clustered by segment-quarter. Specifications are the same as in column 1 of Table 2 except as noted. Observations in column 1 are in levels rather than first differences. In column 2, the dependent variable is the share in quarterly registrations. Column 3 replaces the tax with the log of the tax. Column 5 uses a nested logit specification in which the nests are market segments. The regression includes model fixed effects and the characteristics include fuel costs, engine size, weight, height, length, number of cylinders, and horsepower. The price and tax variables are in logs and the regression includes the log share of registrations in quarterly segment registrations. The price and log share of registrations are instrumented using the characteristics of other vehicles in the same segment sold under other brands as well as the characteristics in other segments sold under the same brand. In column 5 observations are by specification, year, and quarter. Observations in column 6 are by year and vehicle in Panels A and C and by half-year and vehicle in Panel B.

Table 4.

Potential Omitted Variables						
Specification	Add chars X quarter	Include 2005–2010	Control for other subsidies	Add vehicle X quarter	Drop obs near label cutoffs	Drop model- quarter interactions
	(1)	(2)	(3)	(4)	(5)	(6)
<u>Panel A: France</u>						
Tax coefficient	-0.503 (0.086)	-0.324 (0.098)	-0.320 (0.099)	-0.529 (0.052)	-0.440 (0.134)	-0.554 (0.060)
R squared	0.36	0.34	0.34	0.53	0.45	0.04
Obs	13,221	39,789	39,789	39,789	8,894	13,221
Elasticity	-0.374 (0.064)	-0.229 (0.069)	-0.225 (0.070)	-0.373 (0.036)	-0.398 (0.121)	-0.412 (0.045)
<u>Panel B: Germany</u>						
Tax coefficient	-3.159 (0.735)	-3.201 (1.006)		-0.721 (0.460)	-1.532 (0.649)	-1.454 (0.336)
R squared	0.33	0.32		0.52	0.36	0.01
Obs	10,717	55,465		55,465	7,345	10,717
Elasticity	-0.703 (0.164)	-0.686 (0.215)		-0.154 (0.098)	-0.362 (0.153)	-0.323 (0.075)
<u>Panel C: Sweden</u>						
Tax coefficient	-0.139 (0.055)	-0.134 (0.053)	-0.134 (0.053)	-0.009 (0.033)	-0.084 (0.042)	-0.108 (0.019)
R squared	0.39	0.36	0.36	0.62	0.43	0.01
Obs	8,942	26,551	26,551	26,551	6,661	8,942
Elasticity	-0.564 (0.223)	-0.431 (0.171)	-0.431 (0.171)	-0.028 (0.108)	-0.345 (0.171)	-0.441 (0.079)

Notes : Each panel reports a separate set of country-level specifications. The table reports the estimated coefficient on the tax variable and the elasticity of registrations to taxes, computed as in Table 2. Standard errors are in parentheses, clustered by segment-quarter. Specifications are the same as in column 1 of Table 2 except as noted. Column 1 includes interactions of fuel costs, engine size, weight, height, length, number of cylinders, and horsepower with year and quarter. Column 2 is the same as column 1 except that it includes all observations in the years 2005–2010. Panel A in column 3 adds to column 2 the amount of subsidy offered under France's vehicle retirement program. Panel C adds the subsidy offered in Sweden's green car rebate. Column 4 includes vehicle by quarter-of-year interactions. Column 5 drops observations within 5 g CO₂/km of the color cutoffs on new vehicle labels. Column 6 omits the model-quarter interactions.

Table 5.

<u>Cutoff (g CO₂/km)</u>	<u>Share of total registrations</u>	<u>Change in tax or subsidy (euros)</u>
120	0.042	500
130	0.019	200
160	0.026	200
165	0.014	550
200	0.002	850
250	0.002	1000

Notes : The first column shows the share in total registrations accounted for by vehicles within 2 g CO₂/km of the indicated cutoff. The second column shows the increase in tax or decrease in subsidy for vehicles above the cutoff compared to vehicles below the cutoff.

Table 6.

Effect of French Taxes on Vehicle Emissions Rates								
Specification	Include 2009	Include 2005–2010	Add tax change interactions	Include French brands	Add interactions with production start	(2) at model Level	(4) at model Level	(5) at model Level
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Above cutoff	0.0027 (0.0029)	-0.0006 (0.0009)	-0.0018 (0.0019)	-0.0023 (0.0012)	-0.0006 (0.0010)	0.0003 (0.0067)	0.0154 (0.0046)	0.0034 (0.0065)
Above X 2009		-0.0062 (0.0035)	0.0031 (0.0065)	-0.0024 (0.0047)	-0.0061 (0.0037)	-0.0061 (0.0149)	-0.0625 (0.0256)	-0.0165 (0.0155)
Tax X above X 2009			-0.0225 (0.0160)					
Production X above X 2009					-0.0072 (0.0074)			0.0781 (0.0293)
R squared	0.00	0.07	0.07	0.12	0.07	0.27	0.68	0.34
Obs	1,048	5,019	5,019	1,366	4,922	201	25	195

Notes : The table reports estimates of equation (2). Standard errors in parentheses allow for heteroskedasticity. Observations in columns 1–5 are by specification and year, and observations in columns 6–8 are by model and year. All observations are weighted by registrations in the previous year for the corresponding trim or model. Column 1 includes observations for 2009 and columns 2–8 include 2005–2010. All columns include observations within 2 g CO₂/km of the emissions rate cutoffs in the previous year. The dependent variable is the change in log CO₂ emissions rate across years. Above cutoff is a dummy variable equal to one for specifications or models that are above the tax cutoff in 2007. Above X 2009 is the interaction of the above cutoff dummy with a dummy variable equal to one in 2009 and 2010. Tax X Above X 2009 is the interaction of Above X 2009 with Tax, which is defined as the difference between the tax above the cutoff and the tax below the cutoff, divided by 1,000. Besides the reported coefficients, columns 2–8 include a set of year dummies and column 3 includes Tax, the interaction of Tax with Above Cutoff, and the interaction of the post-2009 dummy with Tax. Columns 4 and 7 include specifications or models sold by Citroen, Peugeot, and Renault. Production is a dummy variable equal to one if the corresponding model began production in 2009 or 2010 according to CSM production data. Columns 5 and 8 add to columns 2 and 6 Production, the interaction of Production with the post-2009 dummy, the interaction of Production with Above Cutoff, and the interaction of Production with Above X 2009.

Table 7.

France: Effect of Reducing Emissions Rates on Manufacturers' Profits						
	(1)	(2)	(3)	(4)	(5)	(6)
	<u>All brands</u>	<u>Lowest emissions rate quintile</u>	<u>Second quintile</u>	<u>Third quintile</u>	<u>Fourth quintile</u>	<u>Highest emissions rate quintile</u>
<u>Panel A: Summary Statistics</u>						
Market share		0.249	0.387	0.175	0.047	0.091
2007 CO ₂ emissions rate	148.4	138.0	141.3	149.2	162.8	176.1
Largest brands		Peugeot, Ford	Renault, Citroen	VW, Toyota	Nissan, Suzuki	Mercedes, BMW
<u>Panel B: Coefficient Estimates</u>						
Tax coefficient	-0.561 (0.113)	-0.672 (0.134)	-0.526 (0.258)	-0.881 (0.155)	-0.905 (0.113)	-0.252 (0.057)
<u>Panel C: Simulation Results for Reduction of 5 g CO₂/km</u>						
Purchase tax (euros/g CO ₂ /km)	10.17	17.13	20.29	8.94	5.46	13.33
Reduction in profits (euros/car)	24.1	40.7	48.5	21.2	13.3	32.6

Notes : The table reports the change in profits and emissions rate from imposing a purchase tax of 10 euros per g CO₂/km. Column 1 uses the full sample and the estimated tax coefficient from column 1 of Table 2. Columns 2–6 report results separately for five categories of brands, where brands are assigned categories based on their average CO₂ emissions rates in 2007. Panel A reports summary statistics for the categories and lists the two largest brands in each category. Panel B reports the estimated tax coefficient for the full sample (column 1) and for each category (columns 2–6). The coefficients in columns 2–6 are estimated in a single regression that includes interactions of the tax variable with a set of dummy variables for each category. Panel C reports the tax and change in profits from reducing emissions rates by 5 g CO₂/km for the full sample and for each category, and the linear purchase tax that achieves such a reduction.

Table 8.

Germany and Sweden: Effect of Reducing Emissions Rates on
Manufacturers' Profits

	(1)	(2)
	<u>Germany</u>	<u>Sweden</u>
	<u>Panel A: Coefficient estimates</u>	
Tax coefficient	-1.447 (0.334)	-0.060 (0.026)
	<u>Panel B: Reducing emissions rates by 5 g CO₂/km</u>	
Circulation tax (euros/g CO ₂ /km)	4.06	62.50
Reduction in annual profits (euros/car/year)	9.8	154.2

Notes : The table reports the change in profits and emissions rates for Germany and Sweden. Panel A reports the coefficient estimates from Table 2, and Panel B reports the results of the simulations, which are performed as in column 1 of Table 6, except that a linear circulation tax is used rather than a linear purchase tax.