

KU LEUVEN

CENTER FOR ECONOMIC STUDIES

DISCUSSION PAPER SERIES
DPS14.34

DECEMBER 2014



Consumer valuation of fuel costs and the effectiveness of tax policy: Evidence from the European car market

Laura GRIGOLON, Mathias REYNAERT and Frank VERBOVEN

Econometrics

Faculty of Economics
And Business



Consumer valuation of fuel costs and the effectiveness of tax policy: Evidence from the European car market

Laura Grigolon, Mathias Reynaert, Frank Verboven*

November 2014

Abstract

To what extent do car buyers undervalue future fuel costs, and what does this imply for the effectiveness of alternative tax policies? To address both questions, we show it is crucial to account for consumer heterogeneity in mileage and other dimensions. We use detailed product-level data for a long panel of European countries, and exploit variation in fuel prices by engine type. We find there is only modest undervaluation of fuel costs. As a consequence, fuel taxes are unambiguously more effective in reducing fuel usage than product taxes based on fuel economy, because fuel taxes better target high mileage consumers.

**Laura Grigolon*: McMaster University. Email: Laura.Grigolon@mcmaster.ca. *Mathias Reynaert*: University of Leuven and University of Antwerp. Email: Mathias.Reynaert@kuleuven.be *Frank Verboven*: University of Leuven and C.E.P.R. (London). Email: Frank.Verboven@kuleuven.be. We thank Bruno De Borger, Jan Bouckaert, James Sallee, Jo Van Biesebroeck and seminar participants at several seminars and conferences for helpful comments. We also gratefully acknowledge financial support from the Science Foundation–Flanders (FWO) and the KU Leuven Research Council (Program Financing grant).

1 Introduction

Governments are using a variety of policies to reduce CO₂ emissions from passenger cars. A central question in this debate is whether it is preferable to focus on fuel taxes or on policies that encourage the sales of fuel efficient cars, such as standards or product taxes based on the cars' fuel economy. Fuel taxes are often favored because they directly influence the car usage decision and hence the amount of CO₂ emissions. Furthermore, they may also indirectly influence the car purchase decision if consumers take into account their expected future fuel costs. However, if driving behavior is rather inelastic and if consumers are myopic and ignore future fuel cost savings when purchasing a car, then it may be more effective to directly influence the car purchase decision through standards or through up-front product taxes on cars with low fuel economy (possibly combined with product subsidies on cars with high fuel economy). In sum, fuel taxes may be more effective because they reduce car usage when driving behavior is not perfectly inelastic, while product taxes may be more effective because they can stimulate consumers to buy more fuel efficient cars if there is an investment inefficiency because of consumer myopia (see Allcott and Greenstone (2012) for a recent detailed review).

In this paper we contribute to this debate in two steps. We first ask whether consumers undervalue or correctly value the discounted future fuel costs when purchasing a new car. We subsequently ask what this implies for the relative effectiveness of fuel taxes versus product taxes based on the cars' fuel economy. We define the effectiveness of both taxes in terms of their reduction in total fuel usage (where the taxes are revenue equivalent).¹ Our main contribution is to empirically demonstrate the crucial importance of accounting for different sources of consumer heterogeneity, in particular mileage heterogeneity. This not only avoids biased parameter estimates of consumers' valuation of future fuel costs. It also accounts for another potentially important advantage of fuel taxes over product taxes: fuel taxes may not only reduce usage, but they also target the right consumers (with a high mileage) to purchase the most fuel efficient cars. This aspect of fuel taxes has not been empirically analyzed before.

To address these questions, we build on the aggregate random coefficients logit demand model of Berry, Levinsohn and Pakes (1995), and introduce a specification that accounts for heterogeneous responses to fuel costs because consumers may differ in their annual mileage. We conservatively assume that driving behavior (mileage) is perfectly inelastic with respect

¹Our focus on effectiveness has the advantage of simplicity, and is of direct interest to policy makers interested in reducing fuel consumption (and driving) in a revenue-neutral way. With additional information about the external costs of fuel consumption and driving, and about the nature of competition in the market, one could also use our estimates to conduct a more complete welfare analysis for the effects of both taxes.

to the fuel price. We show that the relative effectiveness of fuel taxes versus fuel economy-based product taxes depends on two key parameters: consumers' valuation of future fuel costs and mileage heterogeneity.

To estimate the demand model, we use detailed data at the level of the car model and engine variant for a panel of seven European countries during 1998-2011. The institutional features of the European car market provides a new way to identify consumer responses to fuel costs, because for most car models consumers can choose between two engine types: gasoline and diesel. Diesel cars are typically more expensive (by on average about 30%), but they also involve much lower fuel costs because of a higher fuel economy (by about 20%) and a lower diesel fuel price (by about 20 cents per liter, depending on the specific country's fuel tax policy). The consumers' intertemporal choice problem is correspondingly simpler in Europe than in the U.S., since consumers can trade off the higher initial purchase price of a diesel engine against the future fuel cost savings of an otherwise nearly identical car. To empirically identify consumer responses to fuel costs, we exploit fuel price variation by fuel type (i.e. gasoline or diesel) across countries and over time, interacted with variation in fuel economy across products. This complements other identification approaches, mainly for the U.S. market, where diesel cars are practically absent: these studies typically can rely on rich time-series variation in the price of gasoline, interacted with fuel economy variation across products (e.g. Klier and Linn (2010), Li, Timmins and von Haefen (2009), Allcott and Wozny (2012), Busse, Knittel and Zettelmeyer (2013), Langer and Miller (2013)).²

Regarding the valuation of future fuel costs, we find evidence of at most only modest undervaluation: for one euro saving in discounted future fuel costs, consumers are willing to pay €0.87 in the form of a higher initial purchase price. The 95% confidence interval of this willingness to pay is between €0.71 and €1.04, so that modest undervaluation or correct valuation cannot be rejected. To obtain this estimate, it is not only important to account for mileage heterogeneity to avoid a sorting bias towards undervaluation (Bento, Li and Roth, 2012). We demonstrate that it is also necessary to account for heterogeneity in the valuation of other car characteristics, since otherwise we would estimate significant overvaluation of discounted fuel costs (the willingness to pay would be equal to €1.27, with a confidence interval between €1.12 and €1.42).

Regarding the effectiveness of fuel taxes versus product taxes (based on a car's fuel economy), we again find that it is crucial to account for mileage heterogeneity. Without accounting for this, a fuel tax is less effective in reducing total fuel usage than a revenue-equivalent product tax on fuel economy, because of the modest undervaluation of future fuel

²The time-series variation in gasoline prices tends to be larger in the U.S. than in Europe, and the car sales and price information is often observed at a higher frequency (monthly instead of annually).

costs. Accounting for mileage heterogeneity reverses this result. In this case, a fuel tax turns out to be more effective in reducing total fuel usage because it specifically targets the high mileage consumers to substitute to cars with a higher fuel economy. In sum, we establish that a fuel tax is more effective than a product tax in shifting demand to more fuel efficient cars, since the undervaluation effect is dominated by the mileage heterogeneity effect. This conclusion is obtained under the assumption that driving behavior is perfectly inelastic. If we would allow driving behavior to depend on fuel prices, our results would be strengthened, since in this case a fuel tax would also induce consumers to reduce their car usage, conditional on the car purchase.

To further explore the effectiveness of fuel taxes under mileage heterogeneity, we also consider the impact of separately raising the price of diesel fuel or gasoline fuel (rather than simultaneously raising both). This is of broader interest, since countries are increasingly using fuel-specific tax policies to encourage the adoption of cars with alternative fuels (such as methanol). We find that the role of mileage heterogeneity becomes even more important under such fuel-specific tax policies. For example, only raising the price of gasoline fuel implies a large shift to diesel cars with higher fuel economy, and this shift is especially by high mileage consumers. As a result, this policy is especially effective in reducing total fuel usage and hence CO₂ emissions (although this can come at the expense of creating other environmental costs from diesel cars).³ As a final examination of the implications of our model, we consider the effects of harmonizing diesel fuel prices and diesel fuel economy to the level of gasoline. We find that both factors explain more than half of the diesel market share in Europe, though less so in the most recent years.

Our paper relates to several strands in the literature. First, our finding of at most only modest undervaluation of future fuel costs contributes to a long empirical debate since Hausman (1979). He estimated consumers' implicit interest rates in their intertemporal trade-off between paying a higher initial purchase price for air conditioners in exchange for future energy cost savings. Most work on the automobile market appears to find mixed evidence for the degree of undervaluation of future fuel costs; see Greene (2010) and Helfand and Wolverton (2011) for recent reviews. Allcott and Wozny (2012) find evidence of moderate undervaluation, while Busse, Knittel and Zettelmeyer (2013) find more or less correct valuation of future fuel costs relative to the initial purchase price of a car. Bento, Li and Roth (2012) show, both analytically and through simulations, that a failure to account for

³Miravete, Moral and Thurk (2014) stress these other (non-CO₂ related) environmental costs of diesel cars, such as raising local air pollution. They study how the lax European tax policy against these other pollution effects of diesel cars may have protected the domestic European producers against foreign competition. They do not explicitly model mileage heterogeneity, as it is less relevant for their purposes, and focus on the Spanish market.

consumer heterogeneity in willingness to pay for fuel costs is responsible for a sorting bias towards finding undervaluation, but they do not provide an empirical analysis to quantify its importance. Our paper incorporates mileage heterogeneity to avoid this bias, and at the same time it also incorporates other sources of consumer heterogeneity to avoid a reverse sorting bias. To our knowledge, no other work has systematically incorporated this to investigate whether consumers undervalue future fuel costs.⁴

Second, we contribute to the literature on the relative effectiveness of energy taxes versus product taxes and standards to reduce total energy usage and CO₂ emissions. Allcott and Wozny (2012) provide a detailed review on the relative effectiveness of both policy instruments. On the one hand, energy taxes directly affect usage, so they can be effective in reducing energy consumption if usage is sufficiently elastic. On the other hand, a product tax on the energy-inefficient product (or a subsidy on the energy-efficient product in their discussion) can better encourage the demand for energy-efficient products if consumers undervalue future energy cost savings, thereby reducing an investment inefficiency. Our contribution to this debate is to show that even if usage is inelastic, an energy tax can be more effective than a product tax. The reason is that if consumers are heterogeneous in their annual mileage, an energy tax better targets the high mileage consumers than a product tax.⁵ In a related theoretical paper, Allcott, Mullainathan and Taubinsky (2014) compare the combined effects of an energy tax and a product tax. They also stress the role of consumer heterogeneity, but they focus on heterogeneity in the extent of undervaluation (or “attentiveness”). For an empirical analysis, it is natural to start with mileage heterogeneity since it is directly observed. Nevertheless, in future work it would be interesting to extend our work to allow for heterogeneity in attentiveness.

Third, there is a large literature on estimating demand systems for automobiles. Berry, Levinsohn and Pakes (1995), Petrin (2002) and others have shown how to make use of aggregate sales data to estimate rich substitution patterns between differentiated cars. While this work has often included a random taste coefficient for “miles per dollar”, we show how to relate this random coefficient more explicitly to consumer mileage heterogeneity. This makes

⁴Verboven (2002) incorporates mileage heterogeneity to explain the consumers’ decision to buy a gasoline or diesel car and obtains estimates on their implicit interest rates when trading off the higher purchase price of a diesel engine against the future fuel cost savings. However, he focuses only on the decision to buy a gasoline or diesel engine, conditional on purchasing a certain car model. Although this avoids biases in estimating implicit interest rates, the framework does not allow to assess how taxes shift consumers to purchase other cars in policy counterfactuals.

⁵Our finding that an energy tax can be more effective than a product tax does not mean that a product tax cannot be effective in itself. In an interesting recent paper, Klier and Linn (2015) use data from three EU countries and find that product taxes can considerably shift sales towards cars with lower emissions. Our own results confirm this, but in addition we show that energy taxes are even more effective because they better target consumers with different mileage.

it possible to use the aggregate demand model to address environmental questions such as the effects of tax policies on total fuel usage, whereas previous applications with aggregate data could only assess a more limited set of effects.⁶ There is already an important empirical literature that uses micro-level data to thoroughly investigate the effects of environmental policies on both car purchase and car usage decisions (e.g. Goldberg (1998), Bento, Goulder, Jacobsen and von Haefen (2009), Gillingham (2012), and D’Haultfoeuille, Givord and Boutin (2014)). However, applications with micro-data are typically limited to a single country and a short time period. Our aggregate demand model with mileage heterogeneity thus considerably broadens the scope of applications, since aggregate data can be analyzed for a large set of countries and a longer time period.

The remainder of the paper is organized as follows. Section 2 develops the demand model, incorporating the consumers’ intertemporal trade-off between the initial purchase price of the car and future fuel costs, and highlighting the key parameters that drive consumer responses to fuel taxes versus product taxes: a future valuation parameter and mileage heterogeneity. Section 3 describes the data and section 4 provides details on the empirical estimation strategy and the identification issues. Section 5 presents and discusses the demand parameter estimates and implications for consumers’ valuation of future fuel costs. Section 6 presents policy counterfactuals, in particular a comparison between fuel tax and product taxes based on fuel-economy. Section 7 concludes.

2 The model

When consumers decide to purchase a new car, they face the intertemporal trade-off between the initial purchase price of the car and the expected future fuel costs. Forward-looking consumers have a high willingness to pay for fuel efficient cars that save on fuel costs, whereas myopic consumers have a low willingness to pay for such cars. In subsection 2.1 we model this intertemporal choice problem. We model the consumers’ discounted expected future fuel costs in the aggregate random coefficients logit demand model of Berry, Levinsohn and Pakes (1995). Our model accounts for heterogeneous responses because consumers may differ in their annual mileage and hence in their expected future fuel costs. In subsection

⁶For example, Adamou, Clerides and Zachariadis (2014) estimate an aggregate demand system to assess the effects of feebates in Germany (combination of car subsidy for fuel efficient cars with tax for fuel inefficient cars). They can look at the effects on demand and consumer surplus, but not at the effects on fuel usage since they do not explicitly model consumer mileage. Huse and Lucinda (2014) consider the effects of a Swedish subsidy program to cars with sufficiently low CO₂ emissions (where cars are more likely to be eligible if they run on alternative fuels). They assess the effects on demand and emissions, but without accounting for mileage heterogeneity in the demand model.

2.2 we use this demand model to discuss the different impact of two alternative taxes: a fuel tax and a product tax on the cars' fuel economy. This will serve as the basis to motivate our demand specification and subsequent policy counterfactuals.

2.1 Demand

There are T markets, defined as country/year combinations, with I_t potential consumers in each market t . Consumers are assumed to purchase a car only in the market where they are located. To simplify notation, we suppress the market subscript t in this section.

We define a car as a combination of a baseline model j and engine variant k . Consumer i may either choose a car model j with engine variant k , or decide not to buy a car and consume the outside good 0. The decision to buy a car affects the intertemporal budget constraint in two ways. First, consumers pay a capital cost, the initial purchase price p_{jk} . Second, they pay the present discounted value of expected future fuel costs G_{ijk} . The conditional indirect utility of consumer i for car model j and engine variant k is

$$u_{ijk} = x_{jk}\beta_i^x - \alpha_i(p_{jk} + \gamma G_{ijk}) + \xi_{jk} + \varepsilon_{ijk}, \quad (1)$$

where x_{jk} is a vector of observed car and engine characteristics and ξ_{jk} is an unobserved product characteristic. The vector β_i^x captures individual-specific valuations for the product characteristics, α_i is the marginal utility of income, and ε_{ijk} is a remaining individual-specific valuation for car jk , modeled as an extreme value (logit) random variable. The utility of the outside good is normalized to $u_{i00} = \varepsilon_{i00}$. The parameter γ is Allcott and Wozny (2012)'s "attention weight" or "future valuation" parameter. If $\gamma = 1$, consumers correctly trade off the car's purchase price p_{jk} against the present discounted value of future fuel costs G_{ijk} . If $\gamma < 1$ consumers undervalue future payoffs, and if $\gamma > 1$ consumers overvalue the future payoffs.

Our main focus in (1) is on the specification of G_{ijk} , i.e. consumer i 's present discounted value of expected future fuel costs for model j with engine k . We allow fuel costs G_{ijk} to be consumer-specific because of heterogeneity in annual mileage. Furthermore, G_{ijk} depends on expected fuel prices, on the relevant time horizon and interest rate. More precisely, we define the present value of expected fuel costs over the car's lifetime as follows:

$$G_{ijk} = E \left[\sum_{s=0}^{S-1} (1+r)^{-s} \beta_i^m e_{jk} g_{ks} \right], \quad (2)$$

where (i) β_i^m is consumer i 's expected annual mileage, measured in km; (ii) e_{jk} is the inverse

of the fuel economy of car j with engine k , measured in liter per km (or “gallons per mile”); (iii) g_{ks} is the fuel price of engine type k (either gasoline or diesel fuel) at time s , measured in € per liter; (iv) r is the interest rate at which consumers discount future fuel costs and S is the time horizon over which consumers value fuel cost savings, which may be related to the car’s expected lifetime.⁷ We model annual mileage and expected future fuel costs in a way that makes the expression for G_{ijk} simple and identification transparent.

First, we allow annual mileage β_i^m to be heterogeneous across consumers using the empirical distribution of mileage. This addresses a sorting bias in aggregate demand models that would arise if high mileage consumers choose to buy more fuel efficient cars (as discussed in Bento, Li and Roth, 2012). It also enables us to assess how consumers with different mileage respond to alternative tax policies.⁸ Our specification assumes that annual mileage β_i^m is perfectly inelastic, i.e. independent of the fuel price. Empirical evidence shows that the driving behavior is indeed relatively inelastic; see e.g. Goldberg (1998), Hughes, Knittel and Sperling (2008), Bento, Goulder, Jacobsen and von Haefen (2009), and Gillingham (2012). If we would allow driving behavior to depend on fuel prices, our result that a fuel tax is more effective than a product tax would be strengthened.

Second, we assume that fuel prices for each engine type k , g_{ks} , follow a random walk, so the consumers’ expected fuel price at time s is equal to the current fuel price (at time $s = 0$): $E[g_{ks}] = g_k$. This assumption is consistent with recent findings of Anderson, Kellogg and Sallee (2013). Some studies have instead considered alternative models for expectations about future fuel prices (based on actual future prices or past prices). In a literature review Alquist, Kilian and Vigfusson (2011) conclude that these more complicated models do not appear to outperform models with expectations that are only based on current prices.

Under these assumptions, we can rewrite the present value of expected future fuel costs for consumer i buying car j with engine k as

$$G_{ijk} = \rho \beta_i^m e_{jk} g_k, \tag{3}$$

i.e. consumer i ’s annual mileage β_i^m (in km) times the inverse of fuel economy e_{jk} (in

⁷In practice, vehicle lifetime may not be fixed but may partly depend on mileage.

⁸Other work has accounted for mileage heterogeneity by using micro-level data, see e.g. Bento et al. (2009) and Gillingham (2012) for the U.S. market. Allcott and Wozny (2012) account for mileage heterogeneity in an aggregate nested logit model by measuring average annual mileage per car model. Their approach does not however allow for policy counterfactuals, since it does not account for the possibility that consumers with different mileages substitute to other cars in response to tax policies.

liter/km) times the current fuel price g_k (in €/liter)km times a capitalization coefficient ρ :

$$\rho \equiv \sum_{s=0}^{S-1} (1+r)^{-s} = \frac{1+r}{r} \left[1 - (1+r)^{-S} \right], \quad (4)$$

which converts the annual fuel cost $\beta_i^m e_{jk} g_k$ into a net present value. Intuitively, the capitalization coefficient ρ measures the extent to which consumers trade off the initial purchase price of a car against annual fuel costs, and lies in the interval $[1, S]$. If consumers are fully myopic ($r \rightarrow \infty$), then $\rho = 1$: consumers then give the same weight to current annual fuel costs as to the initial purchase price. In contrast, if consumers do not discount the future ($r \rightarrow 0$), then $\rho = S$. Consumers then weigh the current annual fuel costs by a factor S relative to the purchase price of the car: they count on a “pay-back time” S when investing in a car with a higher fuel economy.

We can substitute (3) into (1) to write consumer i 's conditional indirect utility for car model j and engine variant k as

$$u_{ijk} = x_{jk} \beta_i^x - \alpha_i (p_{jk} + \gamma \rho \beta_i^m e_{jk} g_k) + \xi_{jk} + \varepsilon_{ijk}. \quad (5)$$

This is close to a standard random coefficients utility specification, where $\beta_i = (\beta_i^x, \alpha_i, \beta_i^m)$ are the random coefficients for which means and (co)variances may be estimated. The future valuation parameter γ and the capitalization coefficient ρ are additional parameters, but they are not separately identified from the scale of β_i^m . We will therefore make use of prior information on the empirical distribution of mileage β_i^m in our empirical analysis, so that $\gamma\rho$ becomes identified from the scale of β_i^m .

One can then use (4), to interpret the estimate of $\gamma\rho$ in three different ways. First, as in Hausman (1979), one can retrieve the consumers' implicit interest rate r at which they discount the future, for a given value of the car's expected lifetime S and setting $\gamma = 1$. Second, one can retrieve the consumers' required pay-back time S , assuming they adopt a market interest rate r and again setting $\gamma = 1$. Third, as in Allcott and Wozny (2012), one can impose both a market interest rate r and an expected lifetime S to retrieve the attention weight parameter γ , measuring the extent to which consumers undervalue ($\gamma < 1$) or overvalue the future ($\gamma > 1$). We will focus on the third approach.

We complete the demand model by assuming that each consumer i chooses the model j with engine k that maximizes her utility out of all possible alternatives in the choice set (including the outside option). Furthermore, assume that the random coefficients $\beta_i = (\beta_i^x, \alpha_i, \beta_i^m)$ come from a distribution $F_\beta(\beta; \theta)$, where θ are means and (co)variance parameters to be estimated; assume also that β_i is independent of the individual- and product-

specific taste valuations ε_{ijk} , which come from the type I extreme value distribution. Under these assumptions, the predicted market share for model j with engine k is the probability that jk gives the highest utility:

$$s_{jk}(\xi; \rho, \theta) = \int_{\beta} \frac{\exp(x_{jk}\beta^x - \alpha(p_{jk} + \gamma\rho\beta^m e_{jk}g_k) + \xi_{jk})}{1 + \sum_{j'}^J \sum_{k'=1}^{K_{j'}} \exp(x_{j'k'}\beta^x - \alpha(p_{j'k'} + \gamma\rho\beta^m e_{j'k'}g_{k'}) + \xi_{j'k'})} dF_{\beta}(\beta; \theta). \quad (6)$$

Observed sales can then be set equal to the predicted market share times the number of potential consumers I , i.e. $q_{jk} = s_{jk}(\xi; \rho, \theta) I$. Following Berry, Levinsohn and Pakes (1995) and subsequent work, the market shares can be approximated by Monte Carlo simulation with R draws of $\beta_i = (\beta_i^x, \alpha_i, \beta_i^m)$ from the distribution $F_{\beta}(\beta; \theta)$.

2.2 Consumer responses to taxes

We will estimate the demand model to compare the effectiveness of two alternative taxes: a fuel tax and a product tax on the cars' fuel economy. We will do this comparison in detail in our policy counterfactuals after having estimated the demand model. In this section, we give an overview of the possible effects of both taxes on the composition of new car sales. This provides economic intuition and will highlight the specific features that need careful attention in our empirical specification.

A fuel tax t_k^G is a tax on gasoline and/or diesel fuel g_k , whereas the product tax t_k^E is a tax on the (inverse) fuel economy e_{jk} of a car. The product tax on fuel economy e_{jk} is equivalent to the commonly used product tax on a car's CO₂ emissions, since there is a proportional relationship between both. The two taxes affect a consumer's conditional indirect utility (5) as follows:

$$u_{ijk} = x_{jk}\beta_i^x - \alpha_i (p_{jk} + t_k^E e_{jk}) - \alpha_i \gamma \rho \beta_i^m e_{jk} (g_k + t_k^G) + \xi_{jk} + \varepsilon_{ijk}.$$

Previous work has stressed the relative advantages of both taxes. On the one hand, a fuel tax is preferable because it directly reduces usage and hence (pollution) externalities, as long as driving behavior is not perfectly inelastic. On the other hand, a product tax on the least fuel efficient cars can be preferable if consumers undervalue their future fuel costs when purchasing a durable good. Put differently, a fuel tax mainly serves to correct for externalities by affecting the driving decision, while a product tax on fuel economy mainly corrects for an investment inefficiency from consumer myopia. Allcott and Greenstone (2012) provide interesting further discussion of both effects (where they focus on a subsidy for the energy efficient product, instead of a tax on the energy inefficient product).

In our analysis we assume that utilization is perfectly inelastic, so we rule out the possibility that a fuel tax corrects for an externality by reducing utilization. We instead focus on another role of fuel taxes: in the presence of consumer mileage heterogeneity, a fuel tax especially targets the high mileage consumers. Hence, a fuel tax may be more effective than a product tax even if it does not directly affect utilization decisions.

To more precisely see the role of mileage heterogeneity and consumer myopia in the effectiveness of a fuel tax versus a product tax, it is useful to consider the case where the taxes do not depend on the fuel type k , i.e. $t_k^G = t^G$ and $t_k^E = t^E$ for both gasoline and diesel cars. Using $q_{jk} = s_{jk}(\xi; \rho, \theta) I$ and the expression (6) for $s_{jk}(\xi; \rho, \theta)$, Appendix A shows that a small increase of t^G and t^E has the following effect on the demand for product jk :

$$\begin{aligned}\frac{\partial q_{jk}}{\partial t^G} &= - \int_{\beta} \alpha_i \gamma \rho \beta_i^m s_{ijk} (e_{jk} - \bar{e}^i + s_{i0} \bar{e}^i) dF_{\beta}(\beta) I \\ \frac{\partial q_{jk}}{\partial t^E} &= - \int_{\beta} \alpha_i s_{ijk} (e_{jk} - \bar{e}^i + s_{i0} \bar{e}^i) dF_{\beta}(\beta) I\end{aligned}$$

where $\bar{e}^i = \sum_j \sum_k e_{jk} s_{ijk} / (1 - s_{i0})$ is the expected fuel economy over cars purchased by consumer i . We can make the following two observations.

First, the effect of both taxes depends on the sign of the term $e_{jk} - \bar{e}^i + s_{i0} \bar{e}^i$. If all cars would have the same fuel economy, i.e. $e_{jk} = e$ for all jk , then $e_{jk} - \bar{e}^i = 0$.⁹ Both taxes would then reduce the demand for all cars, proportional to the aggregate consumer responses to the outside good s_{i0} and the price sensitivity parameter α_i . The effect of the taxes is then similar to the effect of an industry-wide price increase. In contrast, if cars differ in their fuel economy, then the demand for some cars can increase despite the tax increase. This will be the case for cars with a sufficiently low inverse fuel economy, i.e. $e_{jk} < \bar{e}^i$, and when there is limited substitution to the outside good (s_{i0} small). In sum, both taxes have in common that they change the composition of new car sales from cars with a low fuel economy (high e_{jk}) to cars with a high fuel economy (low e_{jk}).

Second, the effect of the taxes differ because of two factors. On the one hand, a fuel tax t^G may be less effective than a product tax on fuel economy t^E if there is consumer myopia ($\gamma < 1$). This is the investment inefficiency referred to above. On the other hand, a fuel tax may be more effective because of mileage heterogeneity β_i^m : a fuel tax especially targets the high mileage consumers to substitute to more fuel efficient cars. In the special case where there is no consumer heterogeneity, it can be shown that a revenue-equivalent product tax t^E has exactly the same effect as a fuel tax t^G if and only if $\gamma = 1$. It has a stronger (weaker)

⁹When $e_{jk} = e$, we have $\bar{e}^i = e \sum_j \sum_k s_{ijk} / (1 - s_{i0}) = e$, so that indeed $e_{jk} - \bar{e}^i = 0$.

impact if and only if there is undervaluation (overvaluation).

In sum, the effectiveness of both taxes is an empirical question. The above discussion highlighted the importance of several factors: the price sensitivity parameter α_i and the extent substitution to the outside good s_{i0} determine the general effectiveness of both taxes (similar to their role in a price elasticity of industry demand). The future fuel valuation parameter γ and the extent of mileage heterogeneity β_i^m explain differences in their effectiveness. We will pay particular attention to these various features in our empirical specification.

3 Data

Our main dataset is a rich panel of data from the European car market, obtained from a market research firm (JATO). The dataset includes the sales, prices, and product characteristics for every new passenger car sold during 1998-2011 in seven European countries: Belgium, France, Germany, Italy, the Netherlands, Spain, and the UK. The data cover around 90% of the sales in the European Union.

The unit of observation is at the very detailed level of the car variant jk , i.e. the combination of a car model j equipped with engine k . The car model j is a brand/model/body type combination, e.g. “Volkswagen Golf hatchback”, whereas the engine k consists of the fuel engine type (gasoline or diesel), displacement and horsepower, e.g. “gasoline, 1,390cc, 59kW”. Our highly disaggregate definition of a car variant make it possible to capture all possible variation in fuel efficiency and engine performance. After excluding cars with extremely low sales (e.g. Bentley Arnage), we retain on average about 800 car variants per country/year, i.e. 180 models with on average 4.4 engines. This results in a panel dataset of approximately 80,000 observations (car variants/countries/years).

Sales are defined as new vehicle registrations. Prices are suggested retail prices, including VAT and registration taxes which differ across markets and engines (separately obtained from the European Automobile Manufacturers Association). As discussed below, we include a rich set of fixed effects for car models and countries/years to account for car-specific discounts and variation across countries and over time. Car characteristics include measures of vehicle size (curb weight, width, length and height), engine performance (horsepower and displacement) and (inverse) fuel economy (liter/100km). In addition, based on a brand’s perceived country of origin, we construct a dummy for whether a model is of foreign or domestic origin in each country.¹⁰

We supplement this dataset with information on fuel prices, the empirical distribution

¹⁰For example, the Volkswagen Golf is perceived as domestic by German consumers even though part of the production of Golf takes place in Spain.

of annual miles travelled and other country/year-level information. Gasoline and diesel fuel prices by year and country are from the Directorate General for Economic and Financial Affairs. Fuel prices mainly vary over time, rather than across countries. Figure 1 plots the average annual price of gasoline and diesel fuel during 1998–2011 (in real year 2000 Euros). Both gasoline and diesel fuel prices are increasing to reach peaks in 2008 and 2011. Diesel fuel is on average 16% less expensive than gasoline fuel, but the gap varies during the period. There is also variation in gasoline and diesel prices across countries.

The empirical distribution of annual miles travelled is from the 2007 UK National Travel Survey, a rich nationally representative survey of 20,000 individuals.¹¹ According to this survey, average annual mileage is 14,700 km/year. The distribution of mileage is skewed to the right: 20% of the population drives less than 7,000 km/year, 50% drives less than 10,200 km/year, 80% drives less than 18,000 km/year, while 10% drives more than 25,000 km/year. There is no such detailed information on the empirical mileage distribution in other countries of our dataset. Eurostat reports average annual mileages by country, using somewhat differing methodologies. These averages are in line with the most reliable average of the distribution reported in the UK National Travel Survey. We therefore assume that mileage in the other countries follows the same distribution as in the UK. Finally, we use information on GDP/capita in each country/year to scale car prices and annual fuel costs in the same units across countries, and we use population per country/year to construct the variable for the number of potential consumers.

Summary statistics Table 1 provides summary statistics (mean and standard deviation) for the variables included in our empirical demand model. As one may expect, characteristics that may vary across both models and engines (price, horsepower, annual fuel costs) show a higher dispersion around the mean than characteristics that only vary across models (size, height and the foreign dummy). Furthermore, several characteristics show considerable changes over the sample period: horsepower, fuel efficiency and the fraction of diesel cars all increased between 1998–2011, whereas prices (relative to income) remained fairly stable.

Table 2 provides more detailed summary statistics broken down between gasoline and diesel cars. This breakdown is informative, since the differences between both types of cars form a main source of variation to identify the willingness to pay for fuel cost savings. The share of new diesel cars increased from 31.7% in 1998 to 57.7% in 2011.¹² There were on

¹¹www.gov.uk/government/organisations/departement-for-transport/series/national-travel-survey-statistics

¹²This trend is common to all European countries, but there are also notable differences across countries (not shown in the table). Belgium, France and Spain reached the highest share of diesel cars in 2011 (respectively 75.2%, 69.0% and 69.0%), whereas the Netherlands had the lowest share (30.5%). The shares

average only 1.6 diesel engines per car in 1998 (compared with 2.5 gasoline engines per car), and this number increased to 2.5 diesel engines per car in 2011. Hence, while a considerable fraction of car models was not sold with a diesel engine in 1998, this was no longer the case in 2011. Diesel cars have 19% higher fuel efficiency than gasoline cars (on average 4.8 liter/100km for diesel cars versus 5.7 liter/100km for gasoline cars in 2011). Furthermore, as mentioned above, diesel fuel tends to be 16% less expensive than gasoline fuel costs. At the same time, a diesel car is on average 29% more expensive than its gasoline counterpart. This varying trade-off between the higher up-front purchase price for diesel cars and the higher expected fuel cost savings across car models/countries and years will be of primary importance to identify the extent to which consumers discount the future.

4 Estimation

We have a panel of T markets, defined as country-year combinations, to estimate the taste parameters of the market share system (6). We reintroduce the subscript t to refer to these markets. For each market t and each car model j with engine k , we observe the sales q_{jkt} , prices p_{jkt} , fuel economy e_{jkt} , a vector of other product characteristics x_{jkt} and the fuel price g_{kt} (gasoline or diesel fuel). The observed market shares are computed as the sales q_{jkt} divided by the number of potential consumers I_t , $s_{jkt} = q_{jkt}/I_t$, and these observed shares are set equal to the predicted shares as given by (6).

We first discuss the specification of the taste parameters. Next, we discuss the assumptions regarding the error term ξ_{jk} and the GMM estimator. Finally, we discuss specific computational aspects of the estimator.

Specification of the taste parameters The taste parameters to be estimated are $\gamma\rho$ and θ , where $\gamma\rho$ measures the extent to which consumers trade off the initial purchase price of a car against annual fuel costs, and θ is a vector of distributional parameters for the random coefficients $\beta_i = (\beta_i^x, \alpha_i, \beta_i^m)$ with distribution $F_\beta(\beta; \theta)$. Recall that β_i^x measures the individual-specific valuations for the product characteristics x_{jkt} , α_i is the marginal utility of income, and β_i^m is consumer i 's annual mileage. As discussed in section 2, $\gamma\rho$ is not separately identified from the scale of β_i^m . Furthermore, estimating a large number of distributional parameters θ , i.e. means and (co)variances, is computationally challenging, so we impose a number of restrictions.

of diesel cars in the other countries varied between 45.5% and 55.4%. These differences may stem from unobserved country-specific factors, such as taxes or fuel station networks. To account for this, we interact the diesel variation with country-specific fixed effects.

First, we assume that β_i^m follows the empirical distribution of mileage based on the information discussed in section 3. This avoids a restrictive functional form and also ensures identification of $\gamma\rho$. Second, following the previous literature, we assume β_i^x is normally distributed and we only estimate means and standard deviations of β_i^x , so we restrict their covariances to be equal to zero:

$$\beta_i^x = \bar{\beta}^x + \Sigma^x \nu_i^x,$$

where $\bar{\beta}^x$ are the mean valuations, Σ is a diagonal matrix with standard deviations σ^x on the diagonal, and ν_i^x are standard normal random variables. We nevertheless allow for a nonzero covariance between the intercept β_i^0 and β_i^m through a parameter σ^{0m} . This allows for the possibility that the high mileage consumers are also more likely to purchase a car than to purchase the outside good. A higher mileage decreases utility through higher fuel costs but can increase utility through a higher value on the interaction with the intercept. Depending on which effect dominates we expect to see different shares of the outside good for different mileages. Third, we specify α_i to be inversely proportional to income y_t in market t , so $\alpha_i = \alpha/y_t$.¹³

The specification of the conditional indirect utility of consumer i for car model j and engine variant k then becomes

$$u_{ijkt} = x_{jkt}\beta_i^x - \alpha p_{jkt}/y_t - \alpha\gamma\rho\beta_i^m e_{jkt}g_{kt}/y_t + \xi_{jkt} + \varepsilon_{ijkt}. \quad (7)$$

We will directly estimate and report $\alpha\gamma\rho$ as the “fuel cost parameter”, and then retrieve the future valuation parameter γ from dividing the estimate of $\alpha\gamma\rho$ by the estimate of the price parameter α and a value of the capitalization coefficient ρ (using the interest rate r and time horizon S). The vector of product characteristics x_{jkt} includes horsepower, size (width times length), height, foreign, and a diesel dummy variable interacted by country dummy variables. The latter captures valuation differences for diesel engines across countries, including unobserved differences in taxation or in fuel station networks for diesel cars. We estimate mean valuations β^x for all these variables, and in addition standard deviations σ^x for horsepower and a covariance parameter σ^{0m} . We also considered a more general specification with standard deviations for the other characteristics: this gives similar results,

¹³Similar to Berry, Levinsohn and Pakes (1995), this specification approximates a Cobb-Douglas specification $\alpha \ln(y_t - p_{jkt} - \gamma\rho\beta_i^m e_{jkt}g_{kt}) \approx \alpha y_t - \alpha\gamma p_{jkt}/y_t - \alpha\gamma\rho\beta_i^m e_{jkt}g_{kt}/y_t$, when capitalized car expenditures are small relative to capitalized income. It is particularly convenient in our setting with many countries with differing exchange rates, because prices and fuel expenditures are expressed in local prices relative to local income. We also considered a specification where α_i follows the empirical distribution of income, $\alpha_i = \alpha/y_i$, in a simpler specification with fewer other random coefficients. This produced similar results, but was computationally much slower, so we focus on the model where $\alpha_i = \alpha/y_t$ and a richer set of other random coefficients.

but less precise estimates for some of our main parameters of interest γ . This is consistent with earlier findings of Reynaert and Verboven (2014) on the difficulties in precisely estimating too many random coefficient parameters.

GMM estimator We exploit the panel nature of our dataset and specify the unobserved product characteristic as

$$\xi_{jkt} = \xi_j + \xi_t + \tilde{\xi}_{jkt}, \quad (8)$$

where ξ_j are fixed effects capturing time-invariant unobserved characteristics for a car model j , and ξ_t captures market fixed effects, modeled as country-specific fixed effects interacted with a time trend and a squared time trend. Since some models were introduced or eliminated within a year, we also include a set of fixed effects for the number of months for which the model was available in a country within a given year. The last term, $\tilde{\xi}_{jkt}$, is the residual error term capturing the remaining unobserved characteristics varying across models, engines and markets.

The error terms $\tilde{\xi}_{jkt}$ enter the market share system (6) in a highly nonlinear way. Following Berry (1994), Berry, Levinsohn and Pakes (1995) and the subsequent literature, we use a contraction mapping to invert the market share system and solve for the vector of error terms $\tilde{\xi}_t$ in each market t .

We account for the fact that price is an endogenous variable, which may be correlated with the error term $\tilde{\xi}_{jkt}$. The main identification assumption is that the error term is mean independent of the other product characteristics, $E(\tilde{\xi}_{jkt}|z_t) = 0$, where z_t is a matrix of all other product characteristics (including x_t , but also fuel economy e_t and fuel prices g_t). These conditional moment restrictions imply an infinite number of unconditional moment restrictions

$$E\left(h_{jt}(x_t)\tilde{\xi}_{jkt}\right) = 0$$

where $h_{jt}(x_t)$ is a vector of instruments formed by any function of the exogenous x_t . The GMM estimator for the complete vector of parameters is then the solution to

$$\min_{\phi} \tilde{\xi}(\phi)' h(z)' \Omega h(z) \tilde{\xi}(\phi),$$

where the vectors and matrices are stacked over all markets, ϕ is the vector of parameters ($\gamma\rho$ and θ) and Ω is a weighting matrix using first-step residuals to account for heteroskedasticity.¹⁴ To construct the instruments $h(z)$, Berry, Levinsohn and Pakes (1995) suggest to use

¹⁴To minimize the objective function, we concentrate out the linear parameters β and the market fixed effects ξ_t as discussed in Nevo (1999). Furthermore, following Baltagi (1995), we use a within transformation of the data to eliminate the car model fixed effects ξ_j .

the own-product characteristics, sums of the characteristics of other products of the same firm, and sums of characteristics of other firms' products. We use these instruments in a first stage to obtain initial parameter estimates and construct optimal instruments following Chamberlain (1987) and Berry, Levinsohn and Pakes (1999). In a second stage, we use these optimal instruments to obtain more efficient parameter estimates. Reynaert and Verboven (2014) provide detailed Monte Carlo evidence and an application with actual data to demonstrate that optimal instruments indeed greatly improve the efficiency of the estimator for this model.

Computational considerations We account for a variety of computational issues to which recent work has drawn attention (see for example Goldberg and Hellerstein (2013)'s checklist). First, we approximate the high-dimension market share integral using 500 draws of a quasi-random number sequence for each of the 98 market. The empirical distribution of mileage consists of 100 nodes, such that each mileage is interacted with 5 draws on the other dimensions. Second, we use a tight convergence level of $1e^{-12}$ for the contraction mapping to solve the market share system in the inner loop within the GMM objective function. Third, to minimize the GMM objective function we use a state-of-the-art optimization algorithm (the Interior/Direct algorithm in Knitro), provide analytical derivatives and set a strict tolerance level at $1e^{-6}$. Fourth, we use a set of 50 starting values to search for a global minimum, and verify the solution by checking the first-order and second-order conditions. Finally, as discussed above, we use optimal instruments in a second stage, which greatly improves the efficiency of the estimator, in particular for the standard deviations of the random coefficients.

5 Empirical results

To assess how ignoring heterogeneity of consumer preferences may bias the extent to which consumers trade off the car purchase price against future fuel cost savings, we consider three alternative models. The first model is a simple logit model which imposes $\beta_i^x = \beta^x$ and $\beta_i^m = \beta^m$ (i.e. the mean of the observed mileage distribution). The second model allows for heterogeneity in mileage and hence in the valuation of fuel costs, using the empirical mileage distribution β_i^m , while continuing to restrict $\beta_i^x = \beta^x$. The third model is our full random coefficients model which allows for both mileage heterogeneity and heterogeneity in the valuations of some other car characteristics. This model also allows mileage to be correlated with the utility for a car relative to the outside good (covariance parameter σ^{0m}).

We first discuss the parameter estimates, and then what these imply for consumers'

valuations of future fuel cost savings.

Parameter estimates Table 3 reports the parameter estimates for these three demand models. First consider the simple logit model. Price and annual fuel costs have the expected negative effect on utility, with $\alpha = -4.54$ and $\alpha\gamma\rho = -39.90$. This implies $\gamma\rho = 8.78$, which we discuss in detail below. The estimated mean valuations β^x for horsepower, size (length times width) and height are positive, while the mean valuation for foreign cars is negative. The valuation of a diesel engine (apart from fuel costs) differs across countries (not shown in table): diesel cars have a lower mean utility in Germany, the UK and especially in the Netherlands. This may be due to unobserved higher car taxes, less elaborated diesel fuel station network, the popularity of alternative fuels (such as LPG in the Netherlands) or other unobserved preference differences.

Now consider the random coefficients logit model with only mileage heterogeneity, and no heterogeneity for other characteristics. Price and fuel costs again have the expected negative effect on utility, but the magnitude of the fuel cost effect increases ($\alpha = -5.20$ and $\alpha\gamma\rho = -67.93$). The estimated mean valuations of the other product characteristics β^x all have the same signs and are similar in magnitude as in the simple logit model.

Finally, consider the full random coefficients logit with heterogeneity for other characteristics (σ^x) in addition to fuel costs. As in the other two models, price and fuel costs have a negative effect on utility. Similarly, the mean valuation for size and height is positive, while the mean valuation for foreign cars is negative. The mean valuation for horsepower now becomes negative. But there is substantial heterogeneity around this mean, implying that about 21% of consumers value horsepower positively. Finally, the covariance parameter σ^{0m} is positive. This means that high mileage consumers show a positive preference for buying cars relative to purchasing the outside good.¹⁵

Implications for valuation of future fuel cost savings What do our estimates imply for consumers' valuations of future fuel cost savings? To address this question, the starting point is the coefficient $\gamma\rho$, which converts annual fuel costs into the present discounted value. It is simply the ratio of the estimated fuel cost coefficient $\alpha\gamma\rho$ over the estimated price coefficient α . As discussed above, we can then use the expression of ρ given by (4) to

¹⁵Note that the standard deviations of the random coefficients are precisely estimated due to the use of optimal instruments. Our first-stage estimates, where we used sums of product characteristics of other products as an approximation for the optimal instruments, produced much less significant estimates, similar to what was found in Reynaert and Verboven (2014). As reported above, we also considered a richer specification with other random coefficients. This gave comparable results, though the estimates of α and $\alpha\gamma\rho$ are less precise.

draw conclusions about consumers' intertemporal preferences. More specifically, we follow Allcott and Wozny (2012), and assume a time horizon S and market interest rate r to obtain a value for ρ and retrieve the attention weight or future valuation parameter γ : this measures the extent to which consumers undervalue ($\gamma < 1$) or overvalue the future ($\gamma > 1$).¹⁶

The bottom panel of Table 3 compares the findings for the three different demand models. To compute γ , we set $r = 6\%$ as in Allcott and Wozny (2012) and assume a time horizon $S = 15$: this is at the higher end of Eurostat estimates of expected vehicle life between $S = 10$ and $S = 15$, which makes it more likely to find undervaluation of future payoffs.¹⁷

The logit model implies that $\gamma\rho = 8.78$, which is precisely estimated with a standard error of 0.55. Setting $\gamma = 1$ and $r = 0$, this can be interpreted as a required payback time of 8.78 years for an investment in a more fuel efficient car. Equivalently, using the above values of r and S , the implied attention weight parameter is $\gamma = 0.85$, which implies a moderate undervaluation of future fuel cost savings.

In contrast, the random coefficients logit model with only mileage heterogeneity implies a considerably larger coefficient of $\gamma\rho = 13.05$ and a corresponding attention weight parameter of $\gamma = 1.27$. This implies that consumers overvalue future fuel cost savings according to this model. This finding is consistent with Bento, Li and Roth (2012): they show, both analytically and through simulation, that ignoring heterogeneous responses to fuel cost leads to a sorting bias and a corresponding underestimate of consumers' willingness to pay for future fuel cost savings. Intuitively, high mileage consumers sort into cars with better fuel economy, and ignoring this amounts to omitting a variable that is positively correlated with fuel costs.

Finally, the full model, which also allows for heterogeneity in the valuation of other product characteristics, results in a coefficient of $\gamma\rho = 9.03$, so that the attention weight parameter is $\gamma = 0.88$. This again implies some moderate (but statistically insignificant) undervaluation of future fuel cost savings, as in the simple logit model. Hence, omitting some dimensions of heterogeneity turns out to introduce another bias in the estimation of consumer sensitivity to fuel costs: fuel economy is highly correlated with other characteristics such as horsepower, for which preferences are also heterogeneous. The results show that not accounting for heterogeneity on other characteristics results in an additional bias, which goes in the reverse direction of Bento, Li and Roth (2012)'s sorting bias.

¹⁶As discussed above, we could equivalently set $\gamma = 1$ and either assume a time horizon S to retrieve the consumers' implicit interest rate r , as in Hausman (1979); or assume an interest rate r to retrieve the consumers' required payback time S .

¹⁷Allcott and Wozny (2012) use a richer model for vehicle life, allowing for an annual depreciation rate. In principle, we could also adopt such an approach, but this would make the expressions less transparent and make the model more difficult to estimate. It is in any case not clear whether consumers are sufficiently sophisticated to incorporate all this information.

In all models, the consumer surplus losses per vehicle from misoptimization are small: €50 from undervaluation in the first model, €112 from overvaluation in the second model, and €22 from undervaluation in the third model. Note that this type of low undervaluation may be interpreted as rational inattention from the consumers’ point of view, if it is costly to compute future fuel savings exactly, see Sallee (2014).

In sum, to obtain reliable estimates on consumers’ valuations of future fuel cost savings, it is important to account for both mileage heterogeneity and heterogeneity regarding other product characteristics than fuel costs. We find only moderate undervaluation with an attention weight of $\gamma = 0.88$ for European consumers. This is similar to Busse, Knittel and Zettelmeyer (2013), and higher than the $\gamma = 0.76$ obtained by Allcott and Wozny (2012). As such, the European car market shows only very limited evidence for an “energy paradox” according to which consumers are puzzlingly slow in investing in energy efficiency (Jaffe and Stavins, 1994). A possible explanation is the rather high fuel prices in Europe, which makes it more profitable for consumers to pay attention to future fuel costs. Furthermore, almost every car comes with either a gasoline engine and a diesel engine (with a higher initial purchase price, but also lower future fuel costs): this makes it easier to compare products that are otherwise nearly identical (see also Verboven, 2002). We now turn to the question what these findings imply for the effectiveness of alternative tax policies.

6 Policy counterfactuals

As shown in the previous section, to obtain reliable estimates of consumers’ valuations of future fuel cost savings it is important to account for consumer heterogeneity, both regarding their mileage and other dimensions. What do our findings imply for the effectiveness of alternative government policies?

Overview: fuel tax versus product tax There is a large literature on the environmental effects of alternative tax and standards policies in the automobile market, and we do not aim to provide a comprehensive analysis on how these various policies may affect all different externalities.¹⁸ Instead, we focus on the effectiveness of two tax policies, in terms of their reduction in total fuel usage: a fuel tax t_k^G (on gasoline and/or diesel fuel) and a product tax t_k^E (on a car’s inverse fuel economy e_{jk}). Both taxes are representative for a broader group of policies that governments can use in order to reduce externalities. A fuel tax is equivalent

¹⁸Parry, Walls and Harrington (2007) and Anderson, Parry, Sallee and Fischer (2011) provide a comprehensive overviews of different (pollution and environmental) externalities and different policies to address them.

to a carbon tax and can be seen as a Pigouvian tax that directly prices the externality. A product tax changes the relative prices of products with different efficiencies and will have similar effects as subsidies for efficient vehicles, feebates or fuel economy standards.

In section 2, we already discussed how the effectiveness of both policies depends on various factors. The price sensitivity parameter α_i and the extent of substitution to the outside good s_{i0} determine the general effectiveness of both taxes (similar to their role in a price elasticity of industry demand). The future valuation parameter γ and the extent of mileage heterogeneity β_i^m explain differences in effectiveness between the fuel tax and the product tax. Without mileage heterogeneity ($\beta_i^m = \bar{\beta}^m$ for all consumers i), a product tax is more effective than a fuel tax if and only if $\gamma < 1$, because consumers respond more to an immediate incentive. However, if there is mileage heterogeneity, a fuel tax may be more effective, because it better targets the high mileage consumers to substitute to cars with better fuel economy.¹⁹

We will first compare the effects of both policies on the composition of new sales, i.e. on the market shares by fuel economy quartile. Next, we compare the effects of both taxes on the average fuel economy and total fuel usage of sold cars. Finally, we discuss the effects of discriminatory taxes that are specific to the fuel type (gasoline versus diesel in our case), and where the role of fuel taxes to target heterogeneous consumers is even more important.

In most counterfactuals, we will consider the effects of a fuel tax of 50c per liter. This is of a roughly similar order of magnitude as the fuel tax of 1\$ per gallon considered in several US studies. We will then compare this with the effects of a revenue-equivalent product tax per unit of fuel economy e_{jk} (in liter/100km). A product tax is revenue equivalent if it gives the same revenues as the capitalized expected revenues from a 50c tax per unit of fuel (in liter). We mainly focus on the results for Germany in 2011, the largest country in the most recent year of our dataset. Computational details are given in Appendix A.

Effects on the composition of new car sales We first discuss how a fuel tax and a revenue-equivalent product tax affect the composition of new car sales. We specifically compute the effect of both taxes on the sales of every car jk , and then aggregate these effects to market share effects of four fuel economy quartiles. This is similar to what Busse, Knittel and Zettelmeyer (2013) report, based on an entirely different identification approach. Since our approach is based on the estimates of a structural demand model, we can report interesting additional information. First, Busse et al. only show the impact of a fuel tax on

¹⁹ As discussed in section 2, a fuel tax may also be more effective because it can influence driving behavior if consumer mileage is not perfectly inelastic. We rule this out to focus on the role of mileage heterogeneity in the effectiveness of a fuel tax.

the market shares by fuel economy quartile, whereas we can compare this with the impact of a (revenue-equivalent) product tax. Second, we can compute the market share effects under both the actual estimate of the consumers' future valuation parameter ($\gamma = 0.87$), and under alternative scenarios with full forward looking behavior ($\gamma = 1$) and strong consumer myopia ($\gamma = 0.5$). This detailed comparison is useful, since it introduces some key intuition behind the subsequent effects on fuel economy, total fuel usage and consumer welfare.

Table 4 shows the market share effects by fuel economy quartile, based on the estimates of the full model with heterogeneity in both mileage and in the valuation of other product characteristics. For completeness, Appendix B also shows the results for the simple logit model without mileage heterogeneity, and we briefly comment on this first. As expected from our earlier analysis in section 2, in the absence of mileage heterogeneity a fuel tax of 50c per liter has exactly the same impact on market shares as a revenue-equivalent product tax when consumers are fully forward looking ($\gamma = 1$), while it is less effective when consumers undervalue the future ($\gamma < 1$).²⁰

We now discuss the market share effects for the full model with consumer heterogeneity. According to the top panel of Table 4, under the actual future valuation parameter ($\gamma = 0.87$) a 50c fuel tax increase raises the market share of cars in the highest fuel economy quartile (lowest e_{jk}) by 12% (or 4.4 percentage points, from a market share of 36.6% to 41.0%). The market shares of the other fuel economy quartiles all drop. In particular, the market share of cars in the lowest fuel economy quartile (highest e_{jk}) drop by 16.1%. Hence, a fuel price increase implies a quite considerable change in the composition of car sales from the low to high fuel economy cars. If consumers would be fully forward looking ($\gamma = 1$), then the 50c fuel tax increase would have been only slightly more effective: it would raise the market share of the highest fuel economy quartile by 13.3% and reduce the market share of the lowest fuel economy quartile by 17.7%.²¹ Conversely, with strong consumer myopia ($\gamma = 0.5$) the fuel tax increase would have had a smaller, but still non-negligible impact: +7.3% for the highest fuel economy quartile and -10.1% for the lowest fuel economy quartile.

The bottom part compares these findings with those of a revenue-equivalent product tax increase. This amounts to a tax per unit of fuel economy of €767 per unit of e_{jk} (so this would amount to a product tax of €3,835 for a car that consumes 5 liter per 100km). As expected, the product tax does not become less effective when the future valuation

²⁰In the logit model, this revenue-neutral product tax is equal to €890 per liter/100km. For example, the product tax on a car that consumes 5 liter/100km amounts to €4,450.

²¹These numbers are somewhat lower, but still of a comparable order of magnitude as those in Busse, Knittel and Zettelmeyer (2013), based on an entirely different identification approach. They find a market share change by +21% for highest fuel economy quartile and by -27% for lowest quartile). This larger impact may be because of a generally larger cost sensitivity (with respect to prices and fuel costs) in the U.S., or because of a different identification approach.

parameter γ decreases. Most interestingly, under the actual estimate of $\gamma = 0.87$, the product tax increase appears to have a *lower* effect on the composition of new car sales than the fuel tax. On the one hand, a product tax would be more effective because of the modest consumer myopia. But on the other hand, a product tax does not target low and high mileage consumers differently, and this effect turns out to dominate. As a result, the product tax is also less effective in altering the composition of new car sales if consumers would be fully forward looking ($\gamma = 1$). In contrast, with strong consumer myopia ($\gamma = 0.5$) the product tax becomes more effective than the fuel tax.

Effects on average fuel economy and total fuel usage Table 5 shows what these sales composition effects imply for the average fuel economy and the total annual fuel usage of new sold cars. While we are mainly interested in the effects based on the parameters of the full model, we also compare the results with the simple logit model that does not include mileage and other heterogeneity. Since both models imply only modest (and statistically insignificant) undervaluation, the table no longer shows how the results would change if we would have fully forward-looking behavior ($\gamma = 1$) or strong myopia, but the results are available in Appendix B.

The first two columns of Table 5 show the effects of the taxes on average fuel economy and total annual fuel usage, conditional on consumers continuing to purchase a car (rather than substituting to the outside good). In the simple logit model without mileage heterogeneity, the fuel tax reduces average fuel economy and total annual fuel usage by exactly the same amount, i.e. by 2.0%. The product tax is more effective since it reduces average fuel economy and total fuel usage by 2.7%. This bigger impact is due the (modest) undervaluation of future fuel costs ($\gamma = 0.87$).

The results are rather different in the full model with mileage and other sources of heterogeneity. First, the fuel tax reduces average fuel economy by 1.3%, but reduces total fuel usage by a much larger 1.8%. This is because the fuel tax mainly targets the high mileage consumers, who are most likely to substitute to cars with a better fuel economy.²² Second, the product tax reduces both average fuel economy and total fuel usage by 1.6%: the product tax is therefore more effective than the fuel tax in reducing average fuel economy, but *less* effective in reducing total fuel usage. This contrasts with our finding for the simple logit, where the product tax seemed to be *more* effective in reducing fuel usage. In sum, despite the consumers' modest undervaluation of future fuel costs, a product tax is on balance less effective than a fuel tax because it cannot target high mileage consumers under mileage

²²This is confirmed from the breakdown by mileage quartile: the lowest mileage quartile reduces fuel usage by only 0.7%, while the largest quartile reduces total fuel usage by 2.5% (not shown in the table).

heterogeneity.

The last two columns show what the taxes imply for total fuel usage, allowing consumers to substitute to the outside good. First notice that the logit model implies an implausibly large substitution to the outside good: the market share of the outside good increases by 10% points after a fuel tax, and even by 13% points after a revenue-equivalent product tax. The full model implies smaller substitution to the outside good, by +5% points for both taxes. This smaller substitution follows from our earlier empirical finding that high mileage consumers tend to have a higher valuation for buying a new car rather than the outside good (positive covariance parameter σ^{0m}).

These differences for outside good substitution translate into different effects of both taxes on total fuel usage. The simple logit model implies an implausibly large reduction in total fuel usage by 32.8% after the fuel tax and by up to 42.8% for the revenue-equivalent product tax (where the larger effect of the product tax again stems from the modest consumer myopia). In contrast, the full model with consumer heterogeneity implies more reasonable reductions in total fuel usage: the fuel tax lowers total annual fuel usage by 14.2% and the product tax reduces it by 12.6%. The smaller effect of the product tax is again because it fails to target consumers with a high mileage.

Effects of discriminatory taxes by fuel type The above counterfactuals focused on the effectiveness of fuel versus product taxes, without discriminating by fuel type. We now consider the effects of discriminatory fuel taxes. This is of broad interest, since countries are increasingly using fuel-specific tax policies to encourage the adoption of cars with alternative fuels (such as methanol). We consider here the effects of discriminatory taxes towards gasoline and diesel fuel. Most European countries have followed such a discriminatory policy during the past decades, with considerably lower taxes for diesel than for gasoline fuel. A justification of this policy was the promotion of diesel cars, which have much lower CO₂ emissions. More recently, there is an increasing pressure to harmonize the gasoline and diesel fuel prices (up to the level of the gasoline fuel price), because of an increased awareness of the negative effects of diesel on local air quality, despite the lower CO₂ emissions.²³

Table 6 shows the effects of a 50c gasoline and/or diesel fuel tax on total fuel usage for the full random coefficients logit model.²⁴ A 50c gasoline fuel tax increase would increase

²³Miravete, Moral and Thurk (2014) stress these negative environmental effects of diesel fuel, and analyze how the favorable fuel taxation towards diesel cars in Europe may have protected European firms from foreign competition.

²⁴The results for the simple logit appear in Appendix B. We no longer show the results for the revenue-equivalent product taxes as this gives no major new insights (other than strengthening the conclusion of its lower effectiveness due to the presence of mileage heterogeneity).

the diesel market share by 19.3% points. This results in a reduction of gasoline fuel usage by 10.7% and in an increase of diesel fuel usage by 5.3%. Conversely, a 50c diesel fuel tax increase would reduce the diesel market share by 15.3%, which results in a reduction of diesel fuel usage by 11.0% and an increase in gasoline fuel usage of 6.0%. These market share effects are much smaller in a logit model without mileage heterogeneity (Appendix B).

These large shifts between gasoline and diesel also have interesting effects on total fuel usage. A 50c gasoline fuel tax reduces total annual fuel usage by 3.1%. This is considerably more than the reduction in total fuel usage when both types of fuel are taxed (−1.8% as seen above, and shown for comparison again on the bottom row of Table 6). This is because a gasoline-only tax induces consumers to substitute to diesel cars, which have a better fuel economy of at least 20%. For the same reason, a 50c fuel tax on only diesel fuel actually *raises* total fuel usage by 1.7%: after such a tax increase, a considerable number of consumers substitute to gasoline cars, which have a lower fuel economy.

In sum, the role of mileage heterogeneity becomes even more important under fuel-specific tax policies. Only raising the price of gasoline fuel implies a large shift to diesel cars with higher fuel economy, and this shift is especially by high mileage consumers. As a result, a discriminatory policy can be especially effective in reducing total fuel usage and hence CO₂ emissions. In the case of diesel fuel this can come at the expense of creating other environmental costs from diesel cars, but for other fuels (such as methanol) discriminatory tax policies may have unambiguously positive effects.

Using the model to explain the market share of gasoline and diesel cars across Europe The previous discussion showed how a discriminatory tax on gasoline or diesel cars has even stronger effects on total fuel usage than a uniform tax in the presence of mileage heterogeneity. We illustrated this for Germany in 2011. We now take a different perspective, and ask to which extent the observed discriminatory fuel taxes and differences in fuel economy can explain the market share of diesel cars across different European countries over time. This serves as a further examination of the implications of our model, and is also of independent interest in light of the importance of diesel cars in Europe.

Table 7 shows the results. The first three columns show the currently observed gaps in fuel prices and fuel economy between gasoline and diesel cars, and the diesel market shares in the seven countries in 1998 and 2011. Diesel fuel was on average 20c per liter less expensive than gasoline fuel in 1998. The highest diesel fuel discounts applied in Belgium, France and the Netherlands (25c–29c per liter), while there was no discount in the U.K. The gap between the gasoline and diesel fuel price somewhat narrowed for all countries in 2011. Furthermore, diesel cars had a better average fuel economy of about 2 liter/100km or 25% in 1998, and

this further improved to 28% in 2011. The lower fuel price and better fuel economy coincides with market shares for diesel cars in the range of 15%–52% in 1998, and 30%–75% in 2011.

The last three columns show how the diesel market share would change if the fuel price gap and fuel economy gap were eliminated. If the price of diesel fuel would be harmonized to the level of gasoline fuel in 1998, the market share of diesel cars would have been 7–14% lower (with the exception of the UK, where there was no fuel price gap). If in addition the fuel economy of diesel and gasoline cars would be equalized, the market share of diesel cars would have been an additional 5–14% lower. In 1998, both factors together explain about 50% of the diesel market share in Belgium, the UK and Spain, and an even larger fraction of 63% in France and Italy, 67% in Germany and 73% in the Netherlands.²⁵ By 2011, the importance of both factors in explaining the diesel market share has diminished because the gap between gasoline and diesel fuel prices narrowed, while diesel cars further gained in popularity. Nevertheless, discriminatory fuel taxes and differences in fuel economy still explain about 40% of the diesel market share across countries.

7 Conclusion

We have analyzed to which extent car buyers undervalue future fuel costs, and what this implies for the effectiveness of alternative tax policies. We specifically demonstrated the importance of accounting for consumer heterogeneity in car utilization and other dimensions.

To estimate the demand model, we used detailed data at the level of the car model and engine variant for a panel of seven European countries during 1998-2011. The institutional features of the European car market provided a new way to identify consumer responses to fuel costs, because for most car models consumers can choose between either a gasoline and diesel engine. To empirically identify consumer responses to fuel costs, we thus exploit fuel price variation by fuel type (i.e. gasoline or diesel) across countries and over time, interacted with variation in fuel economy across products.

We find evidence of at most only modest undervaluation of future fuel costs. To obtain this estimate, it was important to account for both mileage heterogeneity and heterogeneity in the valuation of other car characteristics. We then draw implications for the effectiveness of fuel taxes versus product taxes (based on a car's fuel economy). Despite the modest undervaluation, we find that a fuel tax is more effective in reducing fuel usage than a revenue-equivalent product tax on fuel economy, because it specifically targets the high mileage consumers to substitute to cars with a higher fuel economy. Discriminatory taxes by fuel

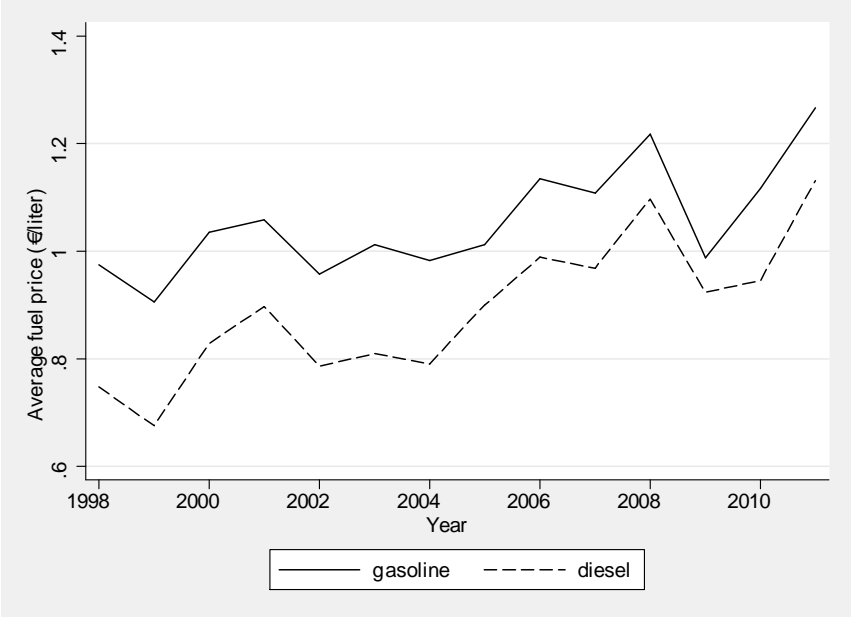
²⁵Note that our model explains a considerably larger part of the diesel market share than Linn (2014), who does not allow for heterogeneous consumer responses to fuel costs.

type are even more effective in targeting consumers with different mileage.

Because our framework makes use of aggregate demand data, it can be used to address a variety of other environmental questions based on datasets for many countries over a long time period. In future research, it would therefore be interesting to apply our framework to investigate the effects of specific policies that countries have followed over the past years, and perform a more complete welfare analysis. This would include the effects on consumers and producers, accounting for several external environmental costs of fuel usage. To make such analysis more realistic, it would also be interesting to extend our aggregate framework to incorporate the possibility that driving behavior is not perfectly inelastic.

8 Figures and Tables

Figure 1: Yearly Average Gasoline and Diesel Prices



The figure shows average yearly prices (in real 2000 €) of gasoline and diesel between 1998 and 2011. Source: DG ECFIN.

Table 1: Summary Statistics

Variable	All years		1998	2011
	Mean	St. Dev.	Mean	Mean
Sales (1,000 units)	1.8	5.1	2.7	1.3
Price/Income	1.1	0.7	1.1	1.1
(Inverse) fuel economy (li/100km)	7.3	2.1	7.8	6.3
Yearly Fuel Costs/Income ($\times 100$)	4.3	1.5	4.7	4.3
Horsepower (in kW)	107.4	54.6	85.1	122.0
Size (1,000 cm ²)	76.7	9.9	74.0	79.5
Height (cm)	148.5	10.7	144.1	150.1
Foreign (0-1)	0.9	0.4	0.8	0.9
Diesel (0-1)	0.4	0.5	0.3	0.5
Months market presence (1-12)	11.4	1.8	11.4	11.5
Number of observations	82,166		4,380	6,898

The table reports means and standard deviations of the main variables for all years and for years 1998 and 2011 separately. The total number of observations (model/engines combinations and markets) is 82,166, where market refer to 7 countries and 14 years.

Table 2: Summary Statistics by Engine Type

	1998		2011	
	Gasoline	Diesel	Gasoline	Diesel
Market Share (%)	68.3	31.7	42.3	57.7
# variants per model	2.5	1.6	3.0	2.5
(Inverse) fuel economy (li/100km)	7.4	6.2	5.7	4.8
Price/Income	0.7	0.9	0.6	0.9

The table reports summary statistics by engine type (gasoline and diesel) in year 1998 and 2011. Fuel economy and price/income are averages weighted by the number of units sold.

Table 3: Parameter Estimates for Alternative Demand Models

	Logit		RC Logit I		RC Logit II	
	Estimate	St.Error	Estimate	St.Error	Estimate	St.Error
	Mean valuations					
Price/Inc. (α)	-4.54	0.19	-5.20	0.20	-4.37	0.22
Fuel Costs/Inc. ($\alpha\gamma\rho$)	-39.90	1.40	-67.93	3.10	-39.37	2.67
Power (kW/100)	2.32	0.14	2.84	0.15	-1.90	0.42
Size (cm ² /10,000)	13.71	0.43	15.92	0.45	18.15	0.50
Height (cm/100)	3.07	0.29	4.08	0.31	4.70	0.34
Foreign	-0.84	0.02	-0.80	0.02	-0.93	0.02
	Standard Deviations of valuations					
Power (kW/100)	-	-	-	-	2.32	0.17
Constant×Mileage	-	-	-	-	4.02	1.12
Mileage distribution	No		Yes		Yes	
	Valuations of Future Fuel Costs					
Fuel Costs/Price ($\gamma\rho$)	8.78	0.55	13.05	0.78	9.00	0.85
Future Valuation γ	0.85	0.05	1.27	0.08	0.87	0.08
Δ Consumer Surplus (€)	-49.67		-111.88		-22.43	

The table reports the parameter estimates and standard errors for the different demand models. The logit assumes homogeneous mileage ($\beta_i^m = \bar{\beta}^m$) and homogeneous valuations for characteristics in x_{jkt} ($\beta_i^x = \bar{\beta}^x$). The random coefficients logit I assumes heterogeneous mileage (β_i^m) and homogeneous valuations for all the other characteristics in x_{jkt} ($\beta_i^x = \bar{\beta}^x$). The random coefficients logit II assumes heterogeneous mileage (β_i^m) and heterogeneous valuations for characteristics in x_{jkt} (β_i^x). Each specification includes model, market/diesel and market/time controls. The total number of observations (combinations of model/engine/market) is 82,166, where markets refer to 7 countries and 14 years. The lower panel reports: (i) the Ratio Fuel Costs/Price ($\gamma\rho$), which converts annual costs into their present discounted value; (ii) the attention weight or future valuation parameter (γ), calculated assuming a market interest rate $r = 6\%$ and an expected car longevity $S = 15$; (iii) the consumer surplus losses per vehicle in €.

Table 4: The Effect of a Fuel Tax and a Product Tax on Market Shares by Fuel Economy Quartile

Change in Market Share	Current	$\gamma = 0.50$		$\gamma = 0.87$		$\gamma = 1.00$	
		% Point	%	% Point	%	% Point	%
Fuel Tax							
Tax per liter		€0.50		€0.50		€0.50	
Fuel Economy Q1 (highest)	37	2.3	7.3	4.4	12.0	5.2	13.3
Fuel Economy Q2	37	-0.4	-1.0	-1.5	-4.1	-2.0	-5.6
Fuel Economy Q3	20	-1.1	-4.8	-1.9	-9.5	-2.1	-11.1
Fuel Economy Q4 (lowest)	6	-0.8	-10.1	-1.0	-16.1	-1.0	-17.7
Revenue Equivalent Product Tax							
Tax per liter/100km		€839.98		€767.37		€739.34	
Fuel Economy Q1 (highest)	37	3.6	11.8	3.8	10.4	3.8	9.8
Fuel Economy Q2	37	-0.2	-0.5	-0.7	-1.9	-0.8	-2.3
Fuel Economy Q3	20	-1.8	-8.1	-1.8	-9.1	-1.8	-9.4
Fuel Economy Q4 (lowest)	6	-1.6	-20.7	-1.3	-20.9	-1.2	-20.8

The table reports the % point and % changes of a fuel tax and a revenue-equivalent product tax on the market shares of new cars aggregated by quartile of fuel economy. Market shares effects are estimated under strong consumer myopia ($\gamma = 0.50$), under the actual estimate of consumer valuation of fuel costs ($\gamma = 0.87$) and under full forward looking behavior ($\gamma = 1.00$), on the basis of the parameter estimates of RCLogit II in Table 3. The figures refer to Germany in 2011. Q1=quartile 1; Q2=quartile 2; Q3=quartile 3; Q4=quartile 4.

Table 5: The Effect of a Fuel Tax and a Product Tax on Fuel Economy and Fuel Usage

	Conditional on Buying		Unconditional on Buying	
	Fuel Economy	Fuel Usage	Fuel Usage	Outside Good
	% Change	% Change	% Change	% Change
Logit model				
Fuel Tax	-2.0	-2.0	-32.8	10.1
Revenue Eq. Product Tax	-2.7	-2.7	-42.7	13.2
RC Logit II				
Fuel Tax	-1.3	-1.8	-14.2	4.7
Revenue Eq. Product Tax	-1.6	-1.6	-12.6	5.0

The table reports the effect of a fuel tax and a revenue-equivalent product tax on average fuel economy and total annual fuel usage, conditional and unconditional on consumers continuing to buy a car. The simulations are based on the parameter estimates in Table 3. The figures refer to Germany in 2011.

Table 6: The Effect of a Fuel Tax on Gasoline and Diesel Separately

	Fuel Usage			Diesel Share
	All cars	Gasoline cars	Diesel cars	
	% Change	% Change	% Change	% Point Change
Fuel Tax, Gasoline only	-3.1	-10.7	5.6	19.3
Fuel Tax, Diesel only	1.7	6.2	-11.0	-15.3
Fuel Tax, both Gas. and Diesel	-1.8	-2.6	-0.3	4.3

The table reports the effect of a discriminatory fuel tax for gasoline and diesel cars on total fuel usage, conditional on consumers continuing to buy a car. The simulations are based on the parameter estimates of RC Logit II in Table 3. The figures refer to Germany in 2011.

Table 7: Explaining the Diesel Market Shares Across Countries

	Current Situation			Change in Diesel Share: Equalization of		
	Fuel price gap	Fuel Economy Gap	Diesel Share	Fuel Price	Fuel Economy	Both
			Year 1998			
Belgium	-0.25	-1.91	52%	-12%	-12%	-24%
France	-0.26	-2.04	38%	-14%	-10%	-24%
Germany	-0.22	-2.35	15%	-5%	-5%	-10%
Great Britain	0.00	-2.24	15%	0%	-7%	-7%
Italy	-0.19	-2.08	21%	-7%	-6%	-13%
Netherlands	-0.29	-2.12	21%	-9%	-6%	-15%
Spain	-0.11	-1.85	51%	-9%	-14%	-24%
			Year 2011			
Belgium	-0.24	-2.01	75%	-4%	-14%	-18%
France	-0.19	-1.81	69%	-5%	-18%	-23%
Germany	-0.16	-2.23	46%	-5%	-17%	-22%
Great Britain	0.04	-2.04	50%	2%	-17%	-16%
Italy	-0.13	-2.05	56%	-5%	-19%	-24%
Netherlands	-0.28	-2.22	30%	-6%	-12%	-18%
Spain	-0.07	-2.08	69%	-3%	-25%	-28%

The table reports: (i) in the first three columns, the currently observed gaps in fuel prices and fuel economy between gasoline and diesel cars, and the diesel market shares in the seven countries of our dataset in 1998 (upper panel) and 2011 (lower panel); in the last three columns, how the diesel market share would change if the fuel price gap and fuel economy gap were eliminated. The simulations are based on the parameter estimates of RC Logit II in Table 3.

References

- Adamou, Adamos, Sofronis Clerides, and Theodoros Zachariadis**, “Welfare Implications of Car Feebates: A Simulation Analysis,” *The Economic Journal*, 2014, 124 (578), F420–F443.
- Allcott, Hunt and Michael Greenstone**, “Is There an Energy Efficiency Gap?,” *The Journal of Economic Perspectives*, 2012, pp. 3–28.
- **and Nathan Wozny**, “Gasoline Prices, Fuel Economy, and the Energy Paradox,” *Review of Economics and Statistics*, 2012, forthcoming.
- **, Sendhil Mullainathan, and Dmitry Taubinsky**, “Energy Policy with Externalities and Internalities,” *Journal of Public Economics*, April 2014, 112.
- Alquist, Ron, Lutz Kilian, and Robert J. Vigfusson**, “Forecasting the Price of Oil,” CEPR Discussion Papers 8388, C.E.P.R. Discussion Papers May 2011.
- Anderson, Soren T., Ian W. H. Parry, James M. Sallee, and Carolyn Fischer**, “Automobile Fuel Economy Standards: Impacts, Efficiency, and Alternatives,” *Review of Environmental Economics and Policy*, Winter 2011, 5 (1), 89–108.
- **, Ryan Kellogg, and James M. Sallee**, “What Do Consumers Believe About Future Gasoline Prices?,” *Journal of Environmental Economics and Management*, 2013, 66 (3), 383–403.
- Baltagi, Badi H.**, *Econometric Analysis of Panel Data*, John Wiley & Sons, 1995.

- Bento, Antonio M., Lawrence H. Goulder, Mark R. Jacobsen, and Roger H. von Haefen**, “Distributional and Efficiency Impacts of Increased US Gasoline Taxes,” *American Economic Review*, June 2009, *99* (3), 667–99.
- , **Shanjun Li, and Kevin Roth**, “Is there an energy paradox in fuel economy? A note on the role of consumer heterogeneity and sorting bias,” *Economics Letters*, 2012, *115* (1), 44 – 48.
- Berry, Steven T.**, “Estimating Discrete-Choice Models of Product Differentiation,” *The RAND Journal of Economics*, 1994, *25* (2), 242–262.
- , **James Levinsohn, and Ariel Pakes**, “Automobile Prices in Market Equilibrium,” *Econometrica*, 1995, *63* (4), 841–890.
- , – , and – , “Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy,” *American Economic Review*, June 1999, *89* (3), 400–430.
- Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer**, “Are Consumers Myopic? Evidence from New and Used Car Purchases,” *American Economic Review*, February 2013, *103* (1), 220–56.
- Chamberlain, Gary**, “Asymptotic efficiency in estimation with conditional moment restrictions,” *Journal of Econometrics*, 1987, *34* (3), 305 – 334.
- D’Haultfoeuille, Xavier, Pauline Givord, and Xavier Boutin**, “The Environmental Effect of Green Taxation: The Case of the French Bonus/Malus,” *The Economic Journal*, 2014, *124* (578), F444–F480.
- Gillingham, Kenneth**, “Selection on Anticipated Driving and the Consumer Response to Changing Gasoline Prices,” Technical Report, Yale University 2012.

- Goldberg, Pinelopi Koujianou**, “The Effects of the Corporate Average Fuel Efficiency Standards in the US,” *Journal of Industrial Economics*, March 1998, *46* (1), 1–33.
- **and Rebecca Hellerstein**, “A Structural Approach to Identifying the Sources of Local Currency Price Stability,” *Review of Economic Studies*, 2013, *80* (1), 175–210.
- Greene, David L.**, “How Consumers Value Fuel Economy: A Literature Review,” Technical Report EPA-420-R-10-008, U.S. Environmental Protection Agency March 2010.
- Hausman, Jerry A.**, “Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables,” *Bell Journal of Economics*, Spring 1979, *10* (1), 33–54.
- Helfand, Gloria and Ann Wolverton**, “Evaluating the consumer response to fuel economy : a review of the literature,” *International review of environmental and resource economics*, 2011, *5*, 103–146.
- Hughes, Jonathan E., Christopher R. Knittel, and Daniel Sperling**, “Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand,” *The Energy Journal*, 2008, *29* (1), 113–134.
- Huse, Cristian and Claudio Lucinda**, “The market impact and the cost of environmental policy: evidence from the Swedish green car rebate,” *The Economic Journal*, 2014, *124* (578), F393–F419.
- Jaffe, Adam B. and Robert N. Stavins**, “The energy paradox and the diffusion of conservation technology,” *Resource and Energy Economics*, May 1994, *16* (2), 91–122.
- Klier, Thomas and Joshua Linn**, “The Price of Gasoline and New Vehicle Fuel Economy: Evidence from Monthly Sales Data,” *American Economic Journal: Economic Policy*, August 2010, *2* (3), 134–53.

- **and** –, “Using Taxes to Reduce Carbon Dioxide Emissions Rates of New Passenger Vehicles: Evidence from France, Germany, and Sweden,” *American Economic Journal: Economic Policy*, 2015, *7*(1), 1–32.
- Langer, Ashley and Nathan Miller**, “Automakers’ Short-Run Responses to Changing Gasoline Prices,” *Review of Economics and Statistics*, 2013, *95* (4), 1198–1211.
- Li, Shanjun, Christopher Timmins, and Roger H. von Haefen**, “How Do Gasoline Prices Affect Fleet Fuel Economy?,” *American Economic Journal: Economic Policy*, August 2009, *1* (2), 113–37.
- Linn, Joshua**, “Explaining the Adoption of Diesel Fuel Passenger Cars in Europe,” Discussion Paper 14-08, Resources for the Future April 2014.
- Miravete, Eugenio, Maria J. Moral, and Jeff Thurk**, “Protecting the European Automobile Industry through Environmental Regulation: The Adoption of Diesel Engines,” April 2014.
- Parry, Ian W. H., Margaret Walls, and Winston Harrington**, “Automobile Externalities and Policies,” *Journal of Economic Literature*, 2007, *45* (2), 373–399.
- Petrin, Amil**, “Quantifying the Benefits of New Products: The Case of the Minivan,” *Journal of Political Economy*, 2002, *110*, 705–29.
- Reynaert, Mathias and Frank Verboven**, “Improving the performance of random coefficients demand models: The role of optimal instruments,” *Journal of Econometrics*, 2014, *179* (1), 83 – 98.
- Sallee, James**, “Rational Inattention and Energy Efficiency,” *Journal of Law and Economics*, 2014, *forthcoming*.

Verboven, Frank, “Quality-Based Price Discrimination and Tax Incidence: Evidence from Gasoline and Diesel Cars,” *RAND Journal of Economics*, Summer 2002, 33 (2), 275–297.

A Appendix A. Effects of a fuel tax and a product tax

In this Appendix, we first derive the analytic expressions for the demand effects of a fuel tax and a product tax based on fuel economy. We then describe the specific approach to implement the policy counterfactuals.

A.1 Impact of small tax changes on demand

Assume for simplicity that the fuel tax and product tax is uniform, i.e. there is no distinction between gasoline and diesel engine k . The individual choice probability for product jk of a consumer i can be written as:

$$s_{ijk} \left(t^G, t^E; \beta_i \right) = \frac{\exp(v_{ijk})}{1 + \sum_{j'}^J \sum_{k'=1}^{K_{j'}} \exp(v_{ij'k'})}$$

and total sales of product jk under taxes (t^G, t^E) are:

$$q_{jk} \left(t^G, t^E \right) = \int_{\beta} s_{ijk} \left(t^G, t^E; \beta \right) dF_{\beta}(\beta) I,$$

where the individual utility minus extreme value random variable is defined as

$$v_{ijk} \equiv x_{jk} \beta_i^x - \alpha_i (p_{jk} + t^E e_{jk} + \gamma \rho \beta_i^m e_{jk} (g_k + t^G)) + \xi_{jk}.$$

The own- and cross-effects of a change in individual utility on the individual choice probabilities take the usual form:

$$\begin{aligned} \frac{\partial s_{ijk}}{\partial v_{ijk}} &= s_{ijk} (1 - s_{ijk}) \\ \frac{\partial s_{ijk}}{\partial v_{ij'k'}} &= -s_{ij'k'} s_{ijk}. \end{aligned}$$

The effect of a uniform fuel tax t^G on the individual choice probability is then:

$$\begin{aligned}
\frac{\partial s_{ijk}}{\partial t^G} &= -\alpha_i \gamma \rho \beta_i^m \sum_{j'} \sum_{k'} \frac{\partial s_{ijk}}{\partial v_{j'k'}} e_{j'k'} \\
&= -\alpha_i \gamma \rho \beta_i^m s_{ijk} \left(e_{jk} - \sum_{j'} \sum_{k'} s_{ij'k'} e_{j'k'} \right) \\
&= -\alpha_i \gamma \rho \beta_i^m s_{ijk} \left(e_{jk} - (1 - s_{i0}) \sum_{j'} \sum_{k'} e_{j'k'} \frac{s_{ij'k'}}{1 - s_{i0}} \right) \\
&= -\alpha_i \gamma \rho \beta_i^m s_{ijk} (e_{jk} - \bar{e}^i + s_{i0} \bar{e}^i)
\end{aligned}$$

where

$$\bar{e}^i = \sum_{j'} \sum_{k'} e_{j'k'} \frac{s_{ij'k'}}{1 - s_{i0}}$$

is the expected fuel economy of consumer i .

The effect of the fuel tax on total demand is then given by:

$$\begin{aligned}
\frac{\partial q_{jk}}{\partial t^G} &= \int_{\beta} \frac{\partial s_{ijk}}{\partial t^G} dF_{\beta}(\beta) I \\
&= - \int_{\beta} \alpha_i \gamma \rho \beta_i^m s_{ijk} (e_{jk} - \bar{e}^i + s_{i0} \bar{e}^i) dF_{\beta}(\beta) I.
\end{aligned}$$

We can follow similar steps to compute the effect of product tax t^E , so that the effects of both taxes are summarized as:

$$\begin{aligned}
\frac{\partial q_{jk}}{\partial t^G} &= - \int_{\beta} \alpha_i \gamma \rho \beta_i^m s_{ijk} (e_{jk} - \bar{e}^i + s_{i0} \bar{e}^i) dF_{\beta}(\beta) I \\
\frac{\partial q_{jk}}{\partial t^E} &= - \int_{\beta} \alpha_i s_{ijk} (e_{jk} - \bar{e}^i + s_{i0} \bar{e}^i) dF_{\beta}(\beta) I,
\end{aligned}$$

which are the expressions presented in the main text. This shows several things. First, the effect of tax is similar to a price elasticity of industry demand, except for the term $e_{jk} - \bar{e}^i$. If $e_{jk} - \bar{e}^i = 0$, the effect is just like elasticity of industry demand. If $e_{jk} > \bar{e}^i$, then the

effect is bigger (worst fuel efficient cars loose most). If $e_{jk} - \bar{e}^i < 0$, the effect is smaller and may easily turn positive. Second, the energy tax is different from product tax because of γ and β_i^m . This can be confirmed from revenue equivalent tax below. Note also that the expressions simplify if the outside good is absent (inelastic market demand). Then the sign of the tax effect simply depends on sign of $e_{jk} - \bar{e}^i$.

A.2 Details on the policy counterfactuals

We derive the expressions used in our policy counterfactuals to compute the effects of the fuel tax t_k^G and the product tax t_k^E on revenues, market shares, total fuel usage and average fuel economy. Let $k = 1$ refer to gasoline, and $k = 2$ refer to diesel. Denote the vector of taxes by (t^G, t^E) , where $t^G = (t_1^G, t_2^G)$ is the energy tax vector, and $t^E = (t_1^E, t_2^E)$ is the product tax vector.

Sales We slightly modify some of the expressions in the previous subsection to account for the fact that the fuel tax and product tax can vary per fuel type. The choice probability for product jk of a consumer i with a random coefficient vector $\beta_i = (\beta_i^x, \alpha_i, \beta_i^m)$ facing tax vector t^G and t^E is:

$$s_{ijk} \left(t^G, t^E; \beta_i \right) = \frac{\exp(v_{ijk})}{1 + \sum_{j'}^J \sum_{k'=1}^{K_{j'}} \exp(v_{ij'k'})},$$

where the individual utility minus extreme value random variable is now defined with non-uniform taxes as

$$v_{ijk} \equiv x_{jk} \beta_i^x - \alpha_i (p_{jk} + t_k^E e_{jk} + \gamma \rho \beta_i^m e_{jk} (g_k + t_k^G)) + \xi_{jk}.$$

Total sales of product product jk under taxes (t^G, t^E) are again:

$$q_{jk}(t^G, t^E) = \int_{\beta} s_{ijk}(t^G, t^E; \beta) dF_{\beta}(\beta)I$$

So the current quantity is $q_{jk}(0, 0)$, the predicted quantity after only a fuel tax is $q_{jk}(0, t^G)$ and the predicted quantity after only a product tax is $q_{jk}(t^E, 0)$. Based on these predicted quantities per product jk we can compute the market shares per fuel economy quartile (or any other aggregated quantity or market share).

Tax revenues Conditional on buying product jk , an individual consumer pays a taxes $(t_k^E + \rho\beta_i^m t_k^G) e_{jk}$, i.e. the sum of the product tax plus capitalized future energy taxes. Total tax revenues over all products jk are defined as:

$$R(t^G, t^E) = \int_{\beta} \sum_j \sum_k (t_k^E + \rho\beta_i^m t_k^G) e_{jk} s_{ijk}(t_k^G, t_k^E; \beta) dF_{\beta}(\beta)I,$$

i.e. the expected tax revenue over all cars per consumer, averaged over all consumers. We can then compute the tax revenues from only an energy tax $R(t^G, 0)$ or only a product tax $R(0, t^E)$ as:

$$\begin{aligned} R(t^G, 0) &= \int_{\beta} \sum_j \sum_k \rho\beta_i^m t_k^G e_{jk} s_{ijk}(t_k^G, 0; \beta) dF_{\beta}(\beta)I \\ R(0, t^E) &= \int_{\beta} \sum_j \sum_k t_k^E e_{jk} s_{ijk}(0, t_k^E; \beta) dF_{\beta}(\beta)I. \end{aligned}$$

With uniform taxes, we consider a 50c energy tax, $R(0.5, 0)$, so that the revenue-neutral product tax is the solution of t^G to $R(0, t^G) = R(0.5, 0)$.

Average fuel economy and total energy usage Conditional average fuel economy, given that people purchase a car is

$$E(t^G, t^E) = \sum_j \sum_k e_{jk} \frac{s_{jk}(t_k^G, t_k^E; \beta_i^m)}{1 - s_{00}(t_k^G, t_k^E; \beta_i^m)},$$

where

$$s_{jk}(t^G, t^E) = \int_{\beta} s_{ijk}(t^G, t^E; \beta) dF_{\beta}(\beta)$$

Conditional total annual fuel usage given that people buy a car is

$$F_C(t^G, t^E) = \int_{\beta} \sum_j \sum_k \beta_i^m e_{jk} s_{ijk|B}(t_k^G, t_k^E; \beta) dF_{\beta}(\beta) I,$$

and unconditional total annual energy usage is given by

$$F_U(t^G, t^E) = \int_{\beta} \sum_j \sum_k \beta_i^m e_{jk} s_{ijk}(t_k^G, t_k^E; \beta) dF_{\beta}(\beta) I.$$

This accounts for the fact that consumers may substitute to the outside good after a tax increase, so that they do not consume any fuel.

Based on this, we can compute the percentage change in average fuel economy and the percentage change in fuel usage for both taxes. For example, the percentage change in fuel economy for a fuel tax is $\% \Delta E = E(t^G, 0) / E(0, 0) - 1$.

B Appendix B. Additional Tables

Table B.1: The Effect of a Fuel Tax and a Product Tax on Market Shares by Fuel Economy Quartile - Logit

Change in Market Share:	Current	$\gamma = 0.50$		$\gamma = 0.87$		$\gamma = 1.00$	
		% Point	%	% Point	%	% Point	%
Fuel Tax							
Fuel Tax per liter		€0.50		€0.50		€0.50	
Fuel Economy Q1 (highest)	37	2.0	6.3	3.4	9.3	3.9	10.2
Fuel Economy Q2	37	0.3	0.8	-0.1	-0.3	-0.4	-1.1
Fuel Economy Q3	20	-1.0	-4.5	-1.8	-8.9	-2.1	-11.0
Fuel Economy Q4 (lowest)	6	-1.3	-15.1	-1.5	-24.2	-1.5	-27.7
Revenue Equivalent Product Tax							
Tax per liter/100km		€ 1,352.00		€ 890.03		€ 759.79	
Fuel Economy Q1 (highest)	37	7.1	22.1	4.6	12.7	3.9	10.2
Fuel Economy Q2	37	0.4	1.0	-0.3	-0.7	-0.4	-1.1
Fuel Economy Q3	20	-3.6	-16.4	-2.4	-12.2	-2.1	-11.0
Fuel Economy Q4 (lowest)	6	-3.8	-44.0	-1.9	-31.6	-1.5	-27.7

The table reports the effect of a fuel tax and a revenue-equivalent product tax on the market shares of new cars aggregated by quartile of fuel economy. Market shares effects are estimated under strong consumer myopia ($\gamma = 0.50$), under the actual estimate of consumer valuation of fuel costs ($\gamma = 0.87$) and under full forward looking behavior ($\gamma = 1.00$), on the basis of the parameter estimates of Logit in Table 3. The figures refer to Germany in 2011. All values are in percent. Q1=quartile 1; Q2=quartile 2; Q3=quartile 3; Q4=quartile 4.

Table B.2: The Effect of a Fuel Tax and a Product Tax on Fuel Economy and Fuel Usage Under Different Valuation Parameters

	Conditional on Buying		Unconditional on Buying	
	Fuel Economy % Change	Fuel Usage % Change	Fuel Usage % Change	Outside Good % Change
Logit model with $\gamma = 0.5$				
Fuel Tax	-1.5	-1.5	-17.6	10.0
Revenue Eq. Product Tax	-4.6	-4.6	-53.7	24.0
RC Logit II with $\gamma = 0.5$				
Fuel Tax	-0.8	-1.2	-7.2	2.8
Revenue Eq. Product Tax	-1.9	-1.9	-12.6	5.8
Logit model with $\gamma = 1$				
Fuel Tax	-2.2	-2.2	-38.8	10.0
Revenue Eq. Product Tax	-2.2	-2.2	-38.8	10.0
RC Logit II with $\gamma = 1$				
Fuel Tax	-1.5	-1.9	-16.8	5.3
Revenue Eq. Product Tax	-1.5	-1.5	-12.4	4.6

The table reports the effect of a fuel tax and a revenue-equivalent product tax on average fuel economy and total annual fuel usage, conditional and unconditional on consumers continuing to buy a car for two values of consumer valuation of fuel costs (γ): $\gamma = 0.5$ and $\gamma = 1$. The simulations are based on the parameter estimates in Table 3. The figures refer to Germany in 2011.

Table B.3: The Effect of a Fuel Tax on Gasoline and Diesel Separately

	Total fuel usage			Diesel share
	All cars	Gasoline cars	Diesel cars	
	% change	% change	% change	% point change
Fuel Tax, Gasoline only	-2.7	-1.9	0.0	13.6
Fuel Tax, Diesel only	1.1	0.0	-1.6	-11.4
Fuel Tax	-2.0	-1.9	-1.6	1.8

The table reports the effect of a discriminatory fuel tax for gasoline and diesel cars on total fuel usage, conditional on consumers continuing to buy a car. The simulations are based on the parameter estimates of Logit in Table 3. The figures refer to Germany in 2011.

