

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

Laura Grigolon, Mathias Reynaert, Frank Verboven*

May 2017

Abstract

To what extent do car buyers undervalue future fuel costs, and what does this imply for the effectiveness and welfare impact 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 costs by engine type. Although we find there is modest undervaluation of fuel costs, fuel taxes are still more effective in reducing fuel usage than product taxes based on fuel economy. They also perform much better in terms of total welfare even if usage demand is inelastic. The reason is that fuel taxes better target the right consumers, those with a high mileage, to purchase more fuel efficient cars.

**Laura Grigolon*: McMaster University. Email: Laura.Grigolon@mcmaster.ca. *Mathias Reynaert*: Toulouse School of Economics. Email: Mathias.Reynaert@gmail.com *Frank Verboven*: University of Leuven and C.E.P.R. (London). Email: Frank.Verboven@kuleuven.be. We are particularly grateful to the Editor and three referees for helpful comments. We also thank Bruno De Borger, Jan Bouckaert, James Sallee, Jo Van Biesebroeck and seminar participants at several seminars and conferences. 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 by properly accounting for consumer heterogeneity, which implies an important advantage of fuel taxes over product taxes that has not been considered before. Fuel taxes better target the right consumers, those with a high mileage, to purchase the most fuel efficient cars. This targeting effect makes fuel taxes more effective in reducing total fuel usage than product taxes, even when consumers are moderately short sighted and mileage is perfectly inelastic. We establish this result in two steps. We first ask whether consumers undervalue or correctly value the discounted future fuel costs when purchasing a new car. After estimating the degree of consumer myopia, we then compare the effectiveness of a fuel and product tax in reducing total fuel usage and CO₂ emissions, where the taxes are held revenue equivalent. We also compare the welfare effects of both taxes in correcting for the externality costs of driving. In both steps it is crucial to account for consumer heterogeneity. In the first step accounting for consumer heterogeneity avoids possibly biased parameter estimates of consumers' valuation of future fuel costs. In the second step mileage heterogeneity implies that heavy users (and hence heavy polluters) will be more responsive to fuel taxes than product taxes, with implications for both the effectiveness and welfare effects of the taxes.

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 to the fuel price. We show that the relative effectiveness of fuel taxes versus product taxes in reducing total fuel usage 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 provide a new way to help identifying 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 include car model fixed effects, thereby exploiting fuel cost variation between engines of the same car model. Chugh, Cropper, and Narain (2011) and Verboven (2002) also exploit fuel cost variation between gasoline and diesel cars: the former in a hedonic price study of Indian twin car models, the latter in a demand study that focuses on substitution within both fuel types but abstracts from substitution across products. Our approach complements other recent 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 (2014), Busse, Knittel, and Zettelmeyer (2013), Langer and Miller (2013), Sallee, West, and Fan (2016)). We document that time series variation is not important in our setting to identify consumer responses to fuel economy.¹

Regarding the valuation of future fuel costs, we find evidence of only modest undervaluation: for one euro saving in discounted future fuel costs, consumers are willing to pay €0.91 in the form of a higher initial purchase price, and the hypothesis of correct valuation cannot be rejected. The estimates are slightly lower if we do not account for consumer heterogeneity, and the results remain robust in a variety of other specifications.

Regarding the effectiveness of fuel taxes versus product taxes, we find that it is crucial to account for mileage heterogeneity. Without accounting for mileage heterogeneity, a fuel tax would be slightly less effective than a product tax in reducing total fuel usage, because

¹The 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).

of the modest undervaluation of future fuel costs. However, once we account for mileage heterogeneity, a fuel tax turns out to be more effective because it specifically targets the high mileage consumers to substitute to cars with a higher fuel economy. Note that these conclusions depend on our finding that there is only limited undervaluation of future fuel costs. If there would be stronger consumer myopia, our conclusion would be reversed: a product tax would then be more effective in reducing total fuel usage (despite the mileage heterogeneity).²

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, raising the diesel fuel price without changing the gasoline fuel price implies a drop in diesel market share by 12.3%, but a disproportionately larger drop in diesel fuel usage by 37.8%.³ We also find that the lower fuel costs for diesel cars explain about half of the diesel market share in Europe, though less so in the most recent years.

Finally, we consider the welfare implications of our analysis. We compare the impact of a fuel and product tax on the various welfare components: decision consumer surplus, consumer losses from misoptimization (internality), externality costs and tax revenues. We find that a fuel tax performs much better than a product tax, even when usage demand itself is perfectly inelastic. Fuel taxes require a lower total tax amount to achieve the same externality reduction because they directly target the high mileage consumers. In contrast, product taxes affect all consumers in the same way, which results in lower welfare gains despite the modest undervaluation of future fuel costs. This again stresses the importance of accounting for mileage heterogeneity in car taxation policy.

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

²These conclusions are 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 in favour of fuel taxes, since fuel taxes would also induce consumers to reduce their car usage.

³Miravete, Moral, and Thurk (2014) study how the favorable treatment of diesel cars may have protected the domestic European producers against foreign competition in the Spanish market. They do not explicitly model mileage heterogeneity, and our analysis implies that this implies even greater environmental costs.

for the degree of undervaluation of future fuel costs; see Greene (2010) and Helfand and Wolverton (2011) for recent reviews. Allcott and Wozny (2014) find evidence of moderate undervaluation, while Busse, Knittel, and Zettelmeyer (2013), Sallee, West, and Fan (2016), Chugh, Cropper, and Narain (2011) and Verboven (2002) 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 through simulations that a failure to account for consumer heterogeneity in willingness to pay for fuel costs may be 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 and improve total welfare. Allcott and Wozny (2014) 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.⁵ Our contribution to this debate is to show that even if usage demand 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.⁶

In a related theoretical paper, Allcott, Mullainathan, and Taubinsky (2014) compare the combined welfare effects of an energy tax and a product tax. They also stress the role of

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.

⁵Jacobsen (2013) and Jacobsen and van Benthem (2015) show that product taxes may also distort the used car market and households' scrappage decision. Our data is limited to new vehicle sales so we abstract from these effects in our analysis.

⁶Our 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.

consumer heterogeneity, but they focus on heterogeneity in the extent of undervaluation (or “attentiveness”). In our welfare analysis we focus on the role of mileage heterogeneity, a natural starting point since this distribution is directly observed. Nevertheless, in future work it would be interesting to extend the welfare evaluation to allow for both heterogeneity in attentiveness, mileage and product quality. Our welfare analysis is also complementary to recent work by Jacobsen, Knittel, Sallee, and van Benthem (2016). They develop a sufficient statistic to measure the importance of heterogeneity in externalities when designing taxes. Our structural framework explicitly focuses on the role of mileage heterogeneity and how this impacts the performance of fuel versus product taxes, while allowing for mean undervaluation of consumers.

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 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.⁷ 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

⁷For 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.

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 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 (2014)’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 the interest rate. More precisely, we define the present value of expected fuel costs over the car’s lifetime as follows:

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

where (i) β_i^m is consumer i ’s expected annual mileage, measured in km; (ii) e_{jk} is the fuel consumption of car j with engine k , measured in liter per km (or “gallons per mile” in the US), i.e. the inverse of a car’s fuel economy; (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.⁸ 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 possible 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.⁹ 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.

⁸In practice, vehicle lifetime may not be fixed but may partly depend on mileage and we account for this possibility in a robustness analysis below.

⁹Other 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 (2014) 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.

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 alternative models do not appear to outperform models with expectations that are only based on current prices. In our robustness analysis we will use futures prices as a proxy for consumers' price forecasts of fuel prices, which gives similar results.

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

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

i.e. consumer i 's annual mileage β_i^m (in km) times the car's fuel consumption e_{jk} (in liter/km) times the current fuel price g_k (in €/liter)km times a capitalization coefficient ρ :

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

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 $[0, S]$. If consumers are fully myopic ($r \rightarrow \infty$), then $\rho = 0$: consumers then give no weight to current annual fuel costs. 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 , setting $r = 0$ and $\gamma = 1$. Third, as in Allcott and Wozny (2014), 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:

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

Observed sales can then be equated 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 (in reducing total fuel usage) and the total welfare impact of two alternative taxes: a fuel tax and a product tax on the cars' fuel consumption. 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 a car's fuel consumption e_{jk} . The product tax on a car's fuel consumption 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 cars' fuel consumption 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 consumption 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 in reducing total fuel usage 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 impact 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 consumption 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 consumption, i.e. $e_{jk} = e$ for all jk , then $e_{jk} - \bar{e}^i = 0$. 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 both taxes is then similar to the effect of an industry-wide price increase. In contrast, if cars differ in their fuel consumption, then the demand for some cars can increase despite the tax increase. This will be the case for cars with a sufficiently low fuel consumption, 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 shift the composition of new car sales from cars with a high fuel consumption e_{jk} to cars with a low fuel consumption.

Second, the effect of both taxes differs because of two factors: the future valuation parameter (γ) and the extent of mileage heterogeneity (β_i^m). In the absence of mileage heterogeneity ($\beta_i^m = \beta^m$ for all i), it can be shown that a fuel tax t^G has exactly the same demand effect as a revenue-equivalent product tax t^E when $\gamma = 1$, and a weaker impact when $\gamma < 1$ (the investment inefficiency referred to above). When there is mileage heterogeneity, the comparison between the demand impact of both taxes becomes more complicated.

To summarize this, the price sensitivity parameter α_i and the extent substitution to the outside good s_{i0} determine the general impact of both taxes, whereas the future valuation parameter γ and the extent of mileage heterogeneity β_i^m explain the differences in their impact. In our empirical analysis we show how these parameters translate into differences in the effectiveness (in reducing total fuel usage) and the total welfare impact of both taxes.

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). According to European survey data (reviewed in Goldberg and Verboven (2005)), discounts mainly vary by car model and country, which we will account for with a rich set of car and country fixed effects.¹⁰ Car characteristics include measures of vehicle size (curb weight, width, length and height), engine performance (horsepower and displacement) and fuel consumption (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.¹¹

We supplement this dataset with information on fuel prices, the empirical distribution 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) as well as the difference between the price of gasoline and diesel across countries. The top panel shows that 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. According to the bottom panel, the price of gasoline is higher than that of diesel in almost all countries. The gap between both fuel prices is the highest in the Netherlands and Belgium. The gap often changes in the same direction across countries, but the magnitudes of these changes may differ considerably.

The empirical distribution of annual miles travelled is from the 2007 UK National Travel Survey, a rich nationally representative survey of 20,000 individuals.¹² We plot the distribution in Figure 2. 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. According to this survey, average annual mileage is 14,700 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

¹⁰For the U.S., Langer and Miller (2013) exploit detailed transaction data. They find that discounts are higher for cars with higher fuel costs (conditional on time and vehicle fixed effects), and point out that this may imply a downward bias of the fuel cost coefficient by 13.7% (or higher). If discounts also vary with fuel costs in Europe, our fuel cost valuation parameter may also be biased downwards to some extent.

¹¹For example, the Volkswagen Golf is perceived as domestic by German consumers even though part of the production of Golf takes place in Spain.

¹²www.gov.uk/government/organisations/department-for-transport/series/national-travel-survey-statistics

Survey¹³. 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 mainly 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. The share of new diesel car sales increased from 31.7% in 1998 to 57.7% in 2011.¹⁴ Also the supply of diesel engines has increased proportionately more than that of gasoline engines: there were on 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 (and 3.0 gasoline engines per car). 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.

More generally, the variation in the up-front purchase price and the expected fuel costs across cars will be of primary importance to identify the extent to which consumers discount the future. This variation comes from three sources: differences between diesel and gasoline engines, differences between engines within a fuel type, and fuel price changes over time. To assess the importance of each source, we regress our fuel cost variable on all other control variables included in the main specification, and then partition the variance of the residual

¹³In the appendix, we re-estimate the model estimated below by rescaling the UK-survey distribution around the mean from Eurostat for each country. This gives very similar results with a somewhat higher willingness to pay for fuel economy. We prefer to rely on the UK distribution for the main results below because most of the variance in the EUROSTAT data comes from different underlying measurements by governments.

¹⁴This 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 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.

fuel costs. We find that variation in fuel prices over time explains less than 0.01% of the total variance of fuel costs. Differences between gasoline and diesel engines explain 14% of the variance in fuel costs, and differences between engines with the same fuel type explain the remaining 86% of variance in fuel costs. Hence, most of the variation in fuel costs comes from engine variation (between and within fuel type). This is different from the existing literature, which has mainly relied on differences in fuel prices over time to identify the extent to which consumers discount future fuel cost savings.

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 consumption 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.

First, we assume that β_i^m follows the empirical distribution of mileage. The mileage distribution we observe (discussed in section 3) is a conditional distribution, i.e. conditional on having purchased a car, while we need the unconditional distribution in our empirical model. In Appendix A.3 we show how to map the observed conditional distribution into the unconditional distribution for β_i^m . Incorporating the empirical distribution of mileage 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. Third, we specify α_i to be inversely proportional to income y_t in market t , so $\alpha_i = \alpha/y_t$.¹⁵

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, size and foreign. In the Appendix we report specifications with additional random coefficients: this gives similar results, but less precise estimates for some our main parameters of interest. 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

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

¹⁵Similar 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.

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. Since we include a full set of car model fixed effects, the demand parameters are identified from variation across engine variants within car models.

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 consumption e_t and fuel prices g_t). This assumption requires that all product characteristics, except prices, are uncorrelated with unobserved car variant quality, conditional on the car model fixed effects. This appears reasonable in the short run when manufacturers cannot easily redesign their products. To the extent $\tilde{\xi}_{jkt}$ is correlated with the observed product characteristics, this may potentially bias the parameters.

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.¹⁶ To construct the instruments $h(z)$, we follow Berry, Levinsohn and Pakes (1995) and use the own-product characteristics, sums of the characteristics of other products of the

¹⁶To 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 .

same firm, and sums of characteristics of other firms’ products. These instruments influence the mark-ups of the differentiated products and therefore serve as instruments for price. We also considered cost shifters as alternative instruments for price: following Reynaert and Verboven (2014) we interact steel prices and unit labor costs with the product characteristics as product-level cost shifters. Since steel prices and labor costs were not available for the year and country of production of about 8% of the observations, we do not take this as our main specification, but report the results in the appendix.

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, especially for the variances of the random coefficients. 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 markets. 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

Although the importance of heterogeneity in consumer preferences has been emphasized in other work, the specific implications on how this may affect the estimated valuation of future fuel cost savings has received limited attention. 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 compare 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.

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 $\alpha = -4.52$ and $\alpha\gamma\rho = -39.03$.¹⁷ This implies $\gamma\rho = 8.63$, which we discuss in detail below. The price parameter α is in line with previous results in the literature. In the appendix (Table B.1) we show that not using any instruments leads to a very low estimated price elasticity. We also present the first stage of the instrumental variable regression, and find that the excluded instruments correlate strongly with prices and are jointly significant with an F-stat of 291. Finally, we re-estimate the logit model with cost shifters as additional instrumental variables, and this leads to very similar results.

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.¹⁸

Now consider the random coefficients logit model with only mileage heterogeneity, and no heterogeneity for other characteristics (RC Logit I). Price and fuel costs again have the expected negative effect on utility, and the magnitude of the fuel cost effect slightly decreases ($\alpha = -6.22$ and $\alpha\gamma\rho = -46.48$). 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 (RC Logit II). As in the other two models, price

¹⁷Table B.2 replicates Table 3 with mileages from Eurostat instead of the UK survey.

¹⁸Table B.5 presents the results for the logit model, estimated separately for each of the countries. The parameter estimates have robust signs, and with some exceptions the magnitudes of the effects are broadly comparable.

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 insignificant, but there is substantial heterogeneity around the mean, implying that about half of the consumers value horsepower positively. There is also significant heterogeneity in the valuation of size, but essentially all consumers have a positive valuation. The standard deviation for the valuation of foreign cars is not significantly different from zero.

We considered a detailed sensitivity analysis with respect to various assumptions of our demand model. Table B.2 presents parameter estimates under alternative measures of our fuel cost variable: the use of average fuel prices per market (taking out the variation in fuel prices over time); the use of futures prices instead of spot prices; and the use of mean mileage reported by Eurostat for each country instead of the UK survey. The empirical results remain very similar. Intuitively, this is because fuel cost variation is more important across engines and fuel types than over time or countries, as discussed above.

Table B.3 presents alternative specifications for our base random coefficients model (RC Logit II): a model that includes a random coefficient for fuel type (RC Logit III), a model that adds heterogeneity in the valuation of fuel costs for other reasons than mileage (RC Logit IV),¹⁹ and a model that allows vehicle lifetime to vary according to the consumer’s annual mileage (RC Logit V).²⁰ For all specifications we obtain similar empirical results. We therefore focus our subsequent analysis on the base specification RC Logit II, but will also provide robustness analysis in the counterfactuals for the other models.

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

¹⁹More specifically, we follow Anderson, Kellogg, and Sallee (2013), who report additional heterogeneity in the valuation of fuel costs for other reasons than mileage such as discounting and forecasting. Specifically, we increase the variance of fuel costs across sampled individuals with a factor of 7 by adding additional normal distributed draws to each mileage type. In this way 85% of the variation between sampled individuals comes from other sources than mileage. This additional heterogeneity in fuel cost expectations may also impact our tax comparison as we assess below.

²⁰We assume a vehicle’s maximum lifetime distance is $\beta^{\max} = 250,000\text{km}$, so that the vehicle lifetime of individual i is $S_i = \beta^{\max}/\beta_i^m$ (up to the nearest integer). Furthermore, we assume a scrapping rate ψ_s in year s (based on Table 1 of Jacobsen and van Benthem (2015) and setting $\psi_{20} = 1$). The present value of fuel costs over a car’s lifetime then becomes:

$$G_{ijk} = \sum_{s=1}^{S_i} (1+r)^{-s} (1-\psi_s) \beta_i^m e_{jk} g_{ks}.$$

price coefficient α . As discussed above, we can then use the expression of ρ given by (4) to draw conclusions about consumers' intertemporal preferences. More specifically, we follow Allcott and Wozny (2014), 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 ($\gamma > 1$). Alternatively, as in Hausman (1979), we can set $\gamma = 1$ and assume a time horizon S to retrieve the consumers' implicit interest rate r ; or we can set $\gamma = 1$ and $r = 0$ to retrieve the consumers' required payback time.

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 (2014) 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.²¹

The logit model implies that $\gamma\rho = 8.63$, 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 8.63 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 $\gamma = 0.89$, which implies a moderate undervaluation of future fuel cost savings. Finally, setting $\gamma = 1$, our estimates imply the consumers' implicit interest rate is 7.9% when they use a time horizon $S = 15$, and it is 2.8% when they use $S = 10$.

The random coefficients logit model with only mileage heterogeneity implies slightly more undervaluation: the required payback time ($\gamma\rho$) is 7.47 years, which corresponds to an attention weight parameter of $\gamma = 0.77$ and implicit interest rates in the range of 10.3% ($S = 15$) and 5.7% ($S = 10$).

Finally, the full model, which also allows for heterogeneity in the valuation of other product characteristics, results in a coefficient of $\gamma\rho = 8.84$, so that the attention weight parameter is $\gamma = 0.91$. This again implies some moderate (but statistically insignificant) undervaluation of future fuel cost savings.

In all models, the consumer surplus losses per vehicle from misoptimization are small: €73 in the first model, €328 in the second model, and €40 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, we find only moderate undervaluation with an attention weight of $\gamma = 0.91$ for European consumers. This is similar to Busse, Knittel, and Zettelmeyer (2013), and higher

²¹Allcott and Wozny (2014) 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.

than the $\gamma = 0.76$ obtained by Allcott and Wozny (2014). 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 and welfare impact 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 and the welfare impact of alternative tax policies?

As discussed in section 2, we discuss the impact of two tax policies: a fuel tax t_k^G (on gasoline and/or diesel fuel) and a product tax t_k^E (on a car’s fuel consumption 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 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 subsection 6.1 we compare the effectiveness of a fuel tax and a product tax in terms of their capability to reduce total fuel usage. To make this comparison, we start from a fuel tax increase by 50c per liter, which is about twice the size of the fuel tax of 1\$ per gallon often considered in several US studies. We compare this with a revenue-equivalent product tax per unit of fuel consumption e_{jk} (in liter/100km), i.e. a product tax that generates the same revenues as the capitalized expected revenues from the fuel tax increase by 50c.

In subsection 6.2 we compare the total welfare impact of the fuel tax and product tax. We first calculate the impact of a fuel tax increase on total welfare. We then compare this with an externality-equivalent product tax, which reduces the externality by the same amount.

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.

6.1 Effectiveness of fuel versus product taxes

In this section we are mainly interested in the effectiveness of both taxes in terms of their capability to reduce total fuel usage, but to provide background information we first discuss how both taxes affect the composition of new car sales and the sales-weighted average fuel consumption (where, as discussed earlier, fuel consumption is the inverse of fuel economy). As an extension we also compare the effects of discriminatory taxes by fuel type (gasoline versus diesel), and show how these help to explain the market share of diesel cars across European countries.

In section 2 we already discussed that the price sensitivity parameter α_i and the extent substitution to the outside good s_{i0} determine the general demand impact of both taxes, whereas the future valuation parameter γ and the extent of mileage heterogeneity β_i^m explain the differences in their impact. We now show how this can be translated empirically.

Effects on the composition of new car sales To evaluate how a fuel and product tax affect the composition of new car sales we take into account that these taxes may affect car prices, assuming multi-product Bertrand competition. In Appendix B we show the results when car prices are not affected (as under perfect competition). More specifically, we compute the effect of both taxes on the sales and prices of every car jk , and then aggregate these effects to market share and price effects of four fuel consumption quartiles. This is similar to what Busse, Knittel, and Zettelmeyer (2013) report for fuel taxes, 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, we can compare the impact of a fuel tax on the market shares with the impact of a (revenue-equivalent) product tax. Second, we are able to compute the market share effects not only under the actual estimate of the future valuation parameter ($\gamma = 0.91$), but also under alternative scenarios with full forward looking behavior ($\gamma = 1$) and strong consumer myopia ($\gamma = 0.5$). This more detailed comparison provides useful background intuition to interpret the different impact of both taxes on the average fuel consumption of the sold cars and on total fuel usage.

Table 4 shows the tax effects by fuel consumption quartile, based on the estimates of the full model with heterogeneity in both mileage and in the valuation of other product characteristics. According to the top panel of Table 4, under the actual future valuation parameter ($\gamma = 0.91$) a 50c fuel tax increase raises the market share of cars in the lowest fuel consumption quartile by 2.2 percentage points (from a market share of 36.6% to 38.8%). The market shares of the other fuel consumption quartiles all drop. In particular, the market share of cars in the highest fuel consumption quartile drops by 0.7 percentage points (or

12%). Hence, a fuel price increase implies a quite considerable change in the composition of car sales from the high to the low fuel consumption cars.²² If consumers would be fully forward looking ($\gamma = 1$), the fuel tax increase would only have a slightly larger impact on the composition of car sales. Conversely, with strong consumer myopia ($\gamma = 0.5$) the fuel tax increase would have had a smaller impact. Notice finally that the taxes have only a small impact on average car prices, implying that there is a high degree of tax pass-through. Because of this, the tax effects are similar under perfect competition, as shown in Appendix (Table B.6).²³

The bottom part of Table 4 compares these findings with those of a revenue-equivalent product tax increase. This amounts to a tax of €699 per unit of unit of fuel consumption e_{jk} (so this would amount to a product tax of €3,495 for a car that consumes 5 liter per 100km). Under the actual estimate of $\gamma = 0.91$, the product tax increase has a larger effect on the composition of new car sales than the fuel tax. This is also the case if consumers would be fully forward looking, and even more strongly so under strong consumer myopia ($\gamma = 0.5$).

Effects on sales-weighted average fuel consumption and total fuel usage Table 5 shows how these sales composition effects translate into the sales-weighted average fuel consumption and the total annual fuel usage. To make a clean comparison across the specifications, we set $\gamma = 1$ because for all models γ was indeed close to one. In the Appendix (Table B.7), we show the results under consumer myopia ($\gamma = 0.5$).

It is instructive to start with the results from the simple logit model without mileage and other heterogeneity (first two rows of Table 5). In this case, the fuel tax and the product tax have the same impact: they both improve average fuel consumption by 2.2% and reduce total fuel usage by 31.9%. This is as expected since consumers are forward looking and all have the same mileage. Note that the large magnitude of the total fuel usage reduction is due to an implausibly large substitution to the outside good in the simple logit model.

In the random coefficients logit with mileage heterogeneity (RC Logit I) and additional heterogeneity (RC Logit II), the fuel tax has a smaller impact on the sales-weighted average fuel consumption than the product tax. This is consistent with our earlier finding that a

²²These 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. This 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.

²³The high tax pass-through may be due to the assumption of constant marginal costs or the functional form of the demand model. However, if pass-through would in practice be lower (say because of increasing marginal costs), this would likely not differ much between the fuel and product taxes, given that there is only limited undervaluation.

fuel tax has a smaller impact on the composition of new car sales than a product tax under consumer heterogeneity. More importantly, despite the smaller impact on the cars' average fuel consumption, the fuel tax has a larger impact on total fuel usage. According to RC Logit II, the fuel tax reduces total fuel usage by 18.1%, whereas a product tax only reduces total fuel usage by 12.0%. Intuitively, this is because the fuel tax mainly targets the high mileage consumers, who are most likely to substitute to cars with a lower fuel consumption (better fuel economy).²⁴ This conclusion also holds under the actual estimate of $\gamma = 0.91$ (instead of $\gamma = 1$). However, with stronger consumer myopia the conclusion would be reversed, and a product tax becomes more effective in reducing total fuel usage. For example, if $\gamma = 0.5$, a product tax would reduce total fuel usage by 15.2%, compared with a reduction by only 12.0% under a fuel tax (see Table B.7 in Appendix).

The remaining rows of Table 5 show a sensitivity analysis, based on the estimates of extensions of the random coefficients models presented in Table B.3: RC Logit III (extra random coefficient on fuel type), RC Logit IV (extra heterogeneity in fuel cost valuation for other reasons than mileage), and RC Logit V (vehicle lifetime varying according to consumer's actual mileage). In all models, fuel taxes retain their advantage over product taxes in reducing total fuel usage. However, especially in RC Logit IV, the difference in impact between both taxes becomes less pronounced: the targeting effect is weaker because consumers do not only sort based on mileage but also on other factors.

In sum, in all models where we account for mileage heterogeneity, a fuel tax is more effective than a product tax in reducing total fuel usage (despite the fact that it performs worse in improving the average fuel consumption of the sold cars). Intuitively, this is because a fuel tax better targets the high mileage consumers and consumers are sufficiently forward looking to take these future fuel taxes into account. A product tax would only be more effective if there would be sufficiently strong consumer myopia.

Effects of discriminatory taxes by fuel type The above counterfactuals compared the impact of fuel and product taxes, without discriminating between the 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

²⁴This is confirmed from a breakdown by mileage quartile: the lowest mileage quartile reduces fuel usage by much less than the largest quartile (not shown in the table).

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.²⁵

Table 6 shows the effects of a 50c gasoline and/or diesel fuel tax for the full random coefficients logit model. A 50c gasoline fuel tax increase would increase the diesel market share by 13.7% points. This results in a disproportionate reduction of gasoline fuel usage by 40.5% and in an increase of diesel fuel usage by 24.2%. Conversely, a 50c diesel fuel tax increase would reduce the diesel market share by 12.3%, which results in a disproportionate reduction of diesel fuel usage by 37.8% and in an increase in gasoline fuel usage of 27.7%. There is a disproportionately greater impact on fuel usage than on market shares because especially the high mileage consumers substitute in response to the taxes. As a robustness check (Table B.8), we performed the same counterfactuals for RC Logit III (which includes a random coefficient for fuel type). We obtain similar conclusions, though the substitution between fuel types is lower. Finally, note that the shifts between fuel types are correspondingly much smaller in a logit model without mileage heterogeneity (not shown).

In sum, a discriminatory fuel tax can imply a considerable shift in usage between fuel types due to the presence of mileage heterogeneity. This finding may also be relevant for policies towards other fuels (such as methanol).

Do discriminatory taxes explain the diesel market share in Europe? We now take a different perspective, and ask to which extent the observed discriminatory fuel taxes and differences in fuel consumption 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 consumption 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 lower average fuel consumption of about 2 liter/100km in 1998, and this advantage further improved in 2011. The lower fuel price and lower fuel

²⁵Miravete, Moral, and Thürk (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.

consumption 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 consumption 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 7–11% lower (with the exception of the UK, where there was no fuel price gap). If in addition the fuel consumption of diesel and gasoline cars would be equalized, the market share of diesel cars would have been an additional 5–13% lower. In 1998, both factors together explain about 40% of the diesel market share in Belgium, the UK and Spain, and an even larger fraction of 53% in France and Italy, 67% in Germany and the Netherlands.²⁶ 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 consumption still explain about 30% of the diesel market share across countries.²⁷

6.2 Welfare effects

We now compare the welfare impact of a fuel tax and a product tax. Total welfare is the sum of tax revenues, consumer surplus and the environmental externalities from car usage.²⁸ Following Allcott (2013), when consumers may undervalue future fuel costs, consumer surplus is the sum of decision consumer surplus (the usual consumer surplus, at the level of the undervaluation parameter γ) and belief error (the difference between actual and perceived spending on fuel costs); see Appendix A.3 for details.

The impact of the taxes on tax revenues and on consumer surplus can be directly computed from our demand estimates. To compute the impact on the environmental externalities we proceed as follows. First, we calculate the externality per liter of fuel that justifies the current fuel tax levels (based on the first-order conditions for welfare maximization), and we find this to be in the range of 1€ to 1.04€ per liter, depending on the specification.²⁹ Next,

²⁶Note 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.

²⁷In the Appendix (Table B.9) we show the results in an extended model with an additional random coefficient on fuel type (RC logit III). The results are qualitatively the same, but the magnitudes are smaller because consumers substitute less between fuel types.

²⁸To simplify the exposition, we assume car prices are set competitively, so that there is no producer surplus. Our results are similar if we consider multi-product Bertrand-Nash pricing with constant marginal costs.

²⁹This is higher but of comparable order of magnitude than outside estimates of externality costs, which attempt to directly estimate the cost of CO₂ and other emissions, congestion costs and accident externalities. For example, a recent IMF report (Parry et al. (2014)) estimates the total externality to be \$0.58 for gasoline

we assume the perceived public awareness of the externality per liter increases, and then compare the impact of a fuel increase with the impact of an “externality-neutral” product tax increase, i.e. a product tax increase that achieves the same level of externality reduction.

Table 8 shows the results. We provide a comparison between the simple logit benchmark and the base specification (RC logit II), which accounts for mileage and other consumer heterogeneity. The top panel considers a first scenario: the impact of a tax increase after a general increase in the externality by 20 cents (say because of an increase in car congestion). In the logit model without mileage heterogeneity, a fuel tax and a product tax perform equally well: they would both raise welfare by 265 million Euro. There is however a difference in the sources behind these gains. Under a fuel tax there is a larger increase in tax revenues than under a product tax (+1565 million Euro versus +954 million Euro). But at the same time consumers experience extra losses from their belief error (the so-called internality by Allcott (2013)), while these losses go down under a product tax. Intuitively, a fuel tax raises more tax revenues, but consumers are partly unaware of these extra taxes and the overall impact of both taxes is the same.

In the random coefficients model a fuel tax performs considerably better than an externality neutral product tax in correcting for the increased externality: it raises total welfare by 186 million Euro, versus only by 112 million Euro for the product tax. This is essentially due to the much lower impact of a fuel tax on decision consumer surplus (which decreases by only 3.4 billion Euro versus 5.3 billion Euro under a product tax) because a lower tax amount is required to achieve the same externality reduction when there is mileage heterogeneity. Intuitively, this is because fuel taxes directly target the high mileage consumers, while product taxes affect all consumers in the same way.

The bottom panel of Table 8 considers an alternative scenario: the impact of a diesel tax increase after an increase in the externality of diesel cars. This reflects current policy awareness that diesel cars are not less polluting than gasoline cars (in terms of CO₂ emissions), but actually are likely more polluting (because of NO_x and other emissions). To keep the analysis simple, we separately calculate the externality per usage of gasoline and diesel that rationalize the current gasoline and diesel taxation and then we raise the externality usage of diesel to the same level of gasoline. The results give a similar picture: taxing diesel fuel or diesel cars has the same impact on welfare in the absence of mileage heterogeneity, but taxing diesel fuel is much more preferable than taxing diesel cars under mileage heterogeneity.

In sum, these findings show that it is generally preferable to impose fuel taxes than product taxes to correct for usage externalities, even when usage demand itself is perfectly

and \$0.82 for diesel in Germany 2010, and \$0.72 for gasoline and \$0.95 for diesel in France.

inelastic.

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 consumption across products.

We find evidence of 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 consumption). 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 consumption, because it specifically targets the high mileage consumers to substitute to cars with a lower fuel consumption. Discriminatory taxes by fuel type lead to disproportionately larger shifts in fuel type usage than in market shares, because they especially target high mileage consumers. Furthermore, we find that fuel taxes perform much better in terms of total welfare to achieve a same level of externality reduction.

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.


8 Figures and Tables

Figure 1: Gasoline and Diesel Prices across time and countries



The top panel of the figure shows average yearly prices (in real 2005 €) of gasoline and diesel between 1998 and 2011. The bottom panel shows the difference in gasoline and diesel prices for each of the countries between 1998 and 2011. Source: DG ECFIN.

Figure 2: UK mileage survey



This figure shows the frequency distribution for the UK mileage survey in 2007. The left panel has x-axis bins in 1000 of km travelled per year, the mean of the distribution is 14 761km travelled. The right panel shows the distribution with equal sized 5000km bins and a plotted kernel.

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
Fuel consumption (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	82,151		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
Fuel consumption (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 consumption 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	
	Est.	St.Err.	Est.	St.Err.	Est.	St.Err.
	Mean valuations					
Price/Inc. (α)	-4.52	0.19	-6.22	0.22	-5.33	0.21
Fuel Costs/Inc. ($\alpha\gamma\rho$)	-39.03	1.41	-46.48	0.94	-47.11	9.22
Power (kW/100)	2.28	0.14	2.60	0.17	0.25	0.61
Size (cm ² /10,000)	13.25	0.44	16.69	0.48	16.77	2.02
Height (cm/100)	3.00	0.30	4.45	0.32	5.19	0.33
Foreign	-0.83	0.02	-0.75	0.02	-0.89	0.04
	Standard Deviations of valuations					
Power (kW/100)	-	-	-	-	1.95	0.25
Size	-	-	-	-	4.31	2.04
Foreign	-	-	-	-	0.49	0.43
Mileage distribution	No		Yes		Yes	
	Valuations of Future Fuel Costs					
Fuel Costs/Price ($\gamma\rho$)	8.63	0.55	7.47	0.24	8.84	1.77
Future Valuation γ ($r = 6\%$)	0.89	0.06	0.77	0.02	0.91	0.18
Consumer Loss from Misoptim. (€)	73.07		328.13		39.71	
Implicit Interest rate ($T = 10$)	2.77		5.69		2.31	
Implicit Interest rate ($T = 15$)	7.87		10.32		7.48	

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 82,151, 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 due to misoptimization per vehicle in €; (iv) the implicit interest rate that sets $\gamma = 1$ assuming a lifetime of 10 and 15 years.

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

Change in Market Share	Current	$\gamma = 0.50$		$\gamma = 0.91$		$\gamma = 1.00$	
		% Point	€ Δp	% Point	€ Δp	% Point	€ Δp
Fuel Tax							
Tax per liter		€0.50		€0.50		€0.50	
Fuel Consumption Q1 (lowest)	37	1.5	-15	2.2	-31	2.4	-39
Fuel Consumption Q2	37	-0.3	-14	-0.4	-30	-0.5	-36
Fuel Consumption Q3	20	-0.6	-15	-1.0	-29	-1.1	-32
Fuel Consumption Q4 (highest)	6	-0.5	-16	-0.7	-21	-0.8	-23
Revenue Equivalent Product Tax							
Tax per liter/100km		€832		€699		€662	
Fuel Consumption Q1 (lowest)	37	3.8	-41	3.0	-37	2.9	-37
Fuel Consumption Q2	37	-0.5	-47	-0.4	-44	-0.4	-43
Fuel Consumption Q3	20	-1.7	-52	-1.4	-47	-1.4	-45
Fuel Consumption Q4 (highest)	6	-1.5	-52	-1.2	-44	-1.1	-41

The table reports the effects of a fuel tax and a revenue-equivalent product tax on the market shares and prices of new cars aggregated by quartile of fuel consumption, under Bertrand-Nash pricing. The effects are computed under strong consumer myopia ($\gamma = 0.50$), under the actual estimate of consumer valuation of fuel costs ($\gamma = 0.91$) and under full forward looking behavior ($\gamma = 1.00$), on the basis of the parameter estimates of RC Logit 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 Consumption and Fuel Usage

	Outside Good % Point Change	Fuel Consumption % Change	Fuel Usage % Change
Logit			
Fuel Tax	14.06	-2.16	-31.89
Revenue Eq. Product Tax	14.06	-2.16	-31.89
RC Logit I - Mileage Only			
Fuel Tax	10.78	-0.80	-38.19
Revenue Eq. Product Tax	7.33	-1.49	-22.97
RC Logit II			
Fuel Tax	5.71	-1.16	-18.13
Revenue Eq. Product Tax	4.84	-1.53	-12.00
RC Logit III - Extra RC on fuel cost			
Fuel Tax	5.71	-0.95	-15.91
Revenue Eq. Product Tax	6.79	-1.56	-14.83
RC Logit IV - Extra heterog. in fuel cost			
Fuel Tax	3.78	-1.85	-9.52
Revenue Eq. Product Tax	3.74	-1.96	-8.96
RC Logit V - Lifetime varies with miles			
Fuel Tax	6.23	-1.17	-16.20
Revenue Eq. Product Tax	5.87	-1.29	-12.78

The table reports the effect of a fuel tax and a revenue-equivalent product tax on the share of the outside good, average fuel consumption, and total annual fuel usage. The simulations are based on the parameter estimates in Table 3 and Table B.3. In each model we set $\gamma = 1$, so there is no variation in undervaluation across the different models. 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	-9.7	-40.5	24.2	13.7
Fuel Tax, Diesel only	-3.5	27.7	-37.8	-12.3
Fuel Tax, both Gas. and Diesel	-16.8	-21.8	-11.4	2.4

The table reports the effect of a discriminatory fuel tax for gasoline and diesel cars on total fuel usage. 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 Cons. Gap	Diesel Share	Fuel Price	Fuel Cons.	Both
			Year 1998			
Belgium	-0.25	-1.91	52%	-10%	-11%	-21%
France	-0.26	-2.04	38%	-11%	-9%	-20%
Germany	-0.22	-2.35	15%	-5%	-5%	-10%
Great Britain	0.00	-2.24	15%	0%	-6%	-6%
Italy	-0.19	-2.08	21%	-7%	-6%	-13%
Netherlands	-0.29	-2.12	21%	-8%	-6%	-14%
Spain	-0.11	-1.85	51%	-9%	-13%	-21%
			Year 2011			
Belgium	-0.24	-2.01	75%	-4%	-16%	-20%
France	-0.19	-1.81	69%	-4%	-13%	-16%
Germany	-0.16	-2.23	46%	-4%	-14%	-18%
Great Britain	0.04	-2.04	50%	2%	-14%	-12%
Italy	-0.13	-2.05	56%	-3%	-14%	-18%
Netherlands	-0.28	-2.22	30%	-5%	-12%	-15%
Spain	-0.07	-2.08	69%	-2%	-15%	-17%

The table reports: (i) in the first three columns, the currently observed gaps in fuel prices and fuel consumption 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 consumption gap were eliminated. The simulations are based on the parameter estimates of RC Logit II in Table 3.

Table 8: The Welfare Effects of a Fuel Tax versus a Product Tax

Change in Mln Euro:	Revenue	Consumer surplus	Belief error (internality)	Externality	Welfare
Panel A: general externality increase and corrective taxes					
Logit ($\gamma = 0.89$)					
Fuel Tax	1565	-335	-116	2165	265
Product Tax	954	-335	496	2165	265
RC Logit II ($\gamma = 0.91$)					
Fuel Tax	2246	-3409	-234	1584	186
Product Tax	3512	-5294	310	1584	112
Panel B: diesel externality increase and corrective diesel taxes					
Logit ($\gamma = 0.89$)					
Fuel Tax	1009	-1389	-107	632	146
Product Tax	761	-1389	141	632	146
RC Logit II ($\gamma = 0.91$)					
Fuel Tax	1555	-1547	-185	364	187
Product Tax	2579	-2923	30	364	50

The table reports the effects of a fuel tax and an externality-neutral product tax on welfare. Welfare changes are the sum of changes in tax revenues, decision consumer surplus, belief error, and the environmental externality. Panel A considers the effect of (i) an increase in the price of fuel by 20 cents; (ii) an increase in the externality by 20 cents above the level that justifies the current fuel tax. Panel B considers the effect of (i) an increase in the price of diesel by 20 cents; (ii) an increase in the perceived externality of diesel up to the level of that justifies the current gasoline tax. The simulations are based on the parameter estimates of RC Logit II in Table 3. The figures refer to Germany in 2011.

References

- Adamou, A., S. Clerides, and T. Zachariadis (2014). Welfare implications of car feebates: A simulation analysis. *The Economic Journal* 124(578), F420–F443.
- Allcott, H. (2013, August). The welfare effects of misperceived product costs: Data and calibrations from the automobile market. *American Economic Journal: Economic Policy* 5(3), 30–66.
- Allcott, H. and M. Greenstone (2012). Is there an energy efficiency gap? *The Journal of Economic Perspectives*, 3–28.
- Allcott, H., S. Mullainathan, and D. Taubinsky (2014, April). Energy policy with externalities and internalities. *Journal of Public Economics* 112.
- Allcott, H. and N. Wozny (2014). Gasoline prices, fuel economy, and the energy paradox. *The Review of Economics and Statistics* 96(5), 779–795.
- Alquist, R., L. Kilian, and R. J. Vigfusson (2013). Forecasting the price of oil. In G. Elliott and A. Timmermann (Eds.), *Handbook of Economic Forecasting* 2, pp. 427–507. Amsterdam: North-Holland.
- Anderson, S. T., R. Kellogg, and J. M. Sallee (2013). What do consumers believe about future gasoline prices? *Journal of Environmental Economics and Management* 66 (3), 383–403.
- Baltagi, B. H. (1995). *Econometric Analysis of Panel Data*. John Wiley & Sons.

- Bento, A. M., L. H. Goulder, M. R. Jacobsen, and R. H. von Haefen (2009, June). Distributional and efficiency impacts of increased us gasoline taxes. *American Economic Review* 99(3), 667–99.
- Bento, A. M., S. Li, and K. Roth (2012). Is there an energy paradox in fuel economy? a note on the role of consumer heterogeneity and sorting bias. *Economics Letters* 115(1), 44 – 48.
- Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics* 25(2), 242–262.
- Berry, S. T., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica* 63(4), 841–890.
- Berry, S. T., J. Levinsohn, and A. Pakes (1999, June). Voluntary export restraints on automobiles: Evaluating a trade policy. *American Economic Review* 89(3), 400–430.
- Busse, M. R., C. R. Knittel, and F. Zettelmeyer (2013, February). Are consumers myopic? evidence from new and used car purchases. *American Economic Review* 103(1), 220–56.
- Chamberlain, G. (1987). Asymptotic efficiency in estimation with conditional moment restrictions. *Journal of Econometrics* 34(3), 305 – 334.
- Chugh, R., M. Cropper, and U. Narain (2011). The cost of fuel economy in the Indian passenger vehicle market. *Energy Policy* 39(11), 7174–7183.
- D’Haultfoeuille, X., P. Givord, and X. Boutin (2014). The environmental effect of green taxation: The case of the french bonus/malus. *The Economic Journal* 124(578), F444–F480.

- Gillingham, K. (2012). Selection on anticipated driving and the consumer response to changing gasoline prices. Technical report, Yale University.
- Goldberg, P. K. (1998, March). The effects of the corporate average fuel efficiency standards in the us. *Journal of Industrial Economics* 46(1), 1–33.
- Goldberg, P. K. and R. Hellerstein (2013). A structural approach to identifying the sources of local currency price stability. *Review of Economic Studies* 80(1), 175–210.
- Goldberg, P. K. and F. Verboven (2005, January). Market integration and convergence to the Law of One Price: evidence from the European car market. *Journal of International Economics* 65(1), 49–73.
- Greene, D. L. (2010, March). How consumers value fuel economy: A literature review. Technical Report EPA-420-R-10-008, U.S. Environmental Protection Agency.
- Hausman, J. A. (1979, Spring). Individual discount rates and the purchase and utilization of energy-using durables. *Bell Journal of Economics* 10(1), 33–54.
- Helfand, G. and A. Wolverson (2011). Evaluating the consumer response to fuel economy : a review of the literature. *International review of environmental and resource economics* 5, 103–146.
- Hughes, J. E., C. R. Knittel, and D. Sperling (2008). Evidence of a shift in the short-run price elasticity of gasoline demand. *The Energy Journal* 29(1), 113–134.
- Huse, C. and C. Lucinda (2014). The market impact and the cost of environmental policy: Evidence from the swedish green car rebate. *The Economic Journal* 124(578), F393–F419.

- Jacobsen, M. R. (2013). Evaluating us fuel economy standards in a model with producer and household heterogeneity. *American Economic Journal: Economic Policy* 5(2), 148–87.
- Jacobsen, M. R., C. R. Knittel, J. M. Sallee, and A. A. van Benthem (2016, March). Sufficient statistics for imperfect externality-correcting policies. Working Paper 22063, National Bureau of Economic Research.
- Jacobsen, M. R. and A. A. van Benthem (2015, March). Vehicle scrappage and gasoline policy. *American Economic Review* 105(3), 1312–38.
- Jaffe, A. B. and R. N. Stavins (1994, May). The energy paradox and the diffusion of conservation technology. *Resource and Energy Economics* 16(2), 91–122.
- Klier, T. and J. Linn (2010, August). The price of gasoline and new vehicle fuel economy: Evidence from monthly sales data. *American Economic Journal: Economic Policy* 2(3), 134–53.
- Klier, T. and J. Linn (2015). Using taxes to reduce carbon dioxide emissions rates of new passenger vehicles: Evidence from france, germany, and sweden. *American Economic Journal: Economic Policy* 7(1), 1–32.
- Langer, A. and N. Miller (2013). Automakers’ short-run responses to changing gasoline prices. *Review of Economics and Statistics* 95(4), 1198–1211.
- Li, S., C. Timmins, and R. H. von Haefen (2009, August). How do gasoline prices affect fleet fuel economy? *American Economic Journal: Economic Policy* 1(2), 113–37.
- Linn, J. (2014, April). Explaining the adoption of diesel fuel passenger cars in europe. Discussion Paper 14-08, Resources for the Future.

- Miravete, E., M. J. Moral, and J. Thürk (2014, April). Protecting the european automobile industry through environmental regulation: The adoption of diesel engines.
- Parry, I., D. Heine, E. Lis, and S. Li (2014). *Getting Energy Prices Right: From Principle to Practice*. INTERNATIONAL MONETARY FUND.
- Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy* 110, 705–29.
- Reynaert, M. and F. Verboven (2014). Improving the performance of random coefficients demand models: The role of optimal instruments. *Journal of Econometrics* 179(1), 83 – 98.
- Sallee, J. (2014). Rational inattention and energy efficiency. *Journal of Law and Economics* 57(3), 781 – 820.
- Sallee, J. M., S. E. West, and W. Fan (2016). Do consumers recognize the value of fuel economy? evidence from used car prices and gasoline price fluctuations. *Journal of Public Economics* 135, 61 – 73.
- Verboven, F. (2002, Summer). Quality-based price discrimination and tax incidence: Evidence from gasoline and diesel cars. *RAND Journal of Economics* 33(2), 275–297.

A Appendix A. Computational details

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. Finally, we describe how we incorporate the empirical millage distribution in our estimation.

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 a 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 tax effect 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

In this Appendix, 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 market shares, tax revenues, average fuel consumption and total fuel usage, and the various welfare components: tax revenues, decision consumer surplus, belief error (internality) and the externality.

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 the 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 jk under taxes (t^G, t^E) are again:

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

So the predicted quantity after a change in the fuel tax by Δ^G is $q_{jk} (t^G + \Delta^G, t^E)$ and the predicted quantity after a change in the product tax by Δ^E is $q_{jk} (t^G, t^E + \Delta^E)$. Based on these predicted quantities per product jk we can compute the market shares per fuel consumption quartile (or any other aggregated quantity or market share).

Tax revenues Conditional on buying product jk , an individual consumer pays 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:

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

i.e. the expected tax revenue over all cars per consumer, averaged over all consumers. We can then compute the tax revenues from a change in energy taxes $R (t^G + \Delta^G, t^E)$ or product

taxes $R(t^G, t^E + \Delta^E)$ as:

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

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

With uniform taxes, we consider an increase in the energy tax by 50c, $\Delta^G = 0.5$, so that the revenue-neutral product tax is the solution of Δ^E to $R(t^G + 0.5, t^E) = R(t^G, t^E + \Delta^E)$.

Sales-weighted average fuel consumption and total energy usage Weighted average fuel consumption (given that people purchase a car) is

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

where

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

Total annual fuel usage is given by

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

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 consumption and the percentage change in fuel usage for both taxes. For example, the percentage change in average fuel consumption after a fuel tax change by Δ^G is $\% \Delta E = E(t^G + \Delta^G, t^G) / E(t^G, t^E) - 1$.

Welfare The first welfare component consists of tax revenues, already derived above. The second welfare component is experienced consumer surplus. Following Allcott (2013), this is

equal to decision consumer surplus minus belief error. Decision consumer surplus is given by the standard expression in logit models (aggregated over the distribution of consumer valuations):

$$CS(t^G, t^E) = \int_{\beta} \frac{1}{\alpha_i} \ln \left(1 + \sum_j \sum_k \exp(v_{ijk}) \right) dF_{\beta}(\beta) I,$$

and belief error is the difference between the actual and perceived fuel expenditures:

$$CS^b(t^G, t^E) = (1 - \gamma) \int_{\beta} \sum_j \sum_k \rho \beta_i^m (g_k + t_k^G) e_{jk} s_{ijk} (t_k^G, t_k^E; \beta) dF_{\beta}(\beta) I.$$

Experienced consumer surplus is then $CS^*(t^G, t^E) = CS(t^G, t^E) - CS^b(t^G, t^E)$.

The final welfare component is the externality. Let h_k be the externality cost per liter of fuel type k (which includes the CO₂ costs, but also other pollution costs, congestion and accident externalities). The externality is then given by:

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

Total welfare is the sum of the various components:

$$W(t^G, t^E) = R(t^G, t^E) + CS(t^G, t^E) - CS^b(t^G, t^E) + EXT(t^G, t^E).$$

To perform our welfare analysis, we need a measure of the externality cost per liter h_k for each fuel type k . We compute both h_k such that welfare is maximized at the current fuel tax levels, i.e. such that $\partial W / \partial t_k^G = 0$ for each fuel type k . We then assume that h_k increases by a certain amount, and compare the welfare impact of a fuel tax and an externality-neutral product tax, as discussed in the text.

A.3 Incorporating the empirical mileage distribution

In both estimation and counterfactuals we make use of the empirical mileage distribution.

Suppose there are R mileage types, where for each type $r = 1 \dots R$ there is a fraction λ_r (such that $\sum_{r=1}^R \lambda_r = 1$) and a corresponding mileage β_r^m . Total sales for product jk are

$$q_{jk} = \underbrace{\sum_{i=1}^R \lambda_i s_{rjk}(\beta_r^m)}_{s_j} I,$$

where $s_{rjk}(\beta_r^m)$ is the probability that mileage type r chooses product jk (i.e. the integral of the individual choice probabilities over all heterogeneity other than mileage).

In practice, we do not observe the unconditional mileage fractions λ_r , but rather the fractions l_r conditional on buying a car. Based on these observed conditional fractions, total sales for product jk can also be written as

$$q_{jk} = \sum_{i=1}^R l_i s_{rjk|B}(\beta_r^m) Q,$$

where $s_{rjk|B}(\beta_r^m) = s_{rjk}(\beta_r^m) / \sum_{j=1}^J \sum_k s_{rjk}(\beta_r^m)$ is the conditional probability that mileage type r chooses jk , conditional on buying a car. from equating the above two expressions, it follows that

$$\lambda_i = l_i \cdot (Q/I) / \sum_{j=1}^J \sum_k s_{rjk}(\beta_i^m),$$

where $Q = \sum_{j=1}^J \sum_k q_{jk}$. We use this to write all expressions in terms of the observed conditional mileage fractions.

B Appendix B. Additional Tables

Table B.1: Alternative Instruments and First Stage

	Full Sample						Cost Sample			
	No IV		Sums IV		Sums IV		Cost IV		All IV's	
	Est.	St.Er.	Est.	St.Er.	Est.	St.Er.	Est.	St.Er.	Est.	St.Er.
Price/Inc. (α)	-0.03	0.03	-4.52	0.19	-4.39	0.18	-3.09	0.22	-3.96	0.16
Fuel Cost/Inc.	-53.90	1.02	-39.03	1.41	-41.67	1.29	-45.25	1.28	-42.86	1.26
Power	-7.18	0.29	2.28	0.14	2.37	0.13	1.45	0.16	2.06	0.12
Size	12.9	0.31	13.25	0.44	15.29	0.39	14.52	0.38	15.03	0.38
Height	2.56	0.23	3.00	0.30	4.25	0.28	3.80	0.27	4.10	0.28
Foreign	-1.06	0.02	-0.83	0.02	-0.86	0.02	-0.92	0.02	-0.88	0.02
# Obs.	82 151						76 634			
First Stage - Excluded Instruments										
Sums of all other vehicles' characteristics:										
Fuel Cost/Inc.			-1.23	0.11	-1.25	0.11			-0.92	0.12
Power			-2.02	0.61	-2.00	0.66			-1.96	0.65
Size			3.09	0.35	-3.04	0.37			2.61	0.38
Height			3.66	3.14	4.03	3.35			9.12	3.26
# products			-2.24	0.51	-2.24	0.57			-2.27	0.57
Sums of own firm other vehicles characteristics:										
Fuel Cost/Inc.			2.72	0.43	2.61	0.45			2.52	0.44
Power			-2.80	0.39	-2.83	0.41			-2.62	0.38
Size			8.65	1.84	8.49	1.89			7.96	1.75
Height			7.03	1.81	7.11	1.91			7.04	1.86
# products			-1.47	0.31	-1.46	0.32			-1.43	0.31
Cost Shifters										
Labor Costs							4.08	0.56	2.28	0.43
Local Prod.							-1.87	0.93	-1.69	0.88
We*Steel Pr.							-3.33	0.62	-1.69	0.62
Joint F Stat Exc.Instr. (p-val)										
			291	(0.0)	273	(0.0)	531	(0.0)	237	(0.0)

The table reports the parameter estimates and standard errors for the first stage and second stage of the IV regression. The first panel gives the estimated demand parameters for the logit model, the other panels gives first stage estimates for excluded instruments. Column I gives results for the OLS. Column II replicates the logit model presented in the main paper. Columns III-V give results using a smaller sample for which we obtained cost shifter instrument: labor costs, local production and steel prices. Column III uses the same instruments as the main specification in the paper but in the smaller sample. Column IV shows the results when only relying on cost shifter instruments. Column V includes both cost shifters and the instrument set used in the main paper. Note that all regressions include model fixed effects, the cost shifters thus only vary when versions are produced in different plants, or when a model moves between production plants over time.

Table B.2: Parameter Estimates for Alternative Demand Models Part I

	Logit Av. FC		Logit Futures		Logit Eurostat		RC Logit Eurostat	
	Est.	St.Err.	Est.	St.Err.	Est.	St.Err.	Est.	St.Err.
	Mean valuations							
Price/Inc. (α)	-2.90	0.17	-4.14	0.18	-4.19	0.22	-5.96	0.25
Fuel Costs/Inc. ($\alpha\gamma\rho$)			-44.16	1.21	-56.79	1.13	-48.10	1.82
Av. Fuel Cost ($\alpha\gamma\rho$)	-28.17	1.13						
Power (kW/100)	1.04	0.12	2.21	0.13	2.41	0.16	-0.77	0.34
Size (cm ² /10,000)	14.64	0.35	15.40	0.37	15.34	0.37	19.87	0.47
Height (cm/100)	3.63	0.25	3.86	0.02	3.92	0.27	5.99	0.33
Foreign	-0.92	0.02	-0.85	0.02	-0.84	0.02	-1.85	0.08
	Standard Deviations of valuations							
Power (kW/100)	-	-	-	-	-	-	-2.67	0.10
Size (cm ² /10,000)	-	-	-	-	-	-	0.17	1.18
Foreign	-	-	-	-	-	-	4.74	0.26
Mileage distribution	No		No		No		Recentered	
	Valuations of Future Fuel Costs							
Fuel Costs/Price ($\gamma\rho$)	9.71	0.84	10.75	0.62	13.54	0.71	8.08	0.41
Future Valuation γ	1.00	0.09	1.11	0.06	1.39	0.07	0.83	0.04

The table reports the parameter estimates and standard errors for the different demand models. The Table presents similar results as Table 3 but with different definitions of fuel costs. In Logit Av. FC, we leave out all variation in fuel prices over time and estimate the model with average fuel costs per market. In Logit Futures, we compute fuel costs by rescaling fuel spot prices by the percentage difference in crude oil spot price and the futures prices from NYMEX (taking an average of futures prices over 1 up to 15 years). In Logit Eurostat we change the mean mileage to the mean mileage reported by Eurostat for each country. In RC Logit Eurostat we introduce heterogeneous mileage, with Eurostat means and a rescaling of the UK mileage distribution around that mean. Each specification includes model, market/diesel and market/time controls. The total number of observations (combinations of model/engine/market) is 82,151, where markets refer to 7 countries and 14 years.

Table B.3: Parameter Estimates for Alternative Demand Models Part II

	RC Logit III		RC Logit IV		RC Logit V	
	Est.	St.Err.	Est.	St.Err.	Est.	St.Err.
	Mean valuations					
Price/Inc. (α)	-5.25	0.20	-5.91	0.20	-4.82	1.32
Fuel Costs/Inc. ($\alpha\gamma\rho$)	40.59	6.27	-57.32	3.23	-	-
ρ_i *Fuel Costs/Inc. ($\alpha\gamma$)	-	-	-	-	-4.61	1.27
Power (kW/100)	-0.68	1.50	-0.52	0.10	-2.27	1.06
Size (cm ² /10,000)	17.46	1.22	19.56	0.48	18.26	1.98
Height (cm/100)	5.22	0.33	5.81	0.32	5.44	1.03
Foreign	-0.91	0.02	-1.14	0.04	-0.97	0.11
	Standard Deviations of valuations					
Power (kW/100)	2.33	0.07	2.56	0.06	2.98	0.91
Size (cm ² /10,000)	3.62	1.36	0.70	0.15	3.30	0.68
Foreign	0.21	0.36	3.04	0.24	0.47	2.30
Fuel type	1.65	0.40	-	-		
Mileage distribution	Yes		Yes		Yes	
	Valuations of Future Fuel Costs					
Fuel Costs/Price ($\gamma\rho$)	7.74	1.35	9.71	0.51	-	-
Future Valuation γ ($r = 6\%$)	0.80	0.14	1.00	0.05	0.96	0.52

The table reports the parameter estimates and standard errors for the different demand models. The Table presents similar results as Table 3 but with different assumptions. RC Logit III introduces a random coefficient on fuel type. RC Logit IV increases the variance of the mileage distribution across individuals. RC Logit V allows vehicle lifetime to depend on a consumer's annual mileage. Each specification includes model, market/diesel and market/time controls. The total number of observations (combinations of model/engine/market) is 82,151, where markets refer to 7 countries and 14 years.

Table B.4: Parameter Estimates for Alternative Demand Models Part II

	Implied γ	St.Err.	Implied γ	St.Err.
	$S = 10$		$S = 15$	
Logit				
r = 0%	0.91	0.07	0.61	0.04
r= 3%	1.04	0.08	0.74	0.05
r= 6%	1.17	0.09	0.89	0.06
r = 10%	1.35	0.10	1.09	0.07
RC Logit I - Mileage only				
r = 0%	0.79	0.03	0.53	0.02
r= 3%	0.90	0.03	0.64	0.02
r= 6%	1.01	0.04	0.77	0.02
r = 10%	1.17	0.04	0.95	0.03
RC Logit II				
r = 0%	0.94	0.19	0.62	0.13
r= 3%	1.07	0.21	0.76	0.15
r= 6%	1.20	0.24	0.91	0.18
r = 10%	1.39	0.28	1.12	0.22
RC Logit III on fuel type				
r = 0%	0.82	0.14	0.55	0.10
r= 3%	0.93	0.16	0.67	0.12
r= 6%	1.05	0.18	0.80	0.14
r = 10%	1.21	0.21	0.98	0.17
RC Logit IV - Add noise on fuel cost				
r = 0%	1.03	0.05	0.69	0.04
r= 3%	1.17	0.06	0.84	0.04
r= 6%	1.32	0.07	1.00	0.05
r = 10%	1.52	0.08	1.23	0.06

The table reports implied attention weights γ for different estimates and varying assumptions for the interest rate r and vehicle lifetime S .

Table B.5: Parameter Estimates Country by Country

	Belgium		France		Germany		Great Britain	
	Estimate	St.Error	Estimate	St.Error	Estimate	St.Error	Estimate	St.Error
	Mean valuations							
Price/Inc.	-6.62	0.53	-4.47	0.56	-6.13	0.75	-4.00	0.31
Fuel Costs/Inc.	-23.2	3.53	-43.49	2.94	-86.74	3.20	-26.53	2.41
Power (kW/100)	2.95	0.36	2.70	0.41	4.01	0.54	1.19	0.17
Size (cm ² /10,000)	8.55	0.69	6.66	0.83	1.37	0.09	0.79	0.08
Height (cm/100)	2.37	0.71	3.49	0.66	3.72	0.69	2.08	0.62
	Italy		Netherlands		Spain			
Price/Inc.	-9.43	0.80	-6.8	0.41	-2.19	0.30		
Fuel Costs/Inc.	-51.06	3.75	-44.24	3.77	-56.16	2.84		
Power (kW/100)	6.77	0.62	3.84	0.31	1.47	0.29		
Size (cm ² /10,000)	2.22	0.13	1.66	0.10	2.92	0.95		
Height (cm/100)	39.2	0.91	4.19	0.70	9.09	0.78		

The table reports the results from estimating logit models per country, parallel to the logit model for all countries reported in Table 3.

Table B.6: The Effect of a Fuel Tax and a Product Tax on Market Shares by Fuel Consumption Quartile

Change in Market Share	Current	$\gamma = 0.50$		$\gamma = 0.91$		$\gamma = 1.00$	
		% Point	€ Δp	% Point	€ Δp	% Point	€ Δp
Fuel Tax							
Tax per liter		€0.50		€0.50		€0.50	
Fuel Consumption Q1 (lowest)	37	1.4	0	2.1	0	2.4	0
Fuel Consumption Q2	37	-0.2	0	-0.4	0	-0.5	0
Fuel Consumption Q3	20	-0.6	0	-0.9	0	-1.1	0
Fuel Consumption Q4 (highest)	6	-0.6	0	-0.7	0	-0.8	0
Revenue Equivalent Product Tax							
Tax per liter/100km		€834		€700		€663	
Fuel Consumption Q1 (lowest)	37	3.6	0	3.1	0	2.9	0
Fuel Consumption Q2	37	-0.3	0	-0.4	0	-0.5	0
Fuel Consumption Q3	20	-1.6	0	-1.4	0	-1.4	0
Fuel Consumption Q4 (highest)	6	-1.7	0	-1.2	0	-1.1	0

The table reports the same results as in Table 4 for full pass through (we do not let firms change prices in response to the tax).

Table B.7: The Effect of a Fuel Tax and a Product Tax on Fuel Consumption and Fuel Usage ($\gamma = 0.5$)

	Outside Good % Point Change	Fuel Consumption % Change	Fuel Usage % Change
Logit			
Fuel Tax	7.19	-1.45	-12.01
Revenue Eq. Product Tax	18.71	-3.37	-30.33
RC Logit I - Mileage Only			
Fuel Tax	10.89	-0.77	-29.07
Revenue Eq. Product Tax	17.74	-3.17	-36.08
RC Logit II			
Fuel Tax	4.12	-0.77	-12.01
Revenue Eq. Product Tax	7.09	-2.08	-15.21
RC Logit III - Extra RC on fuel cost			
Fuel Tax	3.98	-0.79	-10.43
Revenue Eq. Product Tax	8.00	-1.86	-15.65
RC Logit IV - Extra heterog. in fuel cost			
Fuel Tax	3.05	-0.99	-6.11
Revenue Eq. Product Tax	6.29	-2.18	-11.39
RC Logit V - Lifetime varies with miles			
Fuel Tax	3.45	-0.74	-7.70
Revenue Eq. Product Tax	7.18	-1.61	-13.26

The table reports the effect of a fuel tax and a revenue-equivalent product tax on the share of the outside good, average fuel consumption, and total annual fuel usage. The simulations are based on the parameter estimates in Table 3 and Table B.3. In each model we set $\gamma = 0.5$, so there is no variation in undervaluation across the different models. The figures refer to Germany in 2011.

Table B.8: 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	-9.9	-33.2	17.1	10.0
Fuel Tax, Diesel only	-3.5	17.5	-27.8	-8.0
Fuel Tax, both Gas. and Diesel	-15.9	-20.7	-10.3	2.4

The table reports the effect of a discriminatory fuel tax for gasoline and diesel cars on total fuel usage as in Table 6 but simulations are based on the parameter estimates of RC Logit III in Table B.3. The figures refer to Germany in 2011.

Table B.9: Explaining the Diesel Market Shares Across Countries

	Current Situation			Change in Diesel Share: Equalization of		
	Fuel price gap	Fuel Cons. Gap	Diesel Share	Fuel Price	Fuel Cons.	Both
			Year 1998			
Belgium	-0.25	-1.91	52%	-7%	-7%	-14%
France	-0.26	-2.04	38%	-8%	-7%	-14%
Germany	-0.22	-2.35	15%	-3%	-4%	-7%
Great Britain	0.00	-2.24	15%	0%	-4%	-4%
Italy	-0.19	-2.08	21%	-4%	-5%	-9%
Netherlands	-0.29	-2.12	21%	-6%	-5%	-10%
Spain	-0.11	-1.85	51%	-5%	-8%	-14%
			Year 2011			
Belgium	-0.24	-2.01	75%	-2%	-7%	-9%
France	-0.19	-1.81	69%	-2%	-8%	-10%
Germany	-0.16	-2.23	46%	-3%	-10%	-12%
Great Britain	0.04	-2.04	50%	1%	-9%	-8%
Italy	-0.13	-2.05	56%	-2%	-10%	-12%
Netherlands	-0.28	-2.22	30%	-3%	-7%	-10%
Spain	-0.07	-2.08	69%	-1%	-9%	-11%

The table reports: (i) in the first three columns, the currently observed gaps in fuel prices and fuel consumption 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 consumption gap were eliminated. The simulations are based on the parameter estimates of RC Logit III in Table B.3.