Taxes vs. Standards as Policy Instruments: Evidence from the Auto Market*

Cristian Huse† Nikita Koptyug‡

May 2017

Abstract

We take advantage of a unique institutional setting which allows consumers to separately value fuel costs and vehicle tax (road tax), which is equivalent to a standard. We estimate a consumer-level structural model of vehicle choice using revealed preference data and controlling for heterogeneity at the micro level. We find that consumers undervalue both fuel costs and vehicle tax, but the undervaluation of vehicle tax (standard) is substantially more severe. We examine potential explanations and document that behavioral explanations, in particular salience of the policy instruments, lie at the root of our findings; for a number of the salient versions of vehicle tax and fuel costs we construct, we cannot reject the null hypothesis of their correct valuation. This also holds when using different measures of news and online search activity as proxies for salience. The results call for complementary policy instruments to restore market efficiency and for measures to make them more salient to consumers.

Key words: automobiles, energy efficiency, energy paradox, energy efficiency gap, fuel tax, standard, fuel economy, CO2 emissions, vehicle tax, road tax.
JEL Classification Codes: D12, L62, Q40, Q41, Q48, Q50, Q58.

---

*We thank participants at the IFN workshop and the International Industrial Organization Conference (IIOC) for feedback. We are also indebted to Reyer Gerlagh, Shanjun Li, Shaun McRae (IIOC discussant), and Sebastian Schwenen for insightful comments. The usual disclaimer applies.

†Stockholm School of Economics. Email address: cristian.huse@gmail.com. Web: http://cristianhuse.webs.com. Correspondence address: Stockholm School of Economics, Box 6501, 113 83 Stockholm, Sweden. Phone: + 46 8736 9167.

‡Research Institute of Industrial Economics (IFN). Email address: nikita.koptyug@ifn.se. Correspondence address: Research Institute of Industrial Economics, Grevgatan 34, Box 55665, 102 15 Stockholm, Sweden.

1 Introduction

Electricity production and fuel for transportation were responsible for over 50 percent of greenhouse gas emissions (GHGs) in the US in 2014 (U. S. Environmental Protection Agency 2016). Behind energy consumption in both sectors lie energy-intensive products, typically durable products such as cars and household appliances. The fact that these products provide their owners a stream of services over time (often years, if not decades), and that the costs associated to their use are a non-trivial share of a typical household's expenditure, begs the question of whether consumers take into account future operating costs when they purchase such products.

Not surprisingly, the valuation of energy efficiency has become key in the study of energy demand since at least Hausman (1979). Oftentimes, consumers are found to undervalue future energy costs, in what has become known as the Energy Paradox (Jaffe and Stavins 1994).

In addition to its relevance for the environment, the trade-off between product prices today and (the different components of) lifetime operating costs are also important for public policy and businesses. In the case of the former, efficient taxation depends on how consumers address this trade-off; as for the latter, firms are expected to develop their products – and price them – in accordance to regulation and expected consumer behavior.

In this paper, we examine the trade-off between product prices and future operating costs looking at consumer choices in the used car market. We address three research questions. First, we examine whether consumers correctly value lifetime operating costs of a used car upon its purchase. The correct valuation is our null hypothesis throughout the paper. Under standard, neoclassical, assumptions consumers would correctly anticipate lifetime operating costs and these should move one-for-one with vehicle prices.

Second, we investigate whether consumers value the different components of operating costs in the same way. In particular, we assess if consumers take into account lifetime fuel costs and lifetime vehicle taxes (or road taxes) in a similar fashion. This is important because while fuel taxes – which restore market efficiency in a setting where market failures are due only to environmental externalities such as emissions – are embedded in the former, standards act in a way similar to the latter. Since both components are essentially different types of operating costs, our null hypothesis is that consumers should value them equally.

Third, we examine the mechanisms underlying any departures from our null hypotheses. In particular, we examine the role of behavioral explanations such as salience and inattention on consumer decisions. Despite recent empirical findings casting doubt on this assumption, many central theoretical results in public economics rely on the assumption that consumers fully optimize with respect to taxes. Our setting is especially well-suited to investigate such issues given the coexistence of a fuel tax and a vehicle tax (equivalent to a standard) in the market we study. Thus, if valuations differ, and these differences can be attributed to differences in salience, policymakers would improve consumer welfare and the effectiveness of policy instruments by increasing their salience. In particular, increasing the salience of policy instruments has the potential to close the energy efficiency gap.

Empirical Strategy. Simply put, we examine whether consumers correctly value the lifetime operating costs of durable products (and its components). To do so, we start from the primitives of a consumer choice model of rational behavior for individual consumers. We construct a dataset of individual, retail-level, transactions in the used car market to examine our hypotheses of interest. Our revealed preference data comes from used car auctions; we follow the activity of individual consumers over time which enables us to recover their choice sets. Moreover, our micro level dataset allows to control for heterogeneity at the very micro level for both products and especially consumers – in particular, we are able to control for unobservable (time-invariant) consumer heterogeneity using consumer fixed-effects.

We specify and estimate a structural econometric model of vehicle choice where a consumer chooses the vehicle that maximizes their conditional indirect utility taking into account both product and household characteristics. In particular, we rely on a random coefficients logit model which accounts for consumer heterogeneity at the micro level and can arbitrarily approximate any choice model (McFadden and Train 2000). We allow for heterogeneity at the individual level in the form of consumer fixed-effects in addition to heterogeneity in both prices and (different components of) operating costs. This is important because it will allow us to depart from the representative consumer framework typically used in the theory and most of the empirics in the related literature. After all, if automobiles are highly differentiated products, this is partly due to substantial heterogeneity in consumer preferences.

With the model in place, we take advantage of a unique institutional setting in the Swedish market where two key policy instruments co-exist and thus can be separately examined. The first component is a fuel tax which is embedded in the expected lifetime fuel costs of a vehicle. A fuel tax is important because it can achieve the first-best solution and thus restore market efficiency in settings where emissions are the only market failure leading to inefficiencies. However, this crucially relies on its correct valuation by consumers.

The second component is a vehicle tax, which is equivalent to a standard, the policy instrument most often compared to – and combined with – the fuel tax. Standards (on emissions, fuel economy) are instruments of choice for a number of policymakers and governments worldwide due to the alleged inefficiency of the fuel tax resulting from the undervaluation of fuel costs, but also due to the relative ease with which they are received by the public opinion, i.e. voters, as compared to a fuel tax. This so happens because the burden of standards is perceived to fall on carmakers rather than consumers, in contrast to what is believed to happen in the case of the fuel tax.

Different sources of variation in the data help identify the parameters of interest. In the case of the fuel cost parameter, identification comes from variation in fuel prices across fuel types (diesel and gasoline) and over time which is interacted with fuel economy across products (which in turn varies over time, though to a lesser extent). This is arguably richer variation than in other markets such as the U.S., where identification is obtained via time series variation in fuel prices interacted with fuel economy across products.

In the case of the vehicle tax parameter, identification comes from variation across fuel types – since vehicles operating on different fuels are taxed differently –, in addition to variation of CO2 emissions (which directly maps onto fuel economy, conditional on fuel type) which varies across products, and also over time.

Main Findings. We find the average valuation of total operating costs (fuel costs plus vehicle tax) to be 0.50 and reject the null hypothesis of correct valuation at the 1 percent significance level, which is consistent with its undervaluation by consumers.[1] Alternatively, and perhaps more accurately in a framework that departs from the standard representative agent framework, we find that 100 percent of consumers undervalue operating costs.[2]

Once operating costs are decomposed into (lifetime) fuel costs and vehicle tax components, it becomes clear how differently they are valued by consumers: the average valuations of fuel costs and vehicle taxes are 0.60 and 0.14, and we reject the null of correct valuation at the 5 and 1 percent significance levels, respectively. That is, while both fuel costs and vehicle tax are undervalued, the degree of undervaluation of the latter is far more severe. As before, the finding of undervaluation also holds in the distributional sense in that 100 percent of consumers undervalue both fuel costs and vehicle tax.

Given the difference in the valuations of the two policy instruments, we then investigate po-

[1]The recent literature typically finds mixed evidence of consumer undervaluation, see Greene (2010) and Helfand and Wolverton (2011) for recent reviews.

[2]Consistent with this approach, in a recent survey Anderson and Sallee (2016) state that *"Even if biases in consumer choice are small on average, they could nevertheless be quite large for some consumers. Too little is known about the effects of fuel-economy labels, information programs, and nudges on vehicle choice in this context."*

tential explanations driving the results, in particular by comparing persistence and salience; the empirical evidence we gather suggests that salience lies at the root of our findings. We examine in detail the role of salience in increasing valuations and potentially closing the energy efficiency gap. To do so, we introduce salient components of both fuel costs and vehicle tax by interacting these variables with proxies for their salience. For instance, interacting lifetime fuel costs with indicators of oil being priced above USD 100, "round" prices, and large increases of fuel prices all induce a higher valuation of fuel costs to the extent that we cannot reject the null hypothesis of correct valuation any longer. That is, the increased salience in fuel costs comes from variables which are functions of fuel prices and exploits the time series behavior of such variables.

In a similar fashion, we introduce salient components of vehicle tax, exploiting temporary tax exemptions and the timing of tax payment within a calendar year. As in the case of fuel costs, we cannot reject the null of correct valuation of vehicle tax in months before it is due to be paid. That is, the salience of vehicle tax is induced by the month (or quarter) when annual vehicle tax is due, and temporary tax exemptions decrease the salience of vehicle tax.

Given the evidence consistent with the role played by salience, we examine two potential channels through which fuel costs might become more salient, namely news and internet search activity.[3] According to our results, increases in news related to oil prices, and both news and search related to gasoline prices, increase the valuation of fuel costs to the point that the null hypothesis of correct valuation cannot be rejected any longer. These findings are suggestive of salience being (at least partially) captured by news/searches for gasoline and oil prices.

Taken together, our findings point to the importance of behavioral explanations in influencing consumer valuations of policy instruments and in closing the energy efficiency gap.

Contribution and Related Literature. Our contribution comes from a combination of institutional setting, data, and empirical strategy. To the best of our knowledge, we are the first to take advantage of a unique institutional setting which allows the joint evaluation of two key policy instruments, namely fuel tax and standards (via their equivalent, vehicle taxes). Most of the literature typically evaluates only the role of the fuel tax; if the findings are consistent with undervaluation of fuel costs, the policy implication of these papers is that standards are preferred to taxes (Parry, Walls and Harrington, 2007). However, our findings suggest that the undervaluation of standards may be even more severe than that of fuel costs.

With the institutional setting in mind, we pursue two important steps to enable our analysis. First, we construct a unique, individual-level, revealed preference dataset focusing on retail consumers.[4] This will allow us to control for important sources of heterogeneity at the micro level, especially consumer heterogeneity. To our knowledge, we are the first to fully exploit the micro – specifically, the consumer – dimension of the auction data often used in the recent literature. This brings us closer in spirit to the empirical literature using micro-level data to investigate the effects of environmental policy on the transport sector, in particular the car market, e.g. Goldberg (1998), Bento et al (2009).

Second, we formulate a structural econometric model at the micro level which we take to data. For instance, since we observe when a bidder is active on the market, we can control for consumer characteristics and allow for richer patterns of consumer heterogeneity as compared to the literature estimating the demand for automobiles.

We contribute to different strands of the literature. First, we contribute to the broader literature on consumer choice of energy-intensive products (Hausman 1979; Dubin and McFadden 1984). Following this seminal work, substantial research has examined the effects of fuel prices on vehicle markets. Different papers have focused on price changes looking at the used car market. The reduced-form branch of this literature was pioneered by Kahn (1986), which tests whether the

[3] Although their interpretation differ, with news arguably a more passive way of making a particular variable salient when compared to search activity, by-and-large the results point in the same direction.

[4] Busse et al (2013b) stress the pitfalls of using wholesale data to draw inferences for retail consumers.

relative prices of used cars fully adjust to changes in the relative net present value of relative fuel prices.[5] Much of the recent empirical work has focused on whether an energy paradox exists, such as Busse, Knittel and Zettelmeyer (2013), Allcott and Wozny (2014), and Sallee, West, and Fan (2016). Consistent with the recent literature looking at the transport sector,[6] these papers use the time series variation in fuel prices interacted with fuel economy to estimate the valuation of energy costs. Busse, Knittel and Zettelmeyer (2013) quantify how changes in fuel prices affect prices and quantities of new and used vehicles in the quartiles of the fuel economy distribution. The discount rates implied by their estimates are consistent with mild undervaluation of energy efficiency, if at all. In turn, Allcott and Wozny (2014) document modest undervaluation of energy costs for a variety of specifications whereas Sallee et al (2016) provides evidence that consumers correctly value fuel economy using variation in odometer readings, which they interact with variation in fuel prices. That is, they exploit the fact that vehicles with different expected remaining lifetimes are affected differently by shocks to fuel prices when it comes to their lifetime fuel costs.[7]

Methodologically, our paper more closely aligns to the structural branch of this literature, which goes back to Hausman (1979) and Dubin and McFadden (1984).[8] Important contributions within this approach also include those of Goldberg (1998), who estimates a discrete-continuous consumer choice problem using cross-sectional data, finding evidence consistent with correct valuation of fuel costs, and of Verboven (1999, 2002), which use aggregate product level data to quantify the valuation of energy costs and the pricing behavior of carmakers, finding mild undervaluation by consumers. More recently, Grigolon, Reynaert and Verboven (2015) quantifies the valuation of energy costs using market-level data for a panel of European markets. Following recent findings that not accounting for heterogeneity in willingness-to-pay for fuel costs biases the valuation of energy costs (Bento, Li and Roth 2012), Grigolon et al (2015) carefully control for consumer heterogeneity, especially in what regards to mileage, finding that consumers modestly undervalue energy costs.

Our result according to which the extent of undervaluation differs across components can be rationalized by models of consumer inattention in which consumers are constrained in their costly deliberation time and devote limited attention to the computation of arguably complex cost components such as lifetime fuel costs and lifetime vehicle tax (see Conlisk 1996 for an early contribution). Since fuel costs are substantially higher than vehicle tax, they are more salient (Gerlagh et al 2016) and thus receive more attention from the part of consumers, resulting in a higher valuation.

Our paper also relates to an important literature focusing on consumer behavior within public and environmental economics. Specifically, we contribute to the recent empirical literature on taxation which documents that individuals optimize imperfectly with respect to many different taxes and thus questions a central tenet in public economics according to which agents fully optimize with respect to taxes.[9] For instance, Chetty, Looney and Kroft (2009) examine the effect of salience – taken to mean the visibility of the tax-inclusive price – on behavioral responses to taxation and find that making taxes more salient by including them in the posted prices displayed to consumers

[5]See also Kilian and Sims (2006) who extend Kahn's approach, and Alberini et al (2016) who use a regression discontinuity design when comparing similar vehicles with different energy labels.

[6]For instance, Li, Timmins and von Haefen (2009) pursue a similar strategy to study vehicle scrappage whereas Klier and Linn (2010) study the effect on new vehicle sales.

[7]Among the extensive literature, related work that also examines the response of vehicle prices to changes in fuel prices includes Sawhill (2008), Langer and Miller (2013). Ohta and Griliches (1986) investigate how changes in fuel prices affect the willingness-to-pay for product characteristics of automobiles.

[8]Note, however, our different treatment of the, continuous choice, utilization component as compared to the latter.

[9]In his seminal analysis of optimal commodity taxation, Ramsey (1927) assumes that agents respond to changes in taxes in the same fashion as changes in price whereas fundamental results by Harberger (1964), Mirrlees (1971), and Atkinson and Stiglitz (1976) on tax incidence, efficiency costs, and optimal income taxation all rely on full optimization with respect to taxes. While Chetty et al (2009) document consumer inattention to sales taxes, Hossain and Morgan (2006), and Abaluck and Gruber (2011) document inattenation to shipping charges and out-of-pocket costs for health care.

in store have larger effects on demand. Finkelstein (2009) finds that policymakers take advantage of the reduced salience induced by electronic toll collection systems whereas Gallagher and Muehlegger (2011) document that sales tax waivers given at the time of purchase (and thus more salient) have a much larger effect on the sales of hybrid electric vehicles as compared to income tax credits of an equivalent amount. More recently, Li, Linn, and Muehlegger (2014) find evidence that consumers respond more strongly to changes in (more salient) fuel taxes than (less salient) net fuel prices. Finally, Sexton (2015) finds that consumers who enroll in automatic bill payment programs – and thus forgo regular inspection of their bills, which makes them less salient – increase their consumption of electricity.

After empirically ruling out persistence in favor of salience as an explanation for our findings, we employ our structural model to explicitly test whether the valuation of the policy instruments increases whenever they are more salient; consumer valuations do not only increase, but we document a number of cases where increasing salience does not allow us to reject the null of correct valuation of lifetime fuel costs and vehicle taxes; we interpret such rejections as evidence consistent with closing the energy efficiency gap. Our findings thus contribute to an extensive literature examining the energy paradox and the energy efficiency gap, see Gerarden, Newell, and Stavins (2017) for a comprehensive survey. While this literature typically offers a number of potential explanations for the gap – often grouped into categories such as market failures, behavioral explanations and modeling flaws, as in Gerarden et al (2017) – our results provide quantitative support for the importance of behavioral explanations, of which salience and consumer inattention are particular cases, in closing the gap, and in addition allow us to rule out a number of competitive explanations.

The paper proceeds as follows. Section 2 presents the institutional background whereas Section 3 presents the data. Section 4 then outlines the empirical strategy. While Section 5 presents results related to the valuation of operating costs, Section 6 focuses on potential explanations for our findings. Section 7 summarizes the results and discusses their implication. The final section concludes

2 Institutional Background

2.1 Policy Instruments

Consumers do not always bear all the consequences of their actions, which leads them to often ignore the external costs of their decisions imposed on other economic agents. The consequence of such *externalities* is that the social cost of consumers' actions exceeds the private cost of the activity, and this inefficient outcome may lead to over-utilization of a resource. Regulators are then faced with the challenge of designing policy instruments to correct the inefficient outcome, in order to induce decision-makers to internalize these externalities.

In the case of the transport sector, to the extent that the externality policymakers aim to address is CO2 emissions (more generally, greenhouse gases) or fuel consumption, a fuel tax typically charged on a per-liter or gallon basis is a Pigouvian tax on emissions and can achieve the first-best solution in the absence of consumer undervaluation of energy efficiency. Otherwise, to achieve a first-best policy one needs to combine a fuel tax with a standard and a car tax (at least in the case of a representative agent model).[10]

Standards can equivalently be imposed on CO2 emissions, fuel consumption or fuel economy.[11]

[10]For alternative externalities, other instruments would be first-best. For instance, road pricing would address the congestion externality, a per-mile tax based on vehicle type might be first-best to address the accident externality. In case local pollutants are also considered, the design of policies becomes more complex (Fullerton and West 2002, 2010; Knittel and Sandler 2012). See Parry, Walls, and Harrington (2007) and Anderson and Sallee (2016) for recent surveys of transport-related externalities.

[11]Recall the a physical relation $\rho(.)$ between a vehicle's CO2 emissions (e_j, in gCO2/km) and fuel consumption

They impose a maximum value on the variable being regulated and, when binding, act as an implicit, revenue-neutral tax on the variable of interest.[12]

The inefficiency of standards occurs due to the rebound effect.[13] That is, standards incentivize the purchase of more efficient cars but the resulting savings generate an increase in utilization (driving). Intuitively, this can be shown as follows. Assume the only transport externality is CO2 emissions, which are denoted by $\mathcal{E}_{ij}$ for driver i and vehicle j, and are measured in gCO2. A fuel tax incides on both vehicle utilization θ_{ij} (measured in VKT, vehicle-kilometers-year, say) and emissions e_j whereas the vehicle tax incides only on the latter. Thus, while the incidence of fuel taxes falls on all transport externalities, those of standards fall only on part of them, which results in their inefficiency.[14]

2.2 Policy Instruments in the Swedish Market

The main policy instruments in the Swedish car market used to internalize externalities arising from car usage are a fuel tax and a vehicle tax.[15] [16] The Swedish fuel tax is comprised of an energy tax and a carbon tax which combined make one of the highest fuel taxes – and thus fuel prices – worldwide. For reference, at USD 5.93 per gallon in 2016:Q2, gasoline prices in Sweden are the 9th highest worldwide. In contrast, at USD 2.57 per gallon gasoline prices are substantially lower in the US (Bloomberg 2016).[17]

Any changes to the fuel tax typically happens yearly, with the gasoline tax historically being higher than the diesel tax.[18] This results in retail prices for gasoline being typically (but not always) higher than those of diesel and the fact that retail fuel prices vary mostly due to changes in net fuel prices rather than fuel taxes – including but not only – during our sample period. Both effects are displayed in Figure 1.

<Figure 1 here>

A vehicle tax is another important policy instrument in the Swedish market. As pointed out at least since Verboven (1999, 2002), vehicle taxes play an important role when it comes to technology adoption (gasoline vs. diesel vehicles) in the European car market.[19] Crucially, a vehicle tax can be seen as equivalent to a standard since in both cases the marginal cost of driving decreases in the variable being regulated, see e.g. Anderson and Sallee (2016) for details and discussion. Alternatively, the incidence of a vehicle tax is the same as that of its corresponding standard, which

(ϕ_j, in liters/km).

[12]Letting τ_s be the *implicit* tax rate, e a vehicle's CO2 emissions, and $\varphi > 0$ the binding standard, the tax imposed is $\tau_s(e - \varphi)$. While in the case of the standard the tax rate is implicit, in the case of the vehicle tax it is *explicit*, see the Appendix for details. Moreover, while the standard consists of a one-off (implicit) payment inciding on new vehicles, vehicle taxes are due yearly in our setting, thus arguably more salient.

[13]Additional externalities, such as congestion, will favor fuel tax.

[14]However, the empirical evidence in the economics literature since at least Goldberg (1998) is that consumers respond little to increasing fuel prices, see e.g. Small and Van Dender (2007), Hughes, Knittel and Sperling (2008), Li, Timmins, and von Haefen (2009).

[15]In contrast with countries such as Denmark, no (one-off) registration tax is charged in Sweden. Registration taxes have been used as an important policy instrument in Denmark. There, the registration tax is either 105 percent or 180 percent of the vehicle price depending on whether the price is below or above 81,700 DKK (about USD 12,150). This arguably contributed for Denmark to be the first country to attain the EU-wide CO2 emission targets for 2010 (130 gCO2/km). However, EU harmonization rules are set to rule out registration taxes in the near future.

[16]The EU CO2 emission standards introduced in 2009 affect only a minority of the used vehicles in our sample.

[17]Bloomberg (2016) provides a fuel price comparison engine available from http://www.bloomberg.com/graphics/gas-prices/#20162:Sweden:USD:g.

[18]As pointed out in Miravete et al (2015), the fact that fuel is taxed by volume instead of energy content since the 1970s following the inception of the European Fuel Tax Directive results in a somewhat lower tax rate for diesel as compared to gasoline.

[19]More recently, Klier and Linn (2015) have also shown the role of vehicle taxes in the reduction of CO2 emissions – to a great extent via an increase in the market share commanded by diesel vehicles – in three European markets.

does not include driving (VKT). To our advantage, since the vehicle tax is effectively charged in monetary units, we are able to value it in the same way we value fuel costs.

<Figure 2 here>

Since the 2006 Swedish vehicle tax reform, CO2 emissions have replaced vehicle weight as the main determinant of the vehicle tax. Vehicle taxes are calculated for each car using pre-specified formulas, with the vehicle tax on a diesel vehicle being approximately 2.5 higher than the tax on a gasoline vehicle with the same CO2 emission level – see the Appendix for details and the exact formulas.[20] In contrast to what happens in most EU countries, vehicle taxes have been higher for diesel as compared to gasoline vehicles in the Swedish market since at least the 1980s. It is then not surprising that they have been singled out as the main inhibitor to diesel adoption in Sweden. In fact, the share commanded by diesel passenger cars in the Swedish market is among the lowest within the EU and experienced only a mild increase in the period 1995-2009 despite the strong wave of *dieselization* experienced in Europe since the 1990s, see Figure 2.[21]

<Figure 3 here>

Since the calculation of vehicle tax has not changed since 2006, the time series variation in vehicle tax is also substantially lower than the variation in fuel prices over time, see Figure 3 for an illustration using baseline gasoline and diesel versions of the best-selling Volvo V70. What is more, despite substantial decreases in CO2 emissions over time, they were not enough to counteract the penalties imposed on diesel vehicles.

Empirical Implementation. Empirically, we will use information on fuel costs and vehicle tax to draw implications for fuel tax and standards, respectively. In the case of the former, the justification comes from the public economics literature; for instance, most of the literature since Ramsey (1927) assumes that agents respond to changes in taxes in the same fashion as changes in price in his analysis of optimal commodity taxation. In the case of the latter, this relies on the equivalence of a vehicle tax and a standard, as argued above. To the extent that the explicit and recurrent (charged yearly) vehicle tax is more salient than the (implicit and one-off) standard, the valuation of the former could be thought of as an upper bound to the latter.

3 Data

We combine a number of datasets. Our unique transaction data allows the construction of consumer-level choice sets based on when a particular individual was active on the market. The use of this revealed preference data of retail transactions – thus not subject to concerns raised in e.g. Busse et al (2013b) regarding the use of wholesale data to make inferences about consumer behavior – allows to control for consumer heterogeneity at the very micro level, in stark contrast with the literature; in particular, it allows the use of consumer (bidder) fixed-effects.

Our car inspection data enables the construction of lifetime mileage estimates at a more detailed level as compared to those (age-invariant ones) typically used in the literature. It is well-known that mileage decreases with age, in addition to being heterogeneous across models and fuels, and our mileage data provides more reliable information in this important dimension.

[20]Vehicle tax is paid according to the Road Traffic Tax Act (2006: 227) which came into force on 1 May 2006, see https://lagen.nu/2006:227. The Act repeals the previous regulation in place since 1988 (1988: 327), see http://rkrattsbaser.gov.se/sfst?bet=1988:327. Exemptions are detailed in Chapter 2 of the Road Traffic Tax Act (2006: 227). The vehicle tax is due in a given month depending on the last digit of a vehicle's registration number.

[21]Miravete et al (2015) document that the share commanded by diesels in Europe increased from as low as 10 percent in 1990 to up to 70 percent by 2000 in some segments. Klier and Linn (2015) argue that vehicle taxes and demand for fuel economy combined (diesel variants have a higher fuel economy than their gasoline counterparts) explain diesel adoption in the continent.

We also use estimates of vehicle hazard rates (scrapping probabilities) for the Swedish car fleet disaggregated by age and fuel.

Finally, we use search volume and news publications indices data from Google Trends. A detailed description of the sample and how we construct it can be found in the Appendix.

Transaction Data. The transaction data used in this study is auction data from a large Swedish platform collected from 12th October 2011 to 13th June 2013.[22] Car auctions on the platform are conducted on a weekly basis, with new vehicles being listed throughout the day every Friday and ending on either Tuesday, Wednesday or Thursday the following week.

Once a car has been listed for auction, anyone is free to follow the auction process and view the detailed vehicle information without having to register an account. The auction is an ascending auction and individuals interested in placing a bid have to register an account and are able to place bids, with a minimum incremental bid of 500 SEK, at no monetary cost. Should a bidder place the winning bid, the bidder is legally bound to go through with the purchase.[23] The history of bids is publicly available, although only the five latest bids are shown, together with the user name of the individual placing the bid, the time at which the bid was placed as well as the bid itself. Once the auction ends, all information regarding the vehicle is removed.

Every car listed for auction is carefully reviewed by an independent mechanic, and a summary of the review in the form of ratings on engine, body, transmission, brakes and interior, together with detailed technical information is made available to the public at the start of each auction.[24] The information is provided in a standardized format and can be broadly categorized into technical, legal and qualitative; technical and legal information is taken from the Swedish Transport Agency's vehicle register. This register contains information about the bodywork, engine, size as well as the inspection period and tax details for each vehicle.[25] The qualitative information contains the mechanic's assessment of the different parts of the vehicle including an overall vehicle assessment and more detailed assessments of the bodywork, engine, breaks, gearbox and interior. Each vehicle is located at one of 15 locations in Sweden and anyone interested may go and test-drive the given vehicle on pre-specified dates.

The exact time at which an auction ends is not known in advance. The auction ends if no new bids are placed within a 3 minute period after a pre-specified time of day.

Importantly, to credibly focus on consumers, our analysis will consider only bidders who have purchased one vehicle and participated in at least two auctions.

Car Inspection Data. In Sweden, every car has to pass a mandatory inspection to assess whether it complies with government safety and environmental standards. The first inspection for passenger cars is carried out at the age of 3 years and the second at the age of 5 years. Thereafter cars undergo inspections annually until they are 30 years of age, after which the car only needs to be inspected every other year. Each car is assigned an individual 4-month inspection window, depending on the last digit of the registration number, and must have carried out an inspection during this period. If the car isn't inspected during the assigned period it is banned from the roads until an inspection has been carried out.

[22] This platform is the leading player in Scandinavia. Anecdotal evidence collected during the course of the project points to the platform having high reputation, its products being sold at "competitive prices" and being widely used by a diverse cross section of consumers, from people with only basic (compulsory) education to university professors.

[23] This means that should a bidder decide to back out after having won the auction, (s)he will have to pay a fee for that particular auction as well as any costs incurred by the seller in putting the vehicle up for auction again. While this means that one does not actually have to purchase the vehicle and incur its entire cost, the cost of compensating the seller for loss of business may still turn out to be substantial.

[24] To minimize conflicts of interest, the auction site provides full refunds and the possibility to terminate the transaction if consumers are able to prove serious hidden faults with the help of their own mechanic.

[25] The Swedish Transport Agency also provides details about the month of payment for vehicle taxes. See the Appendix for details.

Until 2010, the government-owned Swedish Automobile Inspection Association (SAIA, Bilprövningen) was responsible for carrying out all car inspections in Sweden and publicly provided failure rates and average odometer readings disaggregated by model, age, fuel type and year of publication. In 2010, the inspection market was deregulated and the SAIA was privatized. We have collected this data for the years from 2007 to 2010.

Car Survival Data. We collect data on the number of cars in the Swedish car fleet in 2011 disaggregated by car age from Statistics Sweden. Further, we collect data on the number of cars that were scrapped in 2011, disaggregated by car age, from Bil Sweden.

Fuel Price and Volume Data. We collect data on gasoline and diesel prices from fuel retailer Statoil AB's web page. Statoil provides daily recommended prices for it's fuel stations. We collect information on aggregate fuel volumes delivered within Sweden from the Swedish Petroleum and Bio-fuel Institute.

Search Volume and News Publications Data. We collect search index values for searches of the keywords "Oil price", "Gas Price" and "Diesel Price" in addition to news publication indices for the same keywords from Google Trends. Figure 4 displays their time series behavior.

<Figure 4 here>

Combining Datasets. We combine the transaction data (product characteristics and indicator of purchase of product j by consumer i) with the search volume and news publication, and fuel price data by auction ending date. We match the vehicle tax to the month payable through the last digit of each car's registration number. We use the car inspection data to estimate and predict average VKT as a function of car age and we use the car survival data to estimate and predict average car hazard rates as a function of total VKT traveled. A sample choice set is displayed in Table 1.[26]

<Table 1 here>

Table 2 reports summary statistics on bidders, vehicle characteristics and driving characteristics for the final sample of 3,150 bidders and 12,961 auctions.

<Table 2 here>

4 Empirical Strategy

We estimate a discrete choice model of used vehicle purchase using consumer-level data. To do so, we start from the primitives of a consumer choice model in which we observe choice sets, actual choices, in addition to characteristics for every product and consumer. Thus, we are able to model consumer heterogeneity – both observed and unobserved – at its very source.[27] Moreover, the model is flexible and accounts for heterogeneity also at the product dimension. Finally, our use of

[26]Note that we consider only vehicles fueled by diesel and gasoline in our sample, despite the availability of vehicles with other fuel types. This so happens because for the leading alternative fuel type (flexible-fuel vehicles) fuel choice becomes another choice dimension which can significantly complicate the analysis, since the lifetime fuel costs will depend on fuel choice and the expectations of future fuel prices (see Salvo and Huse 2013, and Huse and Lucinda 2014 for examples). Hybrid electric vehicles (such as the Toyota Prius) and Gasoline-CNG hybrids command a negligible share of the market and are also discarded.

[27]This is in contrast with the more traditional literature using product (or market) level data, see Berry, Levinsohn and Pakes (1995) and the recent contribution by Grigolon, Reynaert and Verboven (2015). As mentioned in Goldberg and Verboven (2001), *"Ideally, [our] estimation approach would exploit data at the individual consumer level; such data would allow us to introduce consumer heterogeneity and its interactions with product characteristics in a very flexible manner. Unfortunately, such micro data are not available for any European country(...).*

retail data is consistent with a study whose focus is the behavior of retail consumers, which rids itself from recent criticism directed at the use of wholesale data (Busse et al 2013b).[28]

Model Specification. We estimate the demand for automobiles using a random coefficients logit model for consumer-level data. The starting point is a microeconomic model of rational behavior for individual consumers. Consumers purchase one of the products available to them, namely the one yielding them the highest utility. The econometrician does observe individual choices, prices and a set of characteristics for each of the J products available for a number of markets and periods, in addition to consumer characteristics.[29] Letting i index consumers ($i = 1, ..., N$) and j index products ($j = 1, ..., J$), define the conditional indirect utility of a consumer as

$$U_{ij} = x'_{ij}\beta_i + y'_i\gamma_i - \chi'_{ij}\alpha_i + \varepsilon_{ij}$$

The vector x'_{ij} comprises product characteristics such as price, engine size, vehicle age and model fixed-effects, and ratings by an independent mechanic for engine, body, transmission, brakes, and interior. We include such ratings (attributed to each vehicle by an independent mechanic, see Section 3) to control for heterogeneity across vehicles of same observed characteristics (e.g. model-fuel-vintage-engine size) but different usage – see Appendix C for robustness estimates using odometer readings. The vector y'_i comprises consumer characteristics such as consumer fixed-effects, income or region fixed-effects; $\chi'_{ij} = (c_{ij}, \tau_j)$ are product lifetime operating costs (see below); and ε_{ij} is a mean-zero stochastic term with a type-1 extreme value (T1EV) distribution.

The lifetime operating costs of vehicle services comprise fuel costs (c_{ij}) and taxes (τ_j) which we write as

$$c_{ij} = \sum_{t=0}^{T} \delta^t s_{jt}\pi_{jt}\phi_j\theta_{ijt}$$

$$\tau_j = \sum_{t=0}^{T} \delta^t s_{jt}\tau_{jt}$$

with δ being the discount factor, T the lifetime of a vehicle, and s_{jt} is a vehicle's survival probability (see the Appendix for details).[30] [31] In the case of the fuel costs, $\pi_{jt}\phi_j\theta_{ijt}$ combines fuel price (π_{jt}, in

[28]For wholesale prices to reflect consumer behavior, the pass-through from wholesale to retail prices should be one-for-one. One setting where this holds is under a competitive dealer market. However, Koptyug (2015) documents that the behavior of dealers differs from that of consumers in the Swedish used car market. This suggests that dealers acquire cars and re-sell them charging a markup. If this markup were constant, then any parameter estimates coming from wholesale regressions should be scaled to conform with the behavior of retail consumers. However, market conditions varying across products and over time undermine the constant markup assumption and begs for the use of retail prices.

[29]In practice, since consumers enter and exit the market at different times (weeks), the choice set varies across consumers. However, since this is not crucial for our exposition, we suppress any indices on the choice set for the sake of simplicity.

[30]An alternative strategy consists of obtaining the two discount rates that set the valuations of lifetime fuel costs and vehicle taxes equal to the price of a vehicle, but this would require solving for them implicitly based on a nonlinear model, something computationally more demanding than what we pursue with little additional insight. Our strategy is consistent with a number of papers in the literature; Verboven (2002) and Busse et al (2013) use borrowing rates, Allcott and Wozny (2014) use the average of borrowing costs and stock market returns, and Sallee et al (2016) assume a discount rate.

[31]Although the discount factor $\delta = 1/(1 + r)$ could in principle vary across consumers, we assume this is not the case in our empirical analysis; we set the discount rate r to 5 percent in what we see as a conservative benchmark given borrowing costs in the Swedish market prevailing at the period. This rate is in line with the 2014 Guidelines of the U.S. Environmental Protection Agency (2014). (Using alternative discount rates of 3 and 7 percent does not affect the qualitative results.) Nominal interest rates for car loans, arguably one of the factors underlying consumers' actual discount rates, typically vary according to the applicant's credit history, the amount applied for, and the term of the loan; to gauge their magnitudes, we have submitted queries to a number of loan providers and have observed nominal interest rates in the range 2.5-7 percent, with lower rates typically offered by banks and higher rates often offered by "shark lenders". These findings are consistent with reports of nominal interest rates in the range 4-7 percent reported

SEK/liter), vehicle fuel economy (ϕ_j, in liters/100km) and vehicle-kilometers/year (VKT, in θ_{ijt} km/year) whereas for the tax term, τ_{jt} is the annual vehicle tax (function of fuel, CO2 emissions). Following the literature, we assume fuel prices to follow random walks, which is consistent with evidence documented in Anderson, Kellogg and Sallee (2013).[32]

Generally, the WTP for a characteristic of alternative j is defined as the ratio of the marginal utility of the characteristic to the marginal utility of its cost. As a result, the valuation of the fuel cost and vehicle tax components are given by, respectively, $\upsilon_c := \frac{\partial U_{ij}/\partial c_j}{\partial U_{ij}/\partial p_j} = \frac{\alpha_c}{\alpha_p}$ and $\upsilon_\tau := \frac{\partial U_{ij}/\partial \tau_j}{\partial U_{ij}/\partial p_j} = \frac{\alpha_\tau}{\alpha_p}$.

Given the prominence of heterogeneity in the car market, we allow for different channels through which heterogeneity can affect choices.[33] On the product front, we allow for product (model) fixed-effects, product characteristics, and product ratings. Thus, we control for time-invariant heterogeneity at the product level stemming from, say, a product's reputation, while also controlling for heterogeneity across variants of a given model and/or improvements in product characteristics over time, in addition to the current state of the product.

On the consumer front, we allow heterogeneity to affect choices via consumer fixed-effects, heterogeneity in mileage, and consumer random coefficients. Mileage is disaggregated by the revealed fuel type of the consumer and age of car at purchase in our benchmark specifications, which accounts for the fact that it typically decreases with vehicle age and diesel vehicles consistently are driven more than their gasoline counterparts. Consumer fixed-effects allow controlling for time-invariant consumer heterogeneity (the Appendix reports robustness checks using demographics). Finally, we decompose the individual coefficients as $\beta_{ik} = exp(\beta_k^* + \sigma_k \nu_{ik})$, where β_k^* is common across individuals, υ_{ik} is a consumer-specific random determinant of the taste for characteristic k, which we assume to follow a multivariate normal distribution, $(\upsilon_{1i}, ..., \upsilon_{iK})' \sim \mathcal{N}(\beta^*, \Sigma)$, and $\sigma_k := [\Sigma]_{kk}$ measures the impact of υ on characteristic k.[34] It then follows that the parameter vector to be estimated is given by $(\beta^*, \sigma, \gamma, \alpha)$. Following the literature (see Train 2009), we estimate the model using Simulated Maximum Likelihood, see the Appendix for details.

Our empirical strategy builds on the institutional setting summarized in Section 2 and is consistent with both the literature and data availability. In particular, we assume that utilization conditional on product choice is perfectly inelastic, as often done in the literature, see Grigolon, Reynaert and Verboven (2015) for a recent case. This is consistent with empirical findings of a small and statistically insignificant rebound effect typically found in the literature.[35] We believe this is

by the media in 2013, see https://www.svd.se/bilkopare-lockas-med-olagliga-lan. By deducting inflation expectations available from the surveys carried out by the Swedish *Konjunkturinstitutet* (National Institute of Economic Research), the resulting real interest rates are mostly below 5 percent. Part of the reason why interest rates are relatively low in Sweden comes from the fact that the treatment of personal bankruptcy is stricter than in most countries; if one should fail to make interest rate payments creditors have a right to expropriate the car through the Swedish Enforcement Authority (SEA, *Kronofogden*). The SEA can then auction the car to pay the creditor and any amount outstanding after the auction still has to be paid to the creditor through other means (for instance expropriation of other property). If a person is so heavily indebted that it would take an excessively long time to repay creditors in full, an individual can apply for a complete debt write-off. Should the write-off be granted, a large fraction of the individual's income will be confiscated leaving only a minimal subsistence income. After five years of living on subsistence income, the debt is completely written off even if it was not paid off in full.

[32]See also Allcott and Wozny (2014) and Grigolon, Reynaert and Verboven (2015) for recent papers making the same assumption. Moreover, Alquist, Kilian and Vigfusson (2011)'s survey documents that complex models do not outperform simple models with expectations based only on current energy prices.

[33]As shown in Bento, Li and Roth (2012), one needs to account for heterogeneity in the valuation of fuel costs at the risk of biasing the valuation of energy efficiency downwards.

[34]The lognormal distribution has been proposed as a convenient distribution for random coefficients in discrete choice models by Revelt and Train (1998) and Train (1998). Importantly, its associated willingness-to-pay (WTP) is well-defined and lognormally distributed. See the Appendix for details.

[35]For perspective, Small and Van Dender (2007) estimate a short-run rebound effect of 0.044 using yearly date whereas Klier and Linn estimate an elasticity close to 0.12 using monthly data. Outside of the transport literature, Davis (2008) finds a price elasticity of clothes washing of 0.06.

a good compromise in that in contrast with most of the literature, which uses mileage data disaggregated by market segment (or weight class), we are able to obtain age-fuel-specific estimates of mileage to construct our cost variables.

Despite its flexibility, our estimation strategy abstracts from potentially important features of the auto industry. For instance, cars are durable products, so current ownership of a car is likely to affect the current demand for automobiles, see Hendel and Nevo (2006) and Gowrisankaran and Rysman (2011). Our estimation approach thus represents a pragmatic modeling approximation to actual choice behavior in the industry, being consistent with much of the literature.

Identification. One distinguishing feature of our estimation strategy is that instead of using information on auction outcomes only, e.g. price and characteristics of the product sold, we use data on all the products a consumer could bid for while active on the market to identify the preference parameters. This is only possible by following the activity history of consumers in the sample over time.

The parameters of interest, α_i, are consistently estimated provided $E[\chi_{ij}\varepsilon_{ij}|(x_{ij}, y_i)] = 0$. That is, the components of operating cost are uncorrelated with the error term conditional on consumer and product characteristics. The panel structure of the data allows the use of a number of fixed-effects. For instance, model and fuel fixed-effects to control for (time-invariant) product characteristics unobserved by the econometrician and related to, say, product reputation that may be correlated with the cost components.[36] As model characteristics may well change in ways that are correlated with the cost components, we also control for product characteristics such as engine size, vehicle age, odometer reading, plus transmission type, all-wheel drive and turbo charged fixed-effects. Ratings of body, engine, transmission, brakes and interior are additional ways to control for unobservables from the econometrician's perspective. If endogeneity is still a concern despite the above, we note that the resulting attenuation bias will render an overstated valuation parameter υ, thus making it less likely to reject the null of correct valuation of energy efficiency.

The variations identifying the preference parameters differ slightly between vehicle taxes and fuel costs. Note, however, that both relate to variations in the choice sets (thus fuel economy and CO2 emissions). In the case of vehicle taxes, the first source of variation is the difference in taxation between gasoline and diesel vehicles. Moreover, vehicle taxes differ both due to the difference in CO2 emissions of different models conditional on fuel and difference in CO2 emissions of different vintages of the same model, see Panel B in Figure 3 for an example.[37]

As for fuel costs, the variation is richer in that it combines variation in fuel prices over time combined with changes in fuel economy both over time and across products, which results in rich variation of the fuel cost component. For illustration, retail gasoline prices are about 13.50 SEK at the start of the sample, increase to about 16 SEK and return to about 13.50 SEK towards the end of the sample, see Figure 1.

Testable Implications. Under the null hypothesis of full information and rationality, consumers correctly – and equally – value lifetime fuel and vehicle tax costs. Letting subscripts p, c, τ denote the coefficients for product prices, fuel costs and vehicle taxes, respectively, the first testable implication – regarding the correct valuation of lifetime fuel costs – reads $\upsilon_c := \alpha_c/\alpha_p = 1$ whereas the second testable implication – regarding the correct valuation of lifetime vehicle taxes – reads $\upsilon_\tau := \alpha_\tau/\alpha_p = 1$. Finally, the correct valuation of both fuel costs and vehicle taxes can be expressed

[36] In particular, we follow most of the literature using micro data and assume price exogeneity (see Petrin and Train 2010 for an exception).

[37] CO2 emissions typically decrease over time, even conditioning on engine size or horsepower, due to technological improvements in the auto industry. They also vary cross-sectionally across models and especially fuels for a given product line. Moreover, while the vehicle tax calculation has not changed during the sample period, there is substantial variation within brand-fuel-model, even conditional on vintage. For instance, the best-selling Volvo V70 has CO2 emissions in the range 119-275 gCO2/km (diesel) and 164-275 gCO2/km (gasoline). For the full sample, the mean (median) ratio of maximum to minimum CO2 emissions at the brand-model-fuel level is 1.41 (1.32).

as $\upsilon_c = \upsilon_\tau = 1$.

5 Valuation of Operating Costs

We estimate demand specifications comprised of three components, namely product characteristics, household characteristics, and the specification of consumer heterogeneity. Among product characteristics we consider engine size, vehicle age, odometer reading, and weight; we also include fixed-effects for model, fuel type, transmission type, all-wheel drive, turbo-charged engines, in addition to time (year-month) fixed-effects. We also include a vehicle's (expected) lifetime fuel costs and taxes, in addition to price. Finally, we include the product ratings detailed above.

Consumer heterogeneity is accounted for in two ways. First, we include consumer fixed-effects which control for up time-invariant consumer characteristics.[38] Second, we endow price and all cost components with random coefficients. This allows different consumers to have different sensitivities (thus preferences) to different characteristics and is a standard way to account for unobserved consumer heterogeneity in logit models. Standard errors are robust and clustered at the consumer (bidder) level.

Demand and Valuation Estimates. We report our main demand estimates in Table 3. Specifications 1-3 report estimates of a conditional logit model whereas Specifications 4-6 each report estimates of a random coefficients logit (RCL) model; random draws are lognormally distributed, and both price and cost components are endowed with a random coefficient.[39] Our baseline specification is (6).

Specifications (1)-(3) are conditional logit models which differ with respect to the cost terms. While Specification (1) considers the combination of lifetime fuel and vehicle tax costs (i.e. total operating costs), Specification (2) considers only lifetime fuel costs (thus purposely omitting the vehicle tax term) and Specification (3) decomposes operating costs into fuel costs and vehicle tax terms. The mean price parameters in all three specifications are statistically significant at the 1 percent significance level and fairly similar. The cost component in Specification (2) loads higher than in Specification (1), as expected. Finally, in contrast with a significant heterogeneity parameter (at the 1 percent level) for fuel costs in Specification (3), the vehicle tax parameter is not significant. The associated valuation parameters are 0.50 for Specification (1), 0.62 for Specification (2), and 0.61 for fuel costs and 0.05 for vehicle tax in Specification (3). All cases, the null hypothesis of correct valuation is rejected at the standard significance levels.

<Table 3 here>

One shortcoming of standard logit models is that they do not account for consumer heterogeneity; in particular, heterogeneity in valuation of operating costs, e.g. stemming from heterogeneity in mileage requirements by different consumers. As a result, valuations of operating costs tend to be biased upwards (Bento, Li and Roth 2012). This is why we also estimate RCL models in Specifications 4-6. Specifications (4) and (5) are the RCL-counterparts of Specifications (1) and (2), respectively. Since the heterogeneity (standard deviation) parameters in these specifications are

[38]While we emphasize specifications where we control for consumer heterogeneity via consumer (bidder) fixed-effects, we also perform robustness checks using demographic controls such as average income of a geographic market and geographic market fixed-effects with similar results, see below.

[39]Although in our baseline the draws are uncorrelated, we have also experimented with more general specifications of the matrix of random coefficients before settling for the more parsimonious one. Our choice was driven by (i) convergence issues encountered at times; (ii) the lack of significance of most coefficients in these more general specifications. This lack of significance persisted even when imposing more structure on said matrix, e.g. non-zero covariance terms only between a subset of the coefficients. Our findings are consistent with the fact that it is difficult to identify too many random coefficients in practice. See below for a robustness check where we allow for non-zero correlations between price and each of the cost components.

small and not statistically significant, it is perhaps not surprising that the valuation parameters do not change when compared to their logit counterparts. As before, the null of correct valuation of operating costs is comfortably rejected.

So far, our valuation results point to rejections of the null hypothesis of correct valuation of operating costs. To further examine whether that is indeed the case, in Specification (6) we decompose total lifetime operating costs into fuel costs and vehicle tax components. The first thing to note is that heterogeneity becomes statistically significant (at the 10 percent level) for vehicle tax. That is, the estimates suggest that there is substantial consumer heterogeneity in the responses to vehicle taxes, despite their little variation over time, see Figure 2. As we detail below, this can be attributed to the fact that newer (less than 5 years) and older vehicles are treated differently, with the former receiving a tax exemption prior to the sample period.

When it comes to valuations, while one would expect the fuel cost and vehicle tax parameters to be of similar magnitudes provided consumers value cost components correctly, this is not the case here, as the average valuations of lifetime fuel cost and tax components are 0.60 and 0.14 respectively. When performing inference for those parameters, we obtain two important findings, namely the rejection of the null of correct valuation at the 5 and 1 percent significance levels for fuel costs and vehicle tax, respectively, and the rejection of the null of equal valuation of fuel costs and vehicle tax at the 5 percent significance level (see υ_τ/υ_cin Table 3). That is, on top of undervaluing both components, consumers undervaluation of lifetime vehicle tax is substantially more severe than that of lifetime fuel costs. Moreover, while consumer reaction is relatively homogeneous in the case of fuel costs, it is definitely heterogeneous in the case of vehicle tax.

<Figure 5 here>

The role of heterogeneity in valuations can be further appreciated by plotting the distribution of valuations for our baseline specification, see Figure 5. Valuations are substantially more heterogeneous in Panel B, which displays the valuation of vehicle tax, than in Panel A, which displays the valuation of fuel costs. In particular, the valuation of fuel cost heavily concentrated around the 0.60 value due to the small magnitude and the ultimate lack of statistical significance of the corresponding heterogeneity parameter.

The above findings have important policy implications. First, consistent with previous literature, we document that consumers consistently undervalue the different components of operating costs. The undervaluation of fuel costs has been previously documented in the literature; its main implication is that a fuel tax does not achieve the first best outcome and additional policy instruments – such as a standard – are required.

Second, we document that a policy instrument which is analogous to a standard – and is typically assumed to be correctly valued by both policymakers and researchers – is also undervalued by consumers. What is more, such undervaluation is even more severe than that of the fuel cost component. Taken together, our findings suggest that one needs a menu of policy instruments to adequately internalize transport externalities.

Robustness. We consider a number of alternative specifications departing from our baseline specification. Robustness checks include an increased choice set (top-150 best-selling vehicles instead of top-100), alternative/additional characteristics (vehicle size replacing vehicle weight; vehicle size replacing vehicle weight and HP/weight replacing vehicle size), additional consumer characteristics (market – or region – fixed-effects), alternative lifetime mileage estimates, sampling weights (so fuel market shares in our sample match fuel market shares in the Swedish car fleet), correlated random draws (allowing correlation between the draws of price, fuel cost and vehicle tax to be estimated), alternative vehicle hazard rates (replacing age-based with odometer-based survival probabilities), and sticky fuel prices (to reflect that consumers do not replace information on fuel prices on a frequent basis). Finally, we also increase our discount rate from 5 to 10 percent. The detailed results are reported in the Appendix.

Our first take-away from this exercise is that both fuel costs and vehicle taxes remain undervalued, which allows us to comfortably reject the null of correct valuation for both components. Moreover, fuel costs remain more highly valued than vehicle taxes, despite any changes in mean valuations that may occur – and the rejection of the null hypothesis of equal valuations does hold in most robustness exercises.

Conditional on a 5 percent discount rate, our alternative specifications provide valuation estimates in the range 0.44-0.66 for fuel costs (baseline is 0.60) and in the range 0.02-0.29 for vehicle taxes (baseline is 0.14). The null of correct valuation is rejected in all cases for vehicle tax and in all but one case for the fuel cost component. Moreover, the null of equal valuation of fuel cost and vehicle tax components is rejected in all but two cases.

While we do not claim outright robustness of our baseline specification, ours is not the only case where valuations change somewhat across specification. For instance, despite a substantially larger dataset in which a number of features are taken to be model-vintage averages rather than the raw characteristics as in our case, Allcott and Wozny (2014) provide an extensive set of robustness checks where there is substantial variation in their valuation estimates (0.58-0.87 [0.46-0.68] for their preferred discount rate following their baseline of 0.76 [0.55] for expectations of fuel prices following futures prices [martingales], see their Table 4). Similarly, the estimates of the valuation parameter in Sallee et al (2016) are in the range 0.48-0.98, see their Table 5.

6 Potential Explanations

According to our results, (i) consumers undervalue both lifetime fuel costs and lifetime vehicle tax; and (ii) the undervaluation of vehicle tax is substantially more severe than that of fuel costs. We now examine possible explanations for the different degrees of undervaluation encountered in the results. Following the literature, e.g. Li et al (2014), two potential explanations for this difference in valuations are persistence and salience. In the case of persistence, the reasoning is that if consumers perceive (changes in) fuel costs to be more persistent than (changes in) vehicle taxes, then they are expected to react more strongly to the former rather than the latter.[40]

Since fuel costs $c_{ij} = c_{ij}(\delta, s_j, \pi_j, \phi_j, \theta_{ij})$ are a non-trivial function of a number of variables and potentially both consumer- and vehicle-specific, the study of its dynamic properties poses a number of challenges. For instance, the expected distance driven (VKT, θ_{ij}) depends on both driver and vehicle characteristics, in addition to fuel prices. Moreover, in the medium-run a consumer could replace his vehicle, new vehicles tend to improve their fuel economy over old ones for instance. However, in the short-run, i.e. conditional on a vehicle, fuel economy, or equivalently CO2 emissions, is fixed and VKT responds little to fuel prices, so the time series variation in fuel costs comes mostly from time series variation in (consumer perception of) fuel prices.

While explicitly modeling consumer perceptions of future fuel prices is outside the scope of the paper, we rely on recent evidence in Anderson et al (2013) and Alquist, Kilian and Vigfusson (2011) to focus on fuel prices instead; whereas the former provides evidence that consumers use current prices in their forecasts of future fuel prices, the latter shows that complex models do not outperform simple models with expectations based only on current energy prices. Thus, given that most of the variation in fuel costs comes from fuel prices, our empirical analysis of persistence uses actual fuel prices as a proxy for fuel costs.

Empirically, following a preliminary examination of the series of fuel prices, we estimate parsimonious autoregressive processes for the prices of both diesel and gasoline. It turns out that such models are rich enough to capture the dynamic behavior of fuel prices. Since we would like to account for supply and demand shocks to fuel prices (for instance, due to pass-through from the

[40]In the extreme case that consumers believe that one source of cost faces transitory while the other faces permanent changes, they will fully react to the latter but not to the former.

international oil market), we control for time-varying unobservables in alternative ways using time fixed-effects.

Next, we move on to the analysis of salience. Although not exactly trivial, the calculation of lifetime vehicle taxes, $\tau_j = \tau_j(\delta, s_j, \phi_j)$, is arguably less involving than the calculation of fuel costs. Moreover, in the short-run, and absent a change in the tax code, the only time series variation in vehicle tax is due to the timing of vehicle tax payment (see the Appendix for details). Since tax reforms are infrequent (previous changes in the Swedish market were in 1988 and 2006), there is essentially no time series variation in vehicle tax data.

To better appreciate how our findings relate to salience, consider the model of inattentive agents put forth in Conlisk (1996). The model posits that computing a payoff function is costly for decision-makers in the sense of requiring deliberation due to its very complexity. As a result, the decision-maker applies a costly deliberation technology so that he maximizes the expected payoff net of the cost of deliberating.[41]

Translated to our setting, the correct computation of fuel costs is both more costly (due to its complexity) and also more rewarding to consumers on the market for a used vehicle – recall from Table 2 that the ratio of lifetime fuel costs over lifetime vehicle tax is about 7 at either the mean or median estimates. Since it is not plausible that computing lifetime fuel costs is seven times more costly than lifetime vehicle taxes, it is not unreasonable to think that consumers deliberate more when calculating fuel costs than the vehicle tax given their scarce deliberation time.[42] As a result, the valuation of lifetime fuel costs by consumers is more accurate than that of lifetime vehicle tax, resulting in valuations closer to unity for the former. This is even more likely to be the case for newer (less than 5 years) vehicles which are essentially exempt from vehicle tax (see the Appendix for details).

In practice, the fact that a given component becoming more salient is more likely to make economic agents draw increased attention to it has been studied extensively. For instance, Chetty et al (2009) document that making (sales) taxes more salient reduces the demand for a good whereas Finkelstein (2009) finds that state toll authorities raise tolls more frequently after the introduction of electronic toll collection systems, which make tolls less salient to drivers. More recently, Gallagher and Muehlegger (2011) document that the more salient sales tax rebates had an effect over ten times as large as an equivalent, less-salient, income tax rebate whereas Li et al (2014) find that consumers react more strongly to the more salient changes in fuel taxes as compared to those in net fuel prices.

Empirically, our strategy to account – or proxy – for salience depends on the type of cost considered. In either case, we split a cost component into a standard and a salient component and compare their valuations. In the case of vehicle tax, we exploit the within-year variation induced by different due dates of the vehicle tax to show how increasing the salience of vehicle tax increases its valuation by consumers. For fuel costs, we define standard and salient versions of lifetime fuel costs where the former is switched on whenever a proxy variable is activated. For instance, a salient component of fuel costs consists on the actual fuel costs interacted with an indicator function of oil prices being at least USD 100.

The empirical analysis of salience we perform is essentially reduced-form, in the sense that we do not formulate a structural model of the decision process by consumers active in the used car market. However, in an effort to further document the role of salience and shed light on the potential mechanisms by which salience affects consumers, we replace salience proxies which are a

[41]This model can be cast within the vast bounded rationality literature pioneered by Simon (1955). The intuition underlying such models is that agents face a cost of processing information and thus rationally use simplifying heuristics to tackle complex problems. The framework has been applied in areas such as macroeconomics (e.g. Akerlof and Yellen 1985, Mankiw 1985, Sims 2003, Reis 2006), public finance (e.g. Feldstein 1985, O'Donoghue and Rabin 2006), consumer choice (e.g. Miravete and Palacios-Huerta 2014), environmental economics (e.g. Sallee 2014).

[42]Arguably, in a model with indivisibilities, decision-makers would likely devote even more time to deliberate about fuel costs, thus amplifying the difference in deliberation between fuel costs and vehicle tax.

function of fuel prices with proxies related to news activity and online search activity.

While these mechanisms are likely different from the consumers's standpoint, the results suggest they point in the same direction: news activity (e.g. an indicator for "news about gasoline/gasoline prices is above its long-run average") can arguably be seen as more passive proxy variables in the sense that consumers are likely to be exposed to news via traditional or online media without necessarily looking for any information on fuel prices. This is in stark contrast with the case of online search activity (e.g. an indicator for "search for keyword gasoline/gasoline prices is above its long-run average"), where consumers (potentially after being exposed to news) take an active role to look for (additional) information on fuel prices.

6.1 Persistence

The persistence in fuel prices closely follows that of international oil prices, thus being subject to any shocks facing the international oil market. We examine the persistence of fuel prices by investigating the dynamic behavior of diesel and gasoline prices. To do so, we estimate a number of autoregressive models with alternative controls and sample periods.

<Table 4 here>

The estimates reported in Table 4 use daily recommended fuel prices from Statoil, a leading fuel retailer in the Swedish market, and two sample periods, namely the full sample for which daily fuel prices are available (1st January 2001 to 31st October 2014) and our estimation sample. We report estimates for AR(1) processes, but have also considered alternative and less parsimonious specifications based on more general ARIMA(p,d,q) models; these tend to be more penalized by model selection criteria such as Schwarz's BIC due to the lag of significance of most coefficients with lags beyond one.

Specifications (1) and (2) report estimates for, respectively, gasoline and diesel using the full sample and month-year fixed-effects. In either case, the AR(1) coefficient (0.92 and 0.88 for gasoline and diesel, respectively) is large, suggesting the persistence of the series. Specifications (3) and (4) report results for the same specification using year-week fixed-effects, with substantially lower estimates for the autoregressive coefficients. The estimates suggest that the persistence of the series decreases significantly once unobserved shocks to the oil/fuel market are controlled for.

Specifications (5)-(8) report estimates for the fuel price series using the actual estimation sample. Again, the results are suggestive of high persistence due to the high and statistically significant AR(1) estimates in the case of year-month fixed-effects and decrease substantially if year-week fixed-effects are used instead. As before, the point estimates suggest that gasoline prices are slightly more persistent than diesel ones.

While it is evident that fuel prices – and ultimately fuel costs – exhibit strong persistence, they are still less persistent than vehicle taxes. In fact, since the last changes of the vehicle tax legislation took place in 1988 and 2006 there is just no variation in the vehicle tax of a given car, making the series fully persistent.

Having argued for the stronger persistence of vehicle tax as compared to fuel costs, recall that our estimates point to exactly the opposite of what one would have expected, with a consistently higher valuation to fuel costs as compared to vehicle tax. As a result, persistence does not seem to be the answer for the difference in valuation of the different cost terms.

6.2 Salience

In what follows, we build on our baseline specification to examine the role of salience as a driver of consumer undervaluation of both lifetime fuel costs and lifetime vehicle taxes.[43] We proceed in

[43]Our empirical strategy consists on looking at the fuel costs and vehicle tax separately. Looking at them simultaneously proved challenging given the requirements imposed on the data in what regards the identification of the

two steps. First, we decompose either lifetime fuel costs or lifetime vehicle tax into a "standard" and a "salient" component. While the standard component is the same as in the baseline, the salient component is constructed by interacting the standard component with a proxy variable likely to make a cost component more salient. In the case of vehicle tax, this can be an indicator that the tax payment for a vehicle is due within the next quarter whereas in the case of fuel costs this can be an indicator that oil prices are at least USD 100.

Second, upon re-estimating the baseline where a cost variable comprises both a standard and a salient component, we investigate the extent to which the valuation of the salient component increases with respect to its standard counterpart. If a proxy variable for salience has any explanatory power, it is reasonable to expect that (i) the valuation of the salient component of cost exceeds that of its standard counterpart; and (ii) the valuation of the salient component can help consumers correctly value the cost of interest. In particular, the valuation of the salient component should be closer to unity.

6.2.1 Salience in Vehicle Tax

We rely on the institutional background and examine how consumers react to vehicle taxes as they become increasingly more salient. This can be pursued in two ways, which we will eventually combine. First, we compare the valuation of the vehicle tax component between newer (up to 5 years) and older (from 5 years) vehicles; as the latter are never tax-exempt, their vehicle tax is more salient to consumers upon their purchase.

Second, conditional on not being tax-exempt (alternatively, being at least 5 years out), we compare the valuation of the vehicle tax component as the due date of the vehicle tax approaches. More concretely, we interact the indicator of non-exempt vehicle tax status with indicators of tax being due in the next quarter and in the next month, respectively.

<Table 5 here>

Table 5 reports results for three specifications. Focusing on the valuation parameters, note that there is little variation of the fuel cost parameter (v_c) across specifications – all such estimates are in the range 0.58-0.61 and are moreover similar to the baseline valuation of 0.60. The estimates also share with the baseline the fact that the null of correct valuation is rejected at the 5 percent significance level. While the valuation estimates of the standard component of vehicle tax are largely stable, most of the action happens for the salient component of vehicle tax. Specification (1) distinguishes between newer and older cars, i.e. the salient component switches on for vehicles older than 5 years. Perhaps not surprisingly, the valuation of the salient component is higher than the valuation of the standard one (0.08 vs. 0.01) but the null of equality between them cannot be rejected against a two-sided alternative for standard significance levels.

Specifications (2) and (3) exploit the salience jointly arising from the liability and the timing of taxes. While for Specification (2) the salient component switches on if the vehicle tax is due in the quarter following the vehicle purchase, for Specification (3) it switches on if it is due in the calendar month following the vehicle purchase. In both cases, at 0.07 and 0.10 the standard component of vehicle tax is close to the baseline estimates for the valuation of vehicle tax (0.14). Moreover, the null of correct valuation is rejected at the 1 percent significance level.

Moving to the valuation of the salient component of vehicle tax, we cannot reject the null of correct valuation for neither Specification (2) nor Specification (3) – the corresponding point estimates are 1.01 and 1.66, respectively; while the former essentially means consumer nail exactly the valuation of vehicle tax, the latter may be suggestive of hyperbolic discounting, see e.g. Freder-

parameters in the nonlinear econometric model we estimate.

ick et al (2002) and references therein.[44] Perhaps more importantly, we reject the null of equality between the standard and salient components of vehicle tax at the 1 percent significance level for both specifications.

The above results document that salience plays an important role when it comes to the incidence of vehicle tax. In particular, consumers react in an important way to the timing of vehicle taxes, which suggests that it is in the interest of regulators to highlight both the liability and the timing of tax in the case of durable products such as automobiles. Perhaps more importantly than the undervaluation of non-salient terms, the above findings suggest that tax exemptions tend to depress even further the valuation of tax components and minimally affect the decision to purchase a product. Thus, any policy targets based on vehicle tax are less likely to be achieved than originally planned by the policymaker.

6.2.2 Salience in Fuel Costs

We examine the role of fuel price information in explaining the different valuations of fuel costs and vehicle tax components. To do so, we do as above and split the fuel cost component into a "standard" and a "salient" component and re-estimate our baseline specification. Intuitively, one should expect consumers to react at least as much to the latter than to the former. Moreover, if salience as defined in one of our specifications fully captures the departures from correct valuation, one would also expect the salient component to approach unity.

To narrow the wide possibilities of fuel price information one could use, we resorted to the literature and restricted our analysis to a handful of variables, with results reported in Table 6. In Specification (1), the salient component is switched on whenever oil prices reach at least USD 100.[45] Specification (2) postulates that the salient component switches on whenever fuel prices are within 5 cents of an integer number, e.g. 15 ± 0.05 SEK/liter. Specification (3) switches on separate salient components for price increases and decreases for the months with the 5 percent largest increases and 5 percent largest decreases with respect to the previous month.[46] Finally, Specification (4) switches on the salient component whenever fuel prices are in the fourth quartile of the fuel price distribution, in the spirit of Busse et al (2013).[47]

<Table 6 here>

Focusing on the valuation estimates of vehicle tax (υ_τ) and the standard component of fuel costs (υ_c), the estimates are largely robust when compared to the baseline; the estimates of vehicle tax valuation are in the range 0.05-0.12 (baseline is 0.14) whereas the estimates of fuel costs are in the range 0.54-0.60 (baseline is 0.60). In either case, the null of correct valuation is rejected at least at the 5 percent significance level. The estimates for Specification (1) show that fuel costs are more highly valued if oil prices are at least USD 100. While the increased valuation of the salient

[44]We have also estimated specifications with separate reactions for both the three months prior and the three months following the month the tax is due; we obtained significant estimates only for the two months prior to the due date, with the effect strengthening as the due date approaches.

[45]Our choice of price threshold comes from anecdotal evidence according to which oil prices are more likely to make it to the news when they hit such an important threshold and beyond. For instance, The New York Times explicitly mentions that "(...) oil prices approaching the symbolic threshold of $100 a barrel (...)" and "(...)many analysts expect the psychologically important $100-a-barrel threshold to be breached sometime in the next few weeks." on 9 November 2007, see `http://www.nytimes.com/2007/11/09/business/worldbusiness/09oil.html?_r=2`. Along the same lines, the BBC reports that "The price of Brent crude oil has passed $100 a barrel for the first time since October 2008 (...)" on concerns about the political unrest in Egypt on 31 January 2011, see `http://www.bbc.com/news/business-12328745`. In the last decade or so, Brent crude oil has crossed the USD 100 threshold in most of the three first quarters of 2008, and again in 2011 and most of 2012.

[46]Increasing to the months with the 10 percent largest increases and decreases makes the results stronger.

[47]We have also experimented with decomposing fuel prices into the four quartiles of the fuel price distribution (thus even closer in spirit to Busse et al 2013) with similar results.

component of fuel costs is perhaps not surprising, it is interesting to note that while "standard" fuel costs have a statistically significant standard deviation parameters, suggesting that consumers respond to such changes in a homogeneous way just as in our baseline results, the heterogeneity parameter in the salient component of fuel costs is statistically significant, which is consistent with heterogeneous consumer reactions whenever oil prices cross the USD 100 threshold.

The valuations associated with the demand estimates show marginal changes when it comes to the standard component of fuel costs (decreasing from 0.60 to 0.54) and a substantially higher estimate of 0.85 for the salient component. Importantly, the null of correct valuation is not rejected for the salient component of fuel costs. When comparing valuations, we do reject the null of equal valuation of the salient and standard components of fuel costs (at the 10 percent level) and the null of equal valuation of the standard component of fuel costs and that of vehicle tax.

The estimates for Specifications (2) suggest a similar yet less extreme pattern to those of Specification (1). That is, the standard component of fuel costs and vehicle tax have valuations very similar to those of the baseline specification and the null of correct valuation is always rejected. However, the null of correct valuation cannot be rejected for the salient component.

Specification (3) distinguishes between fuel price increases and decreases. Not only are the point estimates for the valuations of price increases and decreases different, but the former are larger than the latter by a factor of 1.5 (0.73 vs. 0.51, reported in the same row). Moreover, while the null of correct valuation cannot be rejected for the former, it is comfortably rejected for the latter at the 1 percent significance level. When comparing the salient components with their standard counterpart, the findings are intuitive in that one does reject (at the 10 percent significance level) the null of equality between the salient-price increase, component and the standard component, but not the null of equality between the salient-price decrease, component and the standard component. All in all, and consistent with the literature (e.g. Kilian and Sims 2006), consumers tend to value fuel costs more upon price increases than price decreases.

Finally, the estimates for Specification (4) exhibit a similar yet less pronounced pattern where the valuation of the salient term is larger than that of its standard counterpart (0.64 vs. 0.59). However, this time the null of equality of these two components cannot be rejected.

Taken together, we interpret the findings above as evidence that whenever fuel prices (or functions thereof) become more salient, this will result in increased valuations of fuel costs. Frequently, the increased salience might even result in valuations for which the null of correct valuation is not rejected. The relation between the valuation of the standard (non-salient) component of fuel costs and the valuation of vehicle taxes remains largely unaffected.

6.2.3 Salience Induced by News and Web Searches

In the spirit of previous tables, we investigate the role of salience using data from web searches and news related to oil, gasoline and diesel prices. Specifically, we re-estimate our baseline specification adding a term consisting of the interaction of the fuel cost component with a proxy for salience to obtain a salient component of fuel costs.

Table 7 reports estimates related to news activity with the component being salient if the activity is above its mean value for the sample.[48] The standard components of fuel cost valuations are largely similar to the baseline estimate of 0.60, being in the range 0.53-0.70. The null hypothesis of correct valuation for the standard component is rejected at least at the 10 percent significance level, depending on the particular specification. Moreover, the vehicle tax valuations are broadly similar with estimates in the range 0.01-0.10, and the null of correct valuation of vehicle tax is rejected at least at the 10 percent significance level.

<Table 7 here>

[48]We have also estimated the same specification defining the component as salient if it is above the 75 percent percentile activity, obtaining similar results.

The valuation estimates of the salient components of fuel costs vary somewhat; Specifications (1) and (4), which focus on gasoline and oil news activity, respectively, have estimates of 0.95 and 0.89 for which the null of correct valuation is not rejected.

When it comes to Specification (2), which focuses on diesel news, the salient component is actually lower than its standard counterpart (0.58 vs. 0.70), with the null of correct valuation being rejected at the 1 percent significance level. Finally, in the case of Specification (3), which focuses on gasoline or diesel news, the salient component is slightly larger than its standard counterpart (0.64 vs. 0.59). The most robust finding in Table 8 is the fact that the null of equality of standard and salient components cannot be rejected for any specification. Finally, the null of equal valuation of vehicle tax and fuel costs is rejected only in the case of diesel news activity, as per Specification (2).

The higher responsiveness of the salient component in the case of gasoline as compared to that of diesel, and the lack of rejection of the null of correct valuation in the case of gasoline is consistent with demand estimates pointing to a higher price sensitivity of gasoline than diesel, as documented in the Appendix.

Table 8 reports estimates related to searches from Google. As above, we interact fuel costs with an indicator of whether the salient component is switched on. Specifications (1)-(4) switch on the salient component if the keyword search activity is above its long-run average.

<Table 8 here>

Given estimates in the range 0.57-0.61, all standard components of fuel cost valuations are very similar to the baseline estimate of 0.60. The null hypothesis of correct valuation of correct valuation of this component is rejected at either the 5 or 1 percent significance level. Moreover, the vehicle tax valuations are relatively similar, with estimates in the range 0.04-0.07, with the null of correct valuation rejected at the 1 percent significance level.

When it comes to the salient component of fuel costs, Specification (1) focuses on gasoline searches. Looking at the demand estimates, the only significant heterogeneity parameter is exactly that of the salient component in this specification. The associated valuation estimate is 0.70, with the null of correct valuation not being rejected at standard significance levels.

Specifications (2)-(4) focus on diesel, gasoline or diesel, and oil search intensities. With estimates in the range 0.63-0.66, the valuation of the salient components of fuel costs is very similar across the three specifications, and the null of correct valuation is rejected at either the 5 percent (Specification (2)) or 10 percent (Specifications (3)-(4)) significance level. It is only in the case of oil searches – Specification (4) – that the heterogeneity parameter is statistically significant.

Finally, we test the equality between vehicle tax and fuel costs, which is only not rejected for Specification (3), and the equality between standard and salient components of fuel costs, for which the null is never rejected.

Stepping back, we do not see the stronger response to the salient component in the case of gasoline, and the associated lack of rejection of the null of correct valuation as casual. The above findings suggest that the salient components of gasoline and oil prices prices are more pronounced – and captured by search and news activity – than those of diesel and oil prices. As detailed in the Appendix, the empirical evidence from reduced-form demand estimates of fuel sales is that the demand for gasoline is more responsive to prices than the demand for diesel. This could be justified by the fact that diesel drivers – who are typically high-mileage drivers – have a lower rebound effect than gasoline drivers, being less sensitive to both news and search activity.

7 Implications and Discussion

Summary of Findings. We have documented that consumers undervalue both lifetime fuel costs and lifetime vehicle tax, and that the undervaluation of vehicle tax is substantially more severe than that of fuel costs. Our evidence can be rationalized by a model of consumer inattention where the less salient vehicle tax receive less attention (or deliberation) from the part of consumers.

We have also documented a number of cases where we cannot reject the null of correct valuation of fuel costs and vehicle tax once they become more salient. Finally, we have shed light on two potential mechanisms whereby salience of fuel costs may increase, namely increased news about either gasoline or oil prices, and increased online search for gasoline price information; in all such cases, one cannot reject the null of correct valuation of fuel costs.

The particular institutional setting we consider allows the quantification of the valuation of both fuel costs and vehicle tax – a standard-like policy instrument –, two key policy instruments which have been at the center stage of the policy debate in environmental and public economics. This is in stark contrast with most of the literature, whose focus on fuel costs (or fuel taxes) has led to a view that the undervaluation of fuel costs makes standards preferred to taxes. (Parry et al 2007). Our findings suggest that standards are even more undervalued than fuel costs and should be seen as complementary rather than substitutes to the fuel tax, thus departing from the textbook analytical framework typically used to study the efficiency of policy instruments.

Our findings also suggest that while valuations of fuel costs are essentially homogeneous across consumers, the same does not hold for vehicle tax. One potential explanation for this finding are vehicle tax exemptions awarded to new vehicles in recent years in the Swedish market, i.e. vehicle taxes are more salient for some products than others.

When examining the valuation of fuel costs, consumers seem to value them differently depending on the level or changes in fuel prices. We have documented a number of cases where the valuation of fuel costs increased due to more salient fuel prices or functions thereof, from high ($\geq$USD 100) oil prices to large fuel price increases to fuel prices being close to round values. Once again, the findings are consistent with valuations increasing due to salience.

Given the collected evidence suggestive of salience, we investigate two potential mechanisms by which fuel costs might become more salient. While news can be seen as a somewhat passive way whereby salience increases, online search activity (potentially triggered by news itself) can be thought of as one where consumers play a more active role. Our findings suggest that increased news related to oil prices, and both news and search related to gasoline increase the valuation of fuel costs to the point that the null hypothesis of correct valuation cannot be rejected. The fact that similar results do not hold for diesel can be rationalized by a lower rebound effect for this fuel, as we document in the Appendix. In short, diesel drivers are well-known to drive more than their gasoline counterparts, and the empirical evidence we provide shows they are less responsive to any price changes when compared to owners of gasoline vehicles – at least in the short-term.

In sum, we examine the role of salience in the valuation of fuel costs and vehicle tax to find that more often than not, their salient components are not only more highly-valued than their standard counterparts, but the null of correct valuation ($v = 1$) is not rejected. Taken together, we see the results as supporting inattention and salience – more generally, behavioral effects – as the reason behind (the different levels of) undervaluation documented in the data.

Interpretation and Alternative Explanations. The fact that we fail to reject the null hypothesis of correct valuations in a number of situations, and often even close the energy gap, rules out a number of alternative explanations often cited in the literature, see Gerarden et al (2017) for a recent account. Intuitively, for a candidate explanation to be able to explain our findings, such variable has to be correlated with increases in valuations but not to the increase in the salience of fuel costs or vehicle taxes. Such a requirement rules out a number of alternative explanations we now discuss. First, common explanations related to market failures are ruled out due to the fact that these are essentially held constant in our analysis. In particular, information provision about personal utilization and/or peer comparisons (Allcott 2011a; Costa and Kahn 2013; Ayres, Raseman, and Shih 2013; Allcott and Rogers 2014) are unlikely to apply in our setting, as are capital market imperfections and/or liquidity constraints. After all, used vehicles in our data are relatively low-priced, consumers have access to rates which are arguably competitive due to the treatment of personal bankruptcy in Sweden, and our sample period is too short to make changes in capital market imperfections a credible explanation, and even less in times of high salience of policy instruments. Finally, to the extent that imperfect information is time-invariant, it cannot credibly argued to play a major role in our setting.

Second, a number of behavioral explanations are unlikely to explain the gap. For instance, consumer heterogeneity, estimates of energy savings and beliefs of future fuel prices are carefully modeled and/or reasonably assumed constant in our analysis.[49]

Third, justifying low valuations with low fuel prices, low fuel taxes (as in Allcott et al 2014), or the fact that consumers do not take fuel costs into account when purchasing a vehicle (as in Allcott 2011b) is problematic in the Swedish market – Sweden is a market where these variables are among the highest worldwide (and over twice as expensive as in the US, see Section 2) thus making them unlikely to explain the gap.

Finally, when it comes to discount rates, at 9.7 percent per year, the implicit discount rate that leads to the correct valuation of fuel costs in our baseline specification is nearly twice as high as our adopted discount rate, with the implicit discount rate leading to the correct valuation of vehicle tax being even higher.[50]

Quantitative Implications. To provide a sense of the degree of mis-pricing induced by the incorrect valuation of fuel costs and vehicle tax, we calculate car price changes implied by our estimated valuation parameters and compare these to the price changes implied by a correct valuation. We proceed as in Allcott and Wozny (2014) in that we calculate the price change as a hypothetical change in car prices from a 1 SEK increase in either the price of gasoline (Panel A) or diesel (Panel B) measured relative to a car with a fuel consumption of 9 l/100km which emits 210 gCO2/km in the case of gasoline.[51] These levels of fuel consumption and CO2 emissions roughly correspond in our data to a Volvo V70, one of the top-selling car models in Sweden during the sample period. Next, we calculate the hypothetical price change from a 1 SEK increase in the per-gram tax on a gram of CO2 emitted. The results from the fuel price increase are presented graphically in Figure 6 and the corresponding results from the per-gram CO2 tax are presented graphically in Figure 7.

To construct either figure, we assume that the discount rate is the baseline 5 percent, the survival probability is 1, the car is assumed to have 7 years remaining until its terminal age, fuel prices are set to their sample averages (14.73 SEK and 14.60 SEK for gasoline and diesel, respectively), and VKTs are assumed to be 15,000km/year and 20,000km/year for gasoline and diesel cars, respectively.

<Figure 6 here>

[49]The same holds for the modeling of product attributes and the existence of an omitted variable bias; as discussed Section 4, the use of product, fuel and related fixed-effects, in addition to ratings attributed by independent mechanics are likely to mitigate such concerns, on top of being constant in our analysis. In contrast, events such as the Volkswagen scandal, which took place out of the sample period, would likely (dramatically) affect beliefs about energy savings and/or mitigate imperfect information problems, but not in a pattern consistent with the salience of policy instruments.

[50]We have also examined the variation of market discount rates during the sample period by looking at mortgage rates of various maturities, which are highly correlated with interest rates for loans in the used car market, are better documented, recorded at a higher frequency, and have a number of providers. The pattern of such rates is decreasing over time, in stark contrast with the salience patterns observed in either fuel costs or vehicle taxes.

[51]For perspective, 1 l/100km is equivalent to 26.5 gCO2/km in the case of diesel and 23.2 gCO2/km in the case of gasoline. In addition, recall that diesel versions of a given model consume less fuel than their gasoline counterparts, typically resulting in lower CO2 emissions for diesel models. When comparing the two exercises we perform, recall that comparing a 5 l/100km to a 9 l/100km as we do in Figure 6 amounts to comparing vehicles with a difference of roughly 100 gCO2/km. Thus, the monetary values for Figure 7 should be multiplied by 5 to become comparable to those in Figure 6.

From Panel A of Figure 6 we see that a car with a fuel consumption of 5 l/100km would have increased in price by roughly 3,500 SEK, relative to a car with a fuel consumption of 9 l/100km had consumers valued fuel costs at 1. Our baseline valuation estimate of 0.6 suggests however that the adjustment is only around 2,000 SEK. The corresponding numbers for diesel vehicles from Panel B are 4,500 SEK under correct valuation and 2,700 SEK under our estimated baseline valuation parameter for fuel costs. In either case, the mis-pricing corresponds to roughly 2 percent of the cars' estimated price. The difference in price adjustment between gasoline and diesel vehicles comes largely from the higher VKT of diesel vehicles, despite diesel fuel prices being somewhat lower than gasoline prices on average.

<Figure 7 here>

We can perform a similar exercise with vehicle taxes by increasing the per-gram CO2 tax. From Panel A of Figure 7 we see that a car which emits 190 gCO2/km would have increased in price by roughly 115 SEK as compared to a car which emits 210 gCO2/km if the valuation was correct. However, our baseline vehicle tax valuation parameter of 0.14 suggests that the adjustment is only around 16 SEK. From Panel B we see that the corresponding values for diesel vehicles are instead 270 SEK under correct valuation and 40 SEK under our baseline estimated vehicle tax valuation. The higher impact on diesel vehicles stems from a multiplier effect of 2.33 per gram of CO2 as compared to gasoline vehicles (see the Appendix for the exact vehicle tax formulas generating these results).

Policy Implications. Intuitively, economic agents will devote attention to an attribute as long as the associated expected utility exceeds the the cognitive costs of so doing. As a result, agents for which the cognitive costs are high enough will undervalue those attributes – fuel costs and vehicle tax in our setting. Regulators should then formulate policies to address inattention, such as information and/or disclosure programs to target inattentive consumers and enhance the cost-effectiveness of policy instruments. However, we document in our baseline results that undervaluation of fuel costs and vehicle taxes is prevalent in our data, which suggests that even less-targeted policies are likely to improve valuations. First, if consumers are to correctly value any taxes they are due to pay, i.e. internalize any externalities, policymakers should avoid any tax exemptions since they obfuscate important aspects of the tax system to consumers.

Second, allowing the automation of payments and/or making them infrequent should be avoided, as already pointed out in Sexton (2015) for electricity bills; while such measures arguably increase the likelihood of taxes being paid, they also minimize any non-monetary costs associated to them and make them less salient to consumers.

Third, publicizing information about taxes should be encouraged, so it becomes transparent to tax payers which taxes they are due to pay.

Fourth, our findings pointing to the undervaluation of policy instruments – in particular, the fact that standard-like instruments are even more undervalued than fuel costs – suggest the need of a menu of policy instruments for consumers to adequately internalize transport externalities.

Conclusion

This paper examines the energy paradox using unique revealed preference data from retail consumers. in the used car market The use of retail transactions is important due to the potential misleading implications for consumers from analyses based on wholesale transactions. Moreover, the use of consumer data allows to specify and estimate a structural econometric model starting from the primitives of a consumer choice model as well as controlling for heterogeneity at the very micro level. In particular, by considering when individual consumers were actively searching for a used vehicle, we are able to recover the choice sets of each individual consumer, account for time-invariant consumer heterogeneity and estimate the underlying discrete choice model of vehicle purchase.

Our results point to undervaluation of operating costs by consumers. Specifically, consumers undervalue both lifetime fuel costs and lifetime vehicle tax components of the lifetime operating costs of a vehicle. However, the undervaluation of vehicle taxes is substantially more severe than that of fuel costs, which suggests that fuel taxes are more likely to affect consumer behavior. This is important because vehicle taxes are equivalent to standards and have been assumed (even if implicitly) to be correctly valued by both researchers and policymakers.

We document that behavioral explanations, in particular consumer inattention and salience, lie at the root of our findings. Once we introduce a salient component of vehicle tax exploring both tax liability and the timing of tax payment, we cannot reject the null of correct valuation of vehicle tax. This is also the case for the valuation of fuel costs in a number of cases where we construct salient versions of fuel costs. Finally, we provide evidence that certain news- and online search-related variables also have an increased effect on the valuation of fuel costs. All in all, we document a number of stances where increasing salience of policy instruments allows closing the energy efficiency gap.

Our findings suggest that a fuel tax is unlikely to provide a first-best allocation. Rather, policymakers should rely on a menu of instruments in order to make consumers internalize any externalities they create, and make key policy instruments more salient to improve their effectiveness. More concretely, policy measures range from avoiding tax exemptions to make tax payment more frequent and less automated to publicizing information about taxes.

References

[1] Abaluck, J., and J. Gruber (2011). Heterogeneity in Choice Inconsistencies among the Elderly: Evidence from Prescription Drug Plan Choice. *American Economic Review* 101 (3): 377–81.

[2] Akerlof, G. and J. Yellen (1985). Can Small Deviations from Rationality Make Significant Differences to Economic Equilibria. *American Economic Review* 74(4): 708-720.

[3] Alberini, A., M. Bareit, and M. Filippini (2016). What is the Effect of Fuel Efficiency Information on Car Prices? Evidence from Switzerland. *Energy Journal* 37 (3), 315-342.

[4] Allcott, H. (2011a). Social Norms and Energy Conservation. *Journal of Public Economics* 95 (9): 1082–95.

[5] Allcott, H. (2011b). Consumers' Perceptions and Misperceptions of Energy Costs. *American Economic Review* 101 (3): 98–104.

[6] Allcott, H. (2013). The Welfare Effects of Misperceived Product Costs: Data and Calibrations from the Automobile Market. *American Economic Journal: Economic Policy* 5 (3): 30–66.

[7] Allcott, H. and N. Wozny (2014). Gasoline prices, fuel economy, and the energy paradox, *Review of Economics and Statistics* 96 (5), 779-795.

[8] Allcott, H., S. Mullainathan, and D. Taubinsky (2014). Energy Policy with Externalities and Internalities. *Journal of Public Economics* 112: 72–88.

[9] Allcott, H., and T. Rogers (2014). The Short-Run and Long-Run Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation. *American Economic Review* 104 (10): 3003–37.

[10] Alquist, R., L. Kilian, and R. J. Vigfusson (2011). Forecasting the Price of Oil. *CEPR Discussion Paper* #8388.

[11] 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: 383–403.

[12] Anderson, S. T., and J. M. Sallee (2016). Designing Policies to Make Cars Greener: A Review of the Literature. forthcoming in *Annual Review of Economics*.

[13] Atkinson, A. B., and J. E. Stiglitz (1976). The design of tax structure: direct versus indirect taxation. *Journal of Public Economics* 6(1): 55-75.

[14] Ayres, I., S. Raseman, and A. Shih (2013). Evidence from Two Large Field Experiments That Peer Comparison Feedback Can Reduce Residential Energy Usage. *Journal of Law, Economics, and Organization* 29 (5): 992–1022.

[15] Bento, A. M., L. H. Goulder, M. R. Jacobsen, and R. H. von Haefen (2009). Distributional and Efficiency Impacts of Increased US Gasoline Taxes. *American Economic Review* 99 (3), 667-699.

[16] Bento, A., S. Li, and K. Roth (2012). Is there an energy paradox in fuel economy? A note on the role of consumer heterogeneity and sorting bias. *Economics Letters* 115, 44-48.

[17] Berry, S., J. Levinsohn and A. Pakes (1995). Automobile Prices in Market Equilibrium, *Econometrica* 63, 841-890.

[18] Busse, M. R., C. R. Knittel, and F. Zettelmeyer (2013). Are Consumers Myopic? Evidence from New and Used Car Purchases. *American Economic Review* 103 (1), 220-256.

[19] Busse, M. R., N. Lacetera, D. G. Pope, J. Silva-Risso and J. Sydnor (2013b). Estimating the Effect of Salience in Wholesale and retail Car Markets. *American Economic Review: Papers & Proceedings* 103(3): 575-579.

[20] Busse, M., J. Silva-Risso and F. Zettelmeyer (2006). $1000 Cash-Back: the Pass-through of Auto Manufacturing Promotions, *American Economic Review* 96(4): 1253-1270.

[21] Chetty, R., A. Looney and K. Kroft (2009). Salience and Taxation: Theory and Evidence. *American Economic Review* 99(4): 1145-1177.

[22] Conlisk, J. (1996). Why Bounded Rationality? *Journal of Economic Literature* 34 (2): 669-700.

[23] Costa, D. L., and M. E. Kahn (2013). Energy Conservation 'nudges' and Environmentalist Ideology: Evidence from a Randomized Residential Electricity Field Experiment. *Journal of the European Economic Association* 11 (3): 680–702.

[24] Davis, L. W. (2008). Durable Goods and Residential Demand for Energy and Water: Evidence from a Field Trial. *RAND Journal of Economics* 39 (2), 530-546.

[25] Daly, A.J., Hess, S. and K. E. Train (2012). Assuring finite moments for willingness to pay in random coefficients models. *Transportation* 39(1), pp. 19-31.

[26] Dubin, J. A. and D. L. McFadden (1984). An Econometric Analysis of Residential Electric Appliance Holdings and Consumption. *Econometrica* 52 (2), 345-362.

[27] Feldstein, M. (1985). The Optimal Level of Social Security Benefits. *Quarterly Journal of Economics* 100(2): 303-20.

[28] Finkelstein, A. (2009). E-Ztax: Tax Salience and Tax Rates. *Quarterly Journal of Economics*, 124(3): 969-1010.

[29] Frederick, S., G. Loewenstein, and T. O'Donoghue (2002). Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature* 40 (2): 351–401.

[30] Fullerton, D., and S. E. West (2002). Can Taxes on Cars and Gasoline Mimic an Unavailable Tax on Emissions? *Journal of Environmental Economics and Management* 42 (1), 135–157.

[31] Fullerton, D., and S. E. West (2010). Tax and Subsidy Combinations for the Control of Car Pollution. *The B.E. Journal of Economic Analysis & Policy (Advances)* 10 (1).

[32] Gallagher, K. S., and E. Muehlegger (2011). Giving green to get green? Incentives and consumer adoption of hybrid vehicle technology. *Journal of Environmental Economics and Management* 61 (1), 1-15.

[33] Gerarden, T., R. G. Newell, and R. N. Stavins (2015). Deconstructing the Energy-Efficiency Gap: Conceptual Frameworks and Evidence. *American Economic Review: Papers & Proceedings* 105(5): 183–186.

[34] Gerarden, T., R. G. Newell, and R. N. Stavins (2017). Assessing the Energy-Efficiency Gap. Forthcoming in *Journal of Economic Literature*.

[35] Gerlagh, R., I. van den Bijgaart, H. Nijl, and T. Michielsen (2016). Fiscal Policy and CO2 Emissions of New Passenger Cars in the EU. *Environmental and Resource Economics* 1-32.

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

[37] Goldberg, P. K., and F. Verboven (2001). The Evolution of Price Dispersion in the European Car Market. *The Review of Economic Studies* 68 (4), 811-848.

[38] Gowrisankaran, G. and M. Rysman (2011). Dynamics of Consumer Demand for New Durable Goods. Working paper.

[39] Greene, D. L. (2010). How Consumers Value Fuel Economy: A Literature Review. Technical Report EPA-420-R-10-008, U.S. Environmental Protection Agency March 2010.

[40] Grigolon, L., M. Reynaert, and F. Verboven (2015). "Consumer valuation of fuel costs and the effectiveness of tax policy: Evidence from the European car market," KU Leuven Working Paper.

[41] Harberger, A. C. (1964). The Measurement of Waste. *American Economic Review* 54(3): 58-76.

[42] Hausman, J. (1979). Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables. *The Bell Journal of Economics* 10, 33–54.

[43] Helfand, G., and A. Wolverton (2011). Evaluating the Consumer Response to Fuel Economy: A Review of the Literature. *International Review of Environmental and Resource Economics* 5, 103-146.

[44] Hendel, I. and A. Nevo (2006). Measuring the Implications of Sales and Consumer Inventory Behavior. *Econometrica* 74 (6), 1637-1673.

[45] Howarth, R. B. and A. H. Sanstad (1995). Discount Rates and Energy Efficiency. *Contemporary Economic Policy* 13, 101-109.

[46] Hughes, J., C. Knittel, and D. Sperling (2008). Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand. *Energy Journal* 29(1), January 2008.

[47] Huse, C., and C. Lucinda (2014). The Market Impact and the Cost of Environmental Policy: Evidence from the Swedish Green Car Rebate. *Economic Journal* 124, F393-419.

[48] Jaffe, A., and R. Stavins (1994), "The Energy Paradox and the Diffusion of Conservation Technology," *Resource and Energy Economics* 16, 91–122.

[49] Kahn, J. A. (1986). "Gasoline Prices and the Used Automobile Market: A Rational Expectations Asset Price Approach." *Quarterly Journal of Economics* 101 (2): 323–39.

[50] Kilian, L., and E. Sims (2006). "The Effects of Real Gasoline Prices on Automobile Demand: A Structural Analysis Using Micro Data," University of Michigan working paper.

[51] Klier, T. and J. Linn (2010). The Price of Gasoline and New Vehicle Fuel Economy: Evidence from Monthly Sales Data. *American Economic Journal: Economic Policy* 2(3): 134-53.

[52] 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): 212-42.

[53] Knittel, C. R., and R. Sandler (2012). Carbon Prices and Automobile Greenhouse Gas Emissions: The Extensive and Intensive Margins. *In The Design and Implementation of U.S. Climate Policy*, edited by Don Fullerton and Catherine Wolfram. NBER, 287-299.

[54] Langer, A., and N. Miller (2013). Automakers' Short-Run Responses to Changing Gasoline Prices. *Review of Economics and Statistics* 95, 1198–1211.

[55] Li, S., C. Timmins, and R. von Haefen (2009). How Do Gasoline Prices Affect Fleet Fuel Economy? *American Economic Journal: Economic Policy* 1(2): 113-37.

[56] Li, S. J. Linn, and E. Muehlegger (2014), Gasoline Taxes and Consumer Behavior, *American Economic Journal: Economic Policy* 6 (4): 302-42.

[57] Mankiw, G. M. (1985). Small Menu Costs and Large Business Cycles: A Macroeconomic Model. *Quarterly Journal of Economics* 100 (2): 529-538.

[58] Miravete, E. J., M. J. Moral, and J. Thurk (2015). Innovation, Emissions Policy, and Competitive Advantage in the Diffusion of European Diesel Automobiles. *CEPR Discussion Paper* No. 10783.

[59] Miravete, E. J. and I. Palacios-Huerta (2014). Consumer Inertia, Choice Dependence, and Leaning from Experience in a Repeated Decision Problem. *The Review of Economics and Statistics* 96 (3): 524–537.

[60] Mirrlees, J. A. (1971). An Exploration in the Theory of Optimum Income Taxation. *The Review of Economic Studies* 38(2): 175-208.

[61] Newell, R., A. Jaffe, and R. Stavins (1999). The Induced Innovation Hypothesis and Energy-Saving Technological Change. *Quarterly Journal of Economics* 114 (3):941-975.

[62] O'Donoghue, T. and M. Rabin (2006). Optimal Sin Taxes. *Journal of Public Economics* 90(10-11): 1825-1849.

[63] Ohta, M., and Z. Griliches (1986). "Automobile Prices and Quality: Did the Gasoline Price Increhe news articles publishedases Change Consumer Tastes in the U.S.?" Journal of Business and Economic Statistics 4, 187–198.

[64] Parry, I., M. Walls, and W. Harrington (2007). "Automobile Externalities and Policies," *Journal of Economic Literature* 45, 374-400.

[65] Petrin, A., and K. Train (2010). A control function approach to endogeneity in consumer choice models. *Journal of Marketing Research* 47 (1), 3-13.

[66] Ramsey, F. P. (1927). A Contribution to the Theory of Taxation. *Economic Journal* 37(145): 47-61.

[67] Reis, R. (2006). Inattentive Consumers. *Journal of Monetary Economics* 53: 1761- 1800.

[68] Revelt D. and K. Train (1998). Mixed logit with repeated choices: households' choices of appliance efficiency level. *The Review of Economics and Statistics* 80: 647–657.

[69] Sallee, J. (2014). Rational Inattention and Energy Efficiency. *Journal of Law and Economics* 57(3), 781-820.

[70] Sallee, J. M., S. E. West, 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.

[71] Salvo, A., and C. Huse (2013). "Build it, but will they come? Evidence from consumer choice between gasoline and sugarcane ethanol". *Journal of Environmental Economics and Management* 66, 251-279.

[72] Sawhill, J. W. (2008). "Are Capital and Operating Costs Weighted Equally in Durable Goods Purchases? Evidence from the US Automobile Market." Unpublished.

[73] Sexton, S. (2015). Automatic Bill Payent and Salience Effects: Evidence from Electricity Consumption. *Review of Economics and Statistics* 97 (2): 229-241.

[74] Simon, H. (1955). A Behavioral Model of Rational Choice. *Quarterly Journal of Economics* 691(1): 99-118.

[75] Sims, C. (2003). Rational Inattention. Princeton University Working Paper.

[76] Small, K. A., and K. Van Dender (2007). Fuel Efficiency and Motor Vehicle Travel: The Declining Rebound Effect. *The Energy Journal* 28 (1), 25-51.

[77] Train, K. E. (2009). *Discrete Choice Methods with Simulation*, 2nd edition. Cambridge University Press.

[78] Tversky A. (1972). Elimination by Aspects: A Theory of Choice. *Psychological Review* 79: 281– 99.

[79] U.S. Environmental Protection Agency (2014). Guidelines for Preparing Economic Analyses.

[80] U.S. Environmental Protection Agency (2016). Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990–2014. Available from `https://www3.epa.gov/climatechange/Downloads/ghgemissions/US-GHG-Inventory-2016-Main-Text.pdf`

[81] Verboven, F. (1999). "The Markets for Gasoline and Diesel Cars in Europe," Centre for Economic Policy Research discussion paper 2069.

[82] Verboven, F. (2002). "Quality-Based Price Discrimination and Tax Incidence: Evidence from Gasoline and Diesel Cars." *RAND Journal of Economics* 33 (2): 275–97.

Tables and Figures

Table 1: **Sample Choice Set**

Brand	**Model**	**Fuel**	**Price (SEK)**	**Fuel Cost (SEK)**	**Taxes (SEK)**	**Weight (Kg)**	**Age (Yrs)**	**Odometer (km)**	**DV(Purchase)**
Skoda	Octavia	Diesel	99,000	160,608	21,498	1,425	3	181,930	0
Ford	Mondeo	Diesel	69,000	99,973	37,370	1,570	7	104,700	0
Ford	Mondeo	Gasoline	30,000	83,038	14,979	1,520	11	139,320	0
Toyota	Yaris	Gasoline	90,000	79,960	2,947	1,145	4	90,230	0
Toyota	Yaris	Gasoline	90,000	84,578	2,800	1,105	3	75,040	0
Toyota	Yaris	Gasoline	85,000	79,960	2,947	1,145	4	100,480	0
Toyota	Auris	Gasoline	95,000	110,542	10,479	1,380	5	63,130	0
Toyota	Aygo	Gasoline	55,000	66,174	3,102	930	5	53,650	1

Note. This table reports selected variables from the choice set of a randomly chosen user ("-RS6-"). Only auctions in which the user placed at least one bid are shown.

Table 2: **Descriptive Statistics.**

	Mean	Std.	Min.	Q1	Med.	Q3	Max.
Panel A: Bidder Characteristics							
Choice Set (Number of cars)	1,501	2,173	2	138	483	1,864	9,668
Time on Market (Weeks)	3.38	4.28	1	2	2	4	65
Panel B: Vehicle Characteristics							
DV(Diesel)	0.63	0.48	0	0	1	1	1
Price (SEK)	101,759	46,440	5,475	69,350	102,200	131,400	547,500
NPV(Fuel Costs, SEK)	181,886	71,563	30,436	131,397	161,521	226,090	590,729
NPV(Vehicle Tax, SEK)	24,624	14,074	2,922	13,297	22,888	31,585	116,240
Weight (Kg)	1,583	225	911	1,468	1,600	1,690	2,795
Age (Yrs)	4.65	2.71	0.04	3.10	3.67	5.39	19.77
Odometer (km)	115,986	61,230	10	74,500	111,090	146,750	1,105,970
Engine Size (Liters)	2.00	0.43	1.0	1.8	2.0	2.2	6.0

Note. This table reports summary statistics of key bidder and vehicle characteristics. DV(.) denotes dummy variable. Prices and costs are in SEK. For perspective, the SEK/USD exchange rate varied between 6.29 and 7.27 during the sample period, with median, mean and standard deviation given by 6.67, 6.69, and 0.20, respectively.

Table 3: **Demand and Valuation Estimates**

Variables	(1)		(2)		(3)		(4)		(5)		(6)	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Price	-6.46e-06***	-	-6.23e-06***	-	-6.26e-06***	-	-11.95***	0.01	-11.99***	0.01	-11.98***	0.01
	(0.00)		(0.00)		(0.00)		(0.00)	(0.86)	(0.00)	(0.86)	(0.00)	(0.84)
NPV(Fuel+Tax)	-3.24e-06***	-		-			-12.64***	0.02				
	(0.00)						(0.00)	(0.27)				
NPV(Fuel Costs)			-3.88e-06***		-3.83e-06***	-			-12.46***	0.01	-12.49***	0.01
			(0.00)		(0.00)				(0.00)	(0.24)	(0.00)	(0.26)
NPV(Vehicle Tax)					-3.25e-07	-					-15.44***	1.73*
					(0.92)						(0.00)	(0.07)
Average Valuations												
υ_{tc}	0.50***						0.50***					
υ_c			0.62**		0.61*				0.62**		0.60**	
υ_τ					0.05*						0.14***	
υ_τ/υ_c					0.08						0.23**	
Obs.	397,943		397,943		397,943		397,943		397,943		397,943	
Log-Likelihood	-10,861		-10,861		-10,861		-10,861		-10,861		-10,861	

Note. Estimated model is either a Conditional Logit (Specifications 1-3) or a Random Coefficients Logit (RCL) with lognormally distributed parameters on price and all Net Present Value (NPV) terms. The covariance matrix in RCL is constrained to be diagonal as per discussion in text. Auto fixed-effects in every specification include model, fuel type, transmission type, all-wheel drive, turbo charged and year-month fixed effects. Auto controls in every specification include ratings (engine, body, transmission, brakes and interior), engine size, vehicle age, odometer reading and weight. P-values are reported in parentheses. Standard errors are robust and clustered at the consumer level. Significance is given at the 1% level (***), 5% level (**) and 10% level (*). The parameters $\upsilon_{tc}, \upsilon_c, \upsilon_\tau$ denote the valuation of lifetime operating costs, fuel costs, and vehicle tax, respectively. The null hypotheses for all valuation parameters are $H_0 : \upsilon = 1$.

Table 4: **Persistence in Fuel Prices**

Sample	Full sample				Estimation sample			
Fuel Price	Gasoline	Diesel	Gasoline	Diesel	Gasoline	Diesel	Gasoline	Diesel
Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Price_{t-1}	0.92***	0.88***	0.65***	0.60***	0.92***	0.86***	0.70***	0.67***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Month x Year FE's	X	X			X	X		
Week x Year FE's			X	X			X	X
Obs.	5,017	5,005	5,017	5,005	639	639	639	639
R-squared	0.999	0.999	0.999	0.999	0.979	0.963	0.984	0.970

Note. This table displays estimates of persistence of fuel prices for the Swedish market. The dependent variables are daily recommended fuel prices from Statoil, a leading fuel retailer in the Swedish market. The full sample extends from 1st January 2001 to 31st October 2014 whereas the estimation sample matches that of our demand specifications. Price is measured in SEK/liter. Standard errors are robust, p-values of the estimates are reported in parentheses. Significance is given at the 1% level (***), 5% level (**) and 10% level (*).

Table 5: **Salience of Vehicle Tax**

Variables	(1)		(2)		(3)	
	Mean	SD	Mean	SD	Mean	SD
Price	-11.97***	0.01	-11.95***	0.01	-11.96***	0.01
	(0.00)	(0.83)	(0.00)	(0.84)	(0.00)	(0.85)
NPV(Fuel Costs)	-12.47***	0.01	-12.49***	0.01	-12.50***	0.01
	(0.00)	(0.26)	(0.00)	(0.27)	(0.00)	(0.28)
NPV(Vehicle Tax)	-19.10	2.16	-16.04***	1.69	-16.07***	1.89***
	(0.21)	(0.57)	(0.00)	(0.11)	(0.00)	(0.01)
NPV(Vehicle Tax, Age≥ 5)	-15.10**	1.10				
	(0.04)	(0.75)				
NPV(Vehicle Tax, Age≥ 5, Due$_{q+1}$)			-12.03***	0.44		
			(0.00)	(0.64)		
NPV(Vehicle Tax, Age≥ 5, Due$_{m+1}$)					-11.46***	0.12
					(0.00)	(0.59)
Average Valuations						
υ_c	0.61**		0.58**		0.58**	
υ_τ	0.01***		0.07***		0.10***	
$\upsilon_\tau^{Salient}$	0.08***		1.01		1.66	
υ_τ/υ_c	0.01***		0.12***		0.16***	
$\upsilon_\tau^{Salient}/\upsilon_\tau$	8		10***		16***	
Obs.	397,943		397,943		397,943	
Log-Likelihood	-10,861		-10,858		-10,858	

Note. Estimated model is a Random Coefficients Logit (RCL) with lognormally distributed parameters on price and all Net Present Value (NPV) terms. The covariance matrix in RCL is constrained to be diagonal as per discussion in text. Auto fixed-effects in every specification include model, fuel type, transmission type, all-wheel drive, turbo charged and year-month fixed effects. Auto controls in every specification include ratings (engine, body, transmission, brakes and interior), engine size, vehicle age, odometer reading and weight. P-values are reported in parentheses. Standard errors are robust and clustered at the consumer level. Significance is given at the 1% level (***), 5% level (**) and 10% level (*). The parameters $\upsilon_c, \upsilon_\tau$ denote the valuation of lifetime fuel costs, and vehicle tax, respectively. The null hypotheses for all valuation parameters are $H_0 : \upsilon = 1$. The parameter $\upsilon_\tau^{Salient}$ denotes the salient component of lifetime vehicle tax.

Table 6: **Salience of Fuel Costs**

Variables	(1)		(2)		(3)		(4)	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Price	-11.97***	0.01	-11.98***	0.01	-11.98***	0.01	11.98***	0.01
	(0.00)	(0.69)	(0.00)	(0.86)	(0.00)	(0.80)	(0.00)	(0.92)
NPV(Fuel Costs)	-12.60***	0.01	-12.49***	0.01	-12.50***	0.01	-12.50***	0.08
	(0.00)	(0.82)	(0.00)	(0.41)	(0.00)	(0.41)	(0.00)	(0.15)
NPV(Vehicle Tax)	-15.99*	1.36	-15.17***	1.22	-14.68***	1.09	-15.18**	1.21
	(0.06)	(0.51)	(0.01)	(0.46)	(0.01)	(0.24)	(0.02)	(0.81)
NPV(Fuel Costs, Oil $\geq$ \$100)	-12.38***	0.69***						
	(0.00)	(0.00)						
NPV(Fuel Costs, "Round" Prices)			-12.35***	0.04				
			(0.00)	(0.39)				
NPV(Fuel Costs, Large Price Increase)					-12.30***	0.02		
					(0.00)	(0.65)		
NPV(Fuel Costs, Large Price Decrease)					-12.66***	0.01		
					(0.00)	(0.70)		
NPV(Fuel Costs, Price$\geq$75 percentile)							-12.44***	0.14
							(0.00)	(0.48)
Average Valuations								
υ_c	0.54***		0.60**		0.60**		0.59**	
$\upsilon_c^{Salient}$	0.85		0.70		0.73/0.51***		0.64*	
υ_τ	0.05***		0.09**		0.12**		0.09**	
$\upsilon_c^{Salient}/\upsilon_c$	1.58		1.16*		1.21*/0.85		1.08	
υ_τ/υ_c	0.09*		0.14		0.16		0.14	
Obs.	397,943		397,943		397,943		397,943	
Log-Likelihood	-10,858		-10,859		-10,859		-10,860	

Note. Estimated model is a Random Coefficients Logit (RCL) with lognormally distributed parameters on price and all Net Present Value (NPV) terms. The covariance matrix in RCL is constrained to be diagonal as per discussion in text. Auto fixed-effects in every specification include model, fuel type, transmission type, all-wheel drive, turbo charged and year-month fixed effects. Auto controls in every specification include ratings (engine, body, transmission, brakes and interior), engine size, vehicle age, odometer reading and weight. P-values are reported in parentheses. Standard errors are robust and clustered at the consumer level. Significance is given at the 1% level (***), 5% level (**) and 10% level (*). The parameters $\upsilon_c, \upsilon_\tau$ denote the valuation of lifetime fuel costs, and vehicle tax, respectively. The null hypotheses for all valuation parameters are $H_0 : \upsilon = 1$. The parameter $\upsilon_c^{Salient}$ denotes the salient component of lifetime fuel costs. Specification (3) contains NPV of fuel costs split into NPV at time with 5% largest price increases, 5% largest price decreases and remaining periods. Left-most numbers in rows $\upsilon_c^{Salient}$ and $\upsilon_c^{Salient}/\upsilon_c$ are for periods of large increases and right-most numbers for periods with large decreases.

Table 7: **Effect of News Activity**

Variables	(1)		(2)		(3)		(4)	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Price	-11.98***	0.01	-11.96***	0.01	-11.97***	0.01	-11.98***	0.01
	(0.00)	(0.95)	(0.00)	(0.95)	(0.00)	(0.51)	(0.00)	(0.90)
NPV(Fuel Costs)	-12.62***	0.16	-12.32***	0.03	-12.69***	0.61***	-12.49***	0.03
	(0.00)	(0.40)	(0.00)	(0.55)	(0.00)	(0.00)	(0.00)	(0.16)
NPV(Vehicle Tax)	-15.60***	1.03	-23.47	0.64	-14.88***	1.07	-15.25**	1.24
	(0.00)	(0.71)	(0.16)	(0.93)	(0.01)	(0.70)	(0.02)	(0.49)
NPV(Fuel Costs, Gasoline News > Mean)	-12.28	0.72***						
	(0.17)	(0.00)						
NPV(Fuel Costs, Diesel News > Mean)			-13.00***	1.00***				
			(0.00)	(0.00)				
NPV(Fuel Costs, Gas or Diesel > Mean)					-12.43***	0.17		
					(0.00)	(0.48)		
NPV(Fuel Costs, Oil News > Mean)							-12.46***	0.85**
							(0.00)	(0.03)
Average Valuations								
υ_c	0.53***		0.70*		0.59**		0.60**	
$\upsilon_c^{Salient}$	0.95		0.58***		0.64*		0.89	
υ_τ	0.05**		0.01***		0.10**		0.08*	
$\upsilon_c^{Salient}/\upsilon_c$	1.84		0.83		1.59		1.49	
υ_τ/υ_c	0.09		0.01***		0.24		0.14	
Obs.	397,943		397,943		397,943		397,943	
Log-Likelihood	-10,850		-10,855		-10,859		-10,859	

Note. Estimated model is a Random Coefficients Logit (RCL) with lognormally distributed parameters on price and all Net Present Value (NPV) terms. The covariance matrix in RCL is constrained to be diagonal as per discussion in text. Auto fixed-effects in every specification include model, fuel type, transmission type, all-wheel drive, turbo charged and year-month fixed effects. Auto controls in every specification include ratings (engine, body, transmission, brakes and interior), engine size, vehicle age, odometer reading and weight. P-values are reported in parentheses. Standard errors are robust and clustered at the consumer level. Significance is given at the 1% level (***), 5% level (**) and 10% level (*). The parameters $\upsilon_c, \upsilon_\tau$ denote the valuation of lifetime fuel costs, and vehicle tax, respectively. The null hypotheses for all valuation parameters are $H_0 : \upsilon = 1$. The parameter $\upsilon_c^{Salient}$ denotes the salient component of lifetime fuel costs.

Table 8: **Effect of Search Activity**

Variables	(1)		(2)$^{\#}$		(3)		(4)	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Price	-11.97***	0.01	-11.98***	0.01	-11.98***	0.01	-11.98***	0.01
	(0.00)	(0.88)	(0.00)	(0.92)	(0.00)	(0.71)	(0.00)	(0.99)
NPV(Fuel Costs)	-12.48***	0.08	-12.53***	0.35	-12.66***	0.48	-12.48***	0.05
	(0.00)	(0.39)	(0.00)	(0.37)	(0.00)	(0.20)	(0.00)	(0.84)
NPV(Vehicle Tax)	-15.74***	1.44	-16.50*	1.68	-15.35**	1.24	-15.46***	1.33
	(0.01)	(0.31)	(0.09)	(0.49)	(0.02)	(0.46)	(0.01)	(0.34)
NPV(Fuel Costs, Gasoline Search > Mean)	-12.53***	0.65***						
	(0.00)	(0.00)						
NPV(Fuel Costs, Diesel Search > Mean)			-12.45***	0.04				
			(0.00)	(0.38)				
NPV(Fuel Costs, Gas or Diesel > Mean)					-12.41***	0.07		
					(0.00)	(0.12)		
NPV(Fuel Costs, Oil Search > Mean)							-12.50***	0.40**
							(0.00)	(0.04)
Average Valuations								
υ_c	0.60**		0.61**		0.57**		0.61**	
$\upsilon_c^{Salient}$	0.70		0.63**		0.66*		0.64*	
υ_τ	0.07***		0.04***		0.07***		0.07***	
$\upsilon_c^{Salient}/\upsilon_c$	1.18		1.16		1.44		1.06	
υ_τ/υ_c	0.11**		0.08**		0.17		0.12*	
Obs.	397,943		397,943		397,943		397,943	
Log-Likelihood	-10,859		-10,861		-10,859		-10,860	

Note. Estimated model is Random Coefficients Logit (RCL) with lognormally distributed parameters on price and all Net Present Value (NPV) terms. The covariance matrix in RCL is constrained to be diagonal as per discussion in text. Auto fixed-effects in every specification include model, fuel type, transmission type, all-wheel drive, turbo charged and year-month fixed effects. Auto controls in every specification include ratings (engine, body, transmission, brakes and interior), engine size, vehicle age, odometer reading and weight. P-values are reported in parentheses. Standard errors are robust and clustered at the consumer level. Significance is given at the 1% level (***), 5% level (**) and 10% level (*). The parameters $\upsilon_c, \upsilon_\tau$ denote the valuation of lifetime fuel costs, and vehicle tax, respectively. The null hypotheses for all valuation parameters are $H_0 : \upsilon = 1$. The parameter $\upsilon_c^{Salient}$ denotes the salient component of lifetime fuel costs. $^{\#}$Specification 2 does not converge using the baseline number of simulation draws and burn rate. These results are obtained by doubling both the number of repetitions and burn rate to ensure smoother convergence.

Figure 1: **Fuel Price Components**

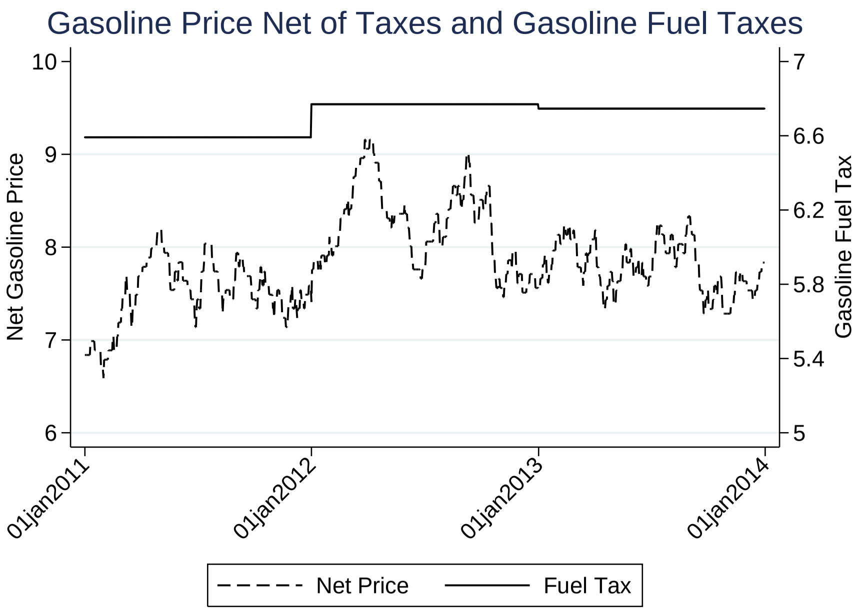

Panel A

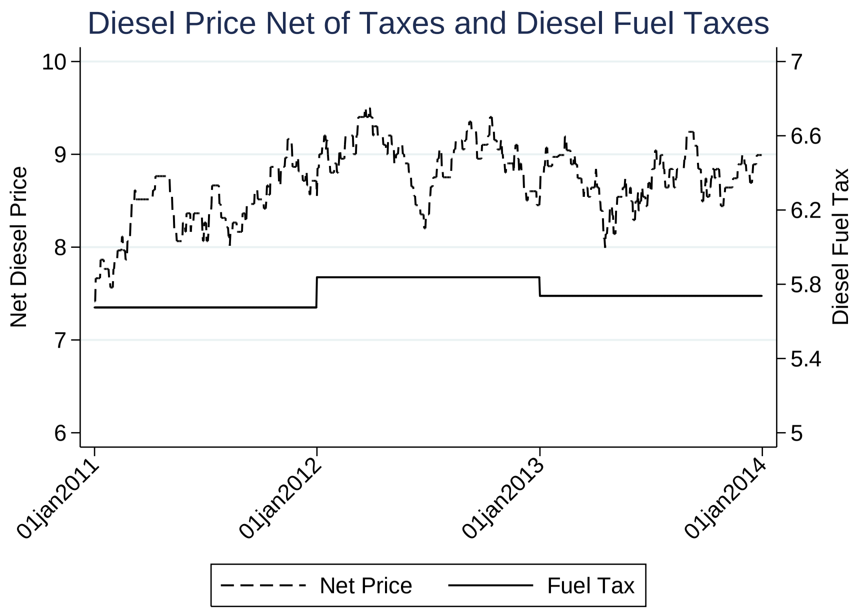

Panel B

Note. This figure displays the time series of net gasoline and diesel prices (exclusive of taxes) together with the time series of the corresponding fuel taxes.

Figure 2: **Share Commanded by Diesel Vehicles in the Passenger Car Fleet – EU Countries (1995, 2009)**

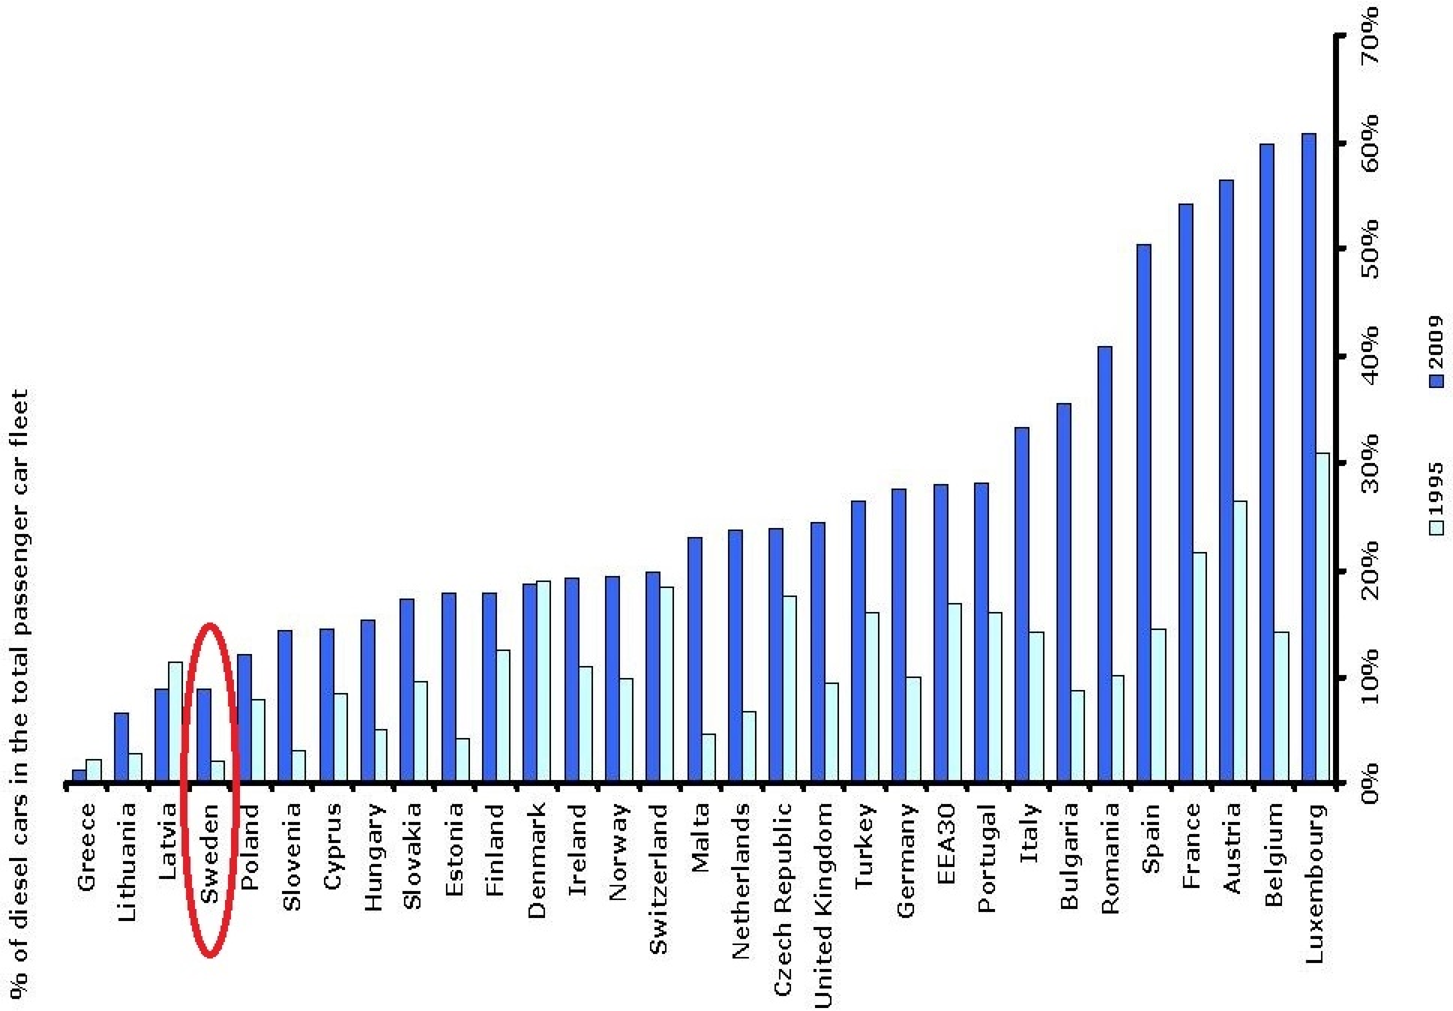

Note. This figure displays the share of the passenger car fleet commanded by diesel vehicles in a number of EU countries in years 1995 and 2009 (Source: http://www.eea.europa.eu/data-and-maps/figures/dieselisation-in-the-eea). Sweden showcases one of the lower penetration of diesel passenger cars which is attributed by market participants to the vehicle tax penalizing diesel vehicles.

Figure 3: **Fuel Prices and Vehicle Taxes**

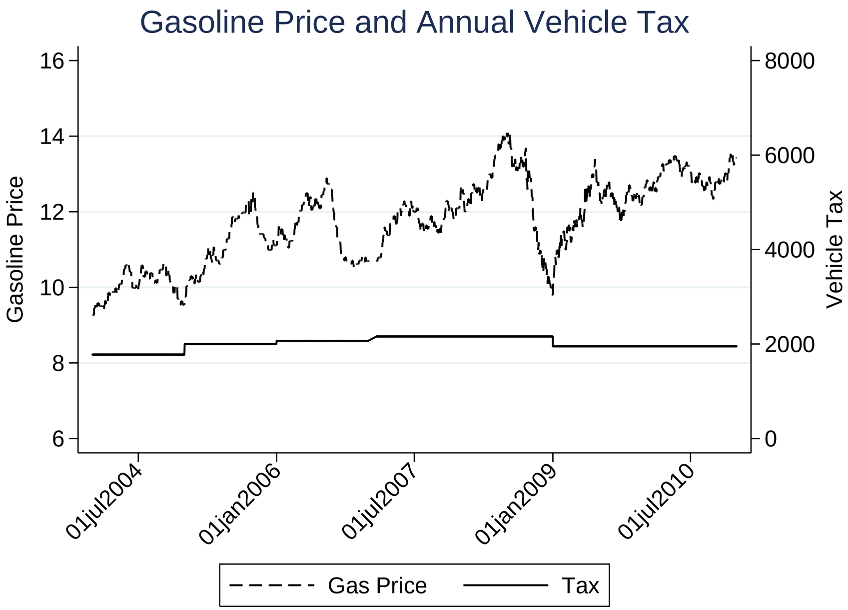

Panel A

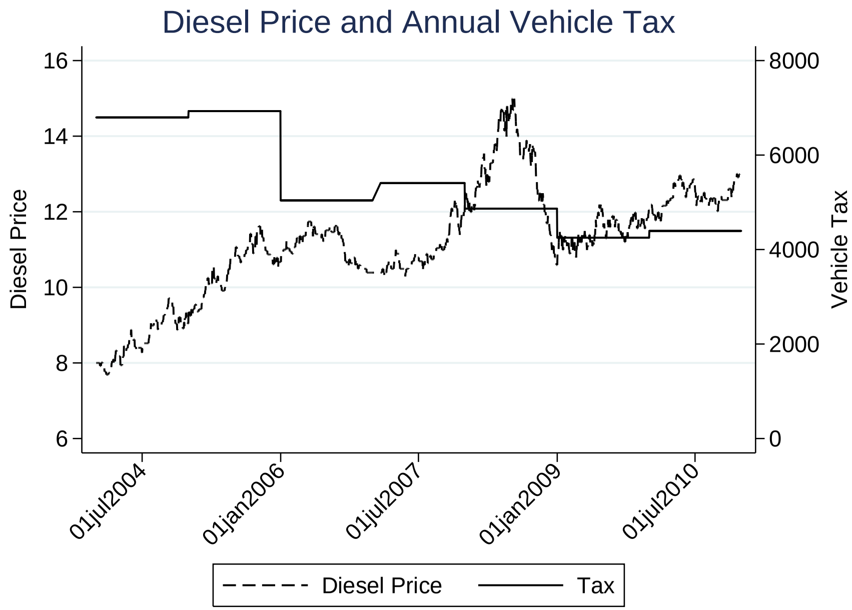

Panel B

Note. This figure displays the time series of gasoline and diesel prices together with the time series of the annual vehicle tax of a baseline model Volvo V70 running on gasoline (Panel A) and diesel (Panel B). Source: New Car Guide.

Figure 4: **Search and News Indices**

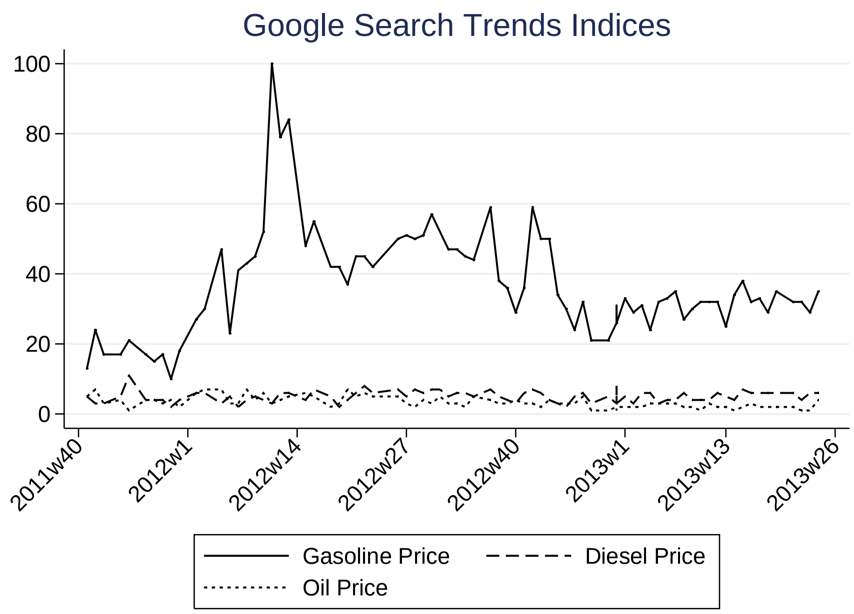

Panel A

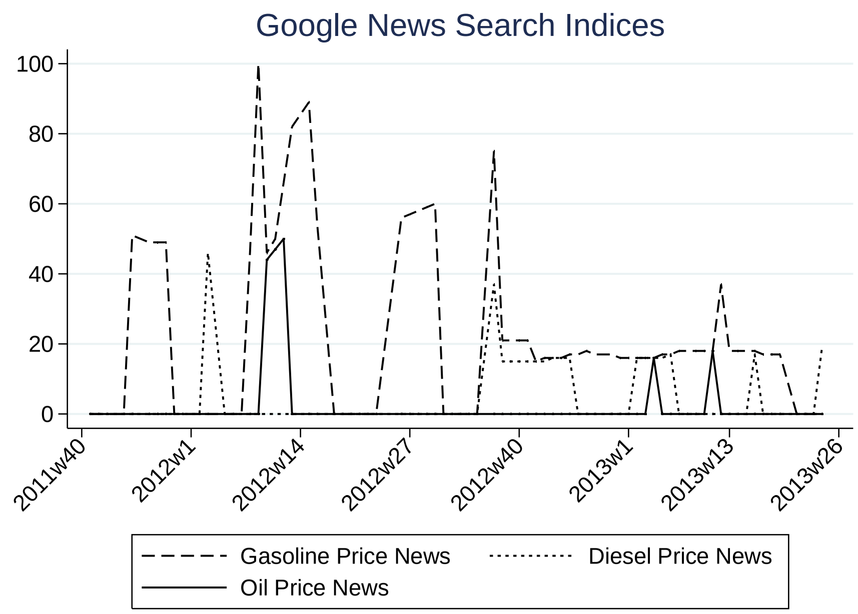

Panel B

Note. This figure displays Google Search Trend Indices (Panel A) and Google News Search Indices (Panel B) at the weekly frequency for the sample period, October 2011 to June 2013.

Figure 5: **Distribution of Valuations of Lifetime Fuel Costs and Vehicle Tax**

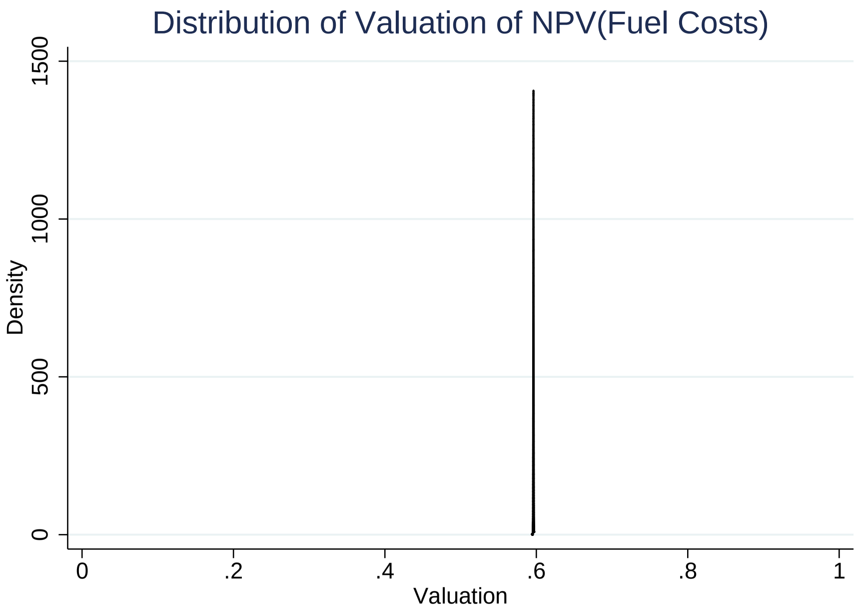

Panel A

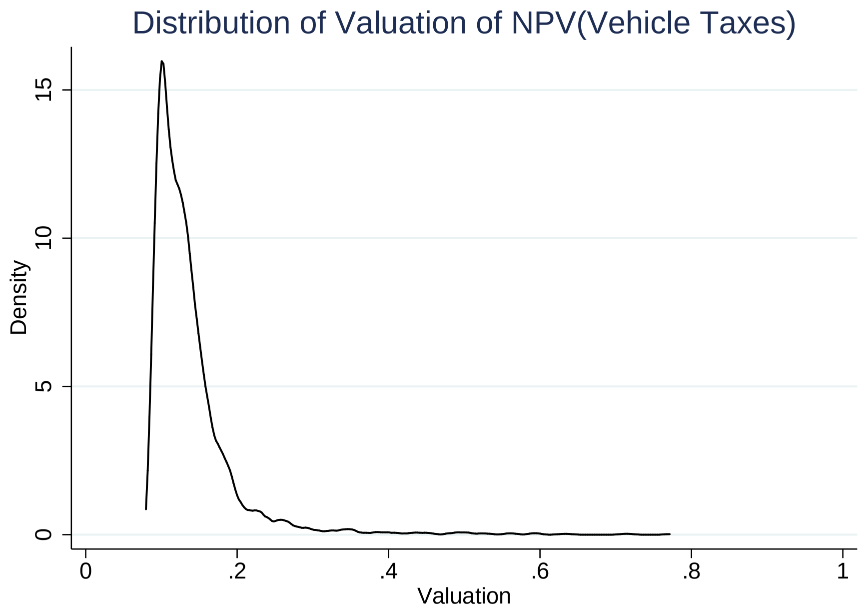

Panel B

Note. This figure displays the distribution of valuations of lifetime fuel costs (Panel A) and vehicle tax (Panel B) as per the baseline specification, Specification (6) in Table 3. The underlying model is a Random Coefficients Logit (RCL) with lognormally distributed parameters on price and all Net Present Value (NPV) terms. The covariance matrix in RCL is constrained to be diagonal as per discussion in text. Auto fixed-effects in every specification include model, fuel type, transmission type, all-wheel drive, turbo charged and year-month fixed effects. Auto controls in every specification include ratings (engine, body, transmission, brakes and interior), engine size, vehicle age, odometer reading and weight.

Figure 6: **Quantitative Implications of Undervaluation**

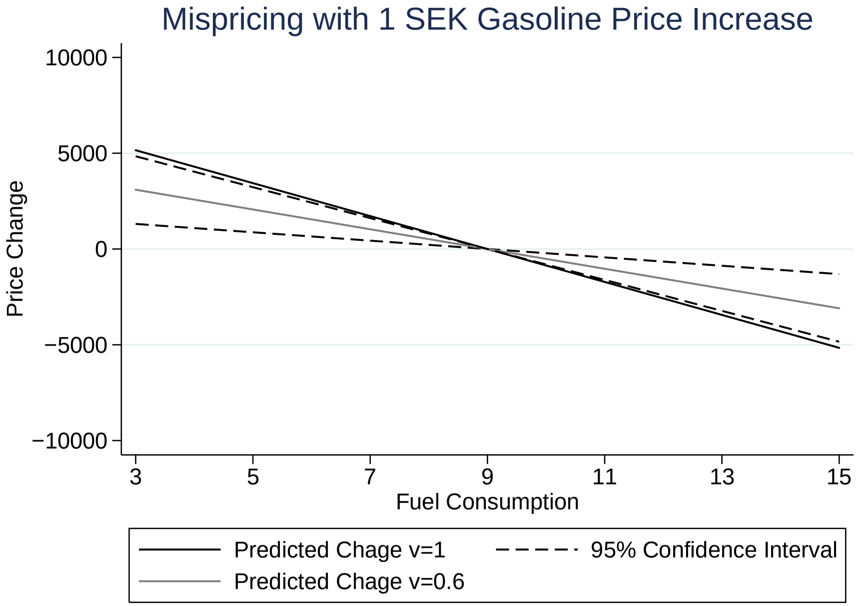

Panel A

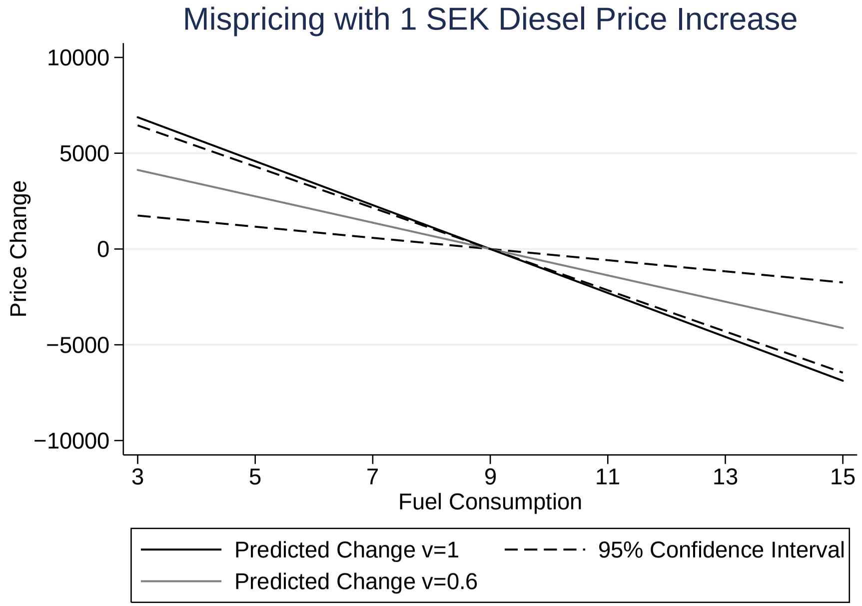

Panel B

Note. This figure quantifies the effect of undervaluation following an increase in 1 SEK/liter on gasoline (Panel A) and diesel (Panel B) vehicles. The price change measures a hypothetical change in car price for a car from a 1 SEK increase in either gasoline (Panel A) or diesel (Panel B) price measured relative to a car with a fuel consumption of 9 l/100km. The discount rate is the baseline 5%, the survival probability is assumed to be 1, the car is assumed to have 7 years remaining until its terminal age, the gasoline price is set to 14.73 and diesel price to 14.60, the respective sample averages, VKT for gasoline cars are assumed to be 15,000km per year and 20,000km for diesel cars.

Figure 7: **Quantitative Implications of Undervaluation**

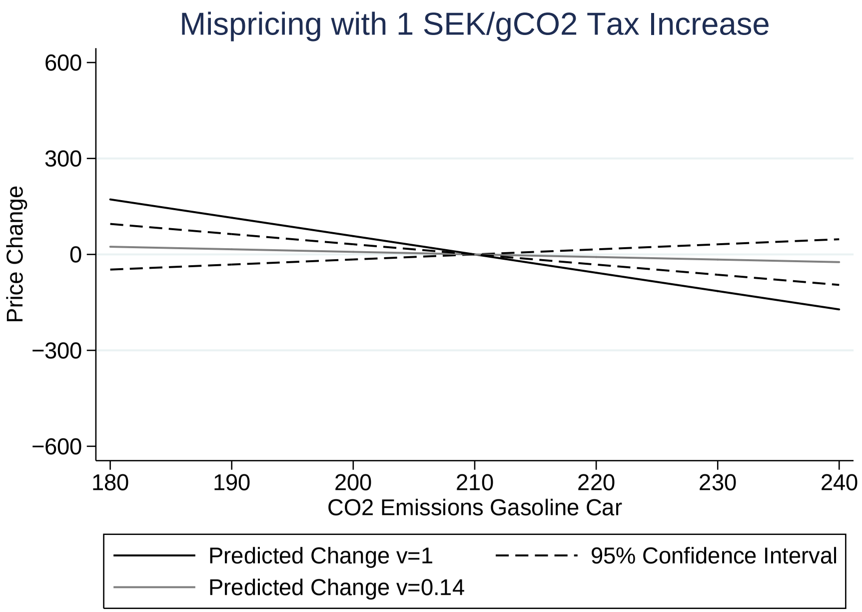

Panel A

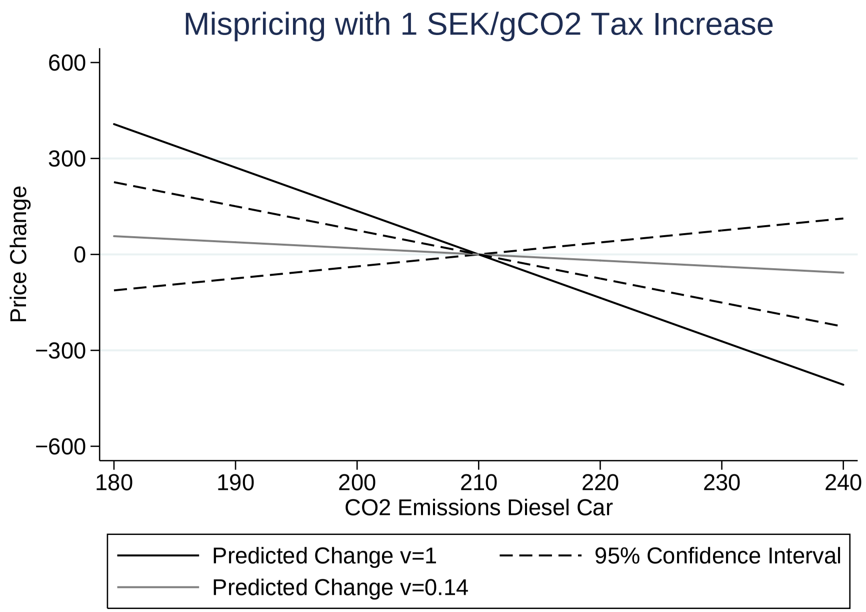

Panel B

Note. This figure quantifies the effect of undervaluation following an increase in 1 SEK/Unit of CO_2 tax on gasoline (Panel A) and diesel (Panel B) vehicles. The price change measures a hypothetical change in car price for a car from a 1 SEK increase in the per-gram tax on a gram of CO_2 relative to a car that emits 210gCO_2/km. The discount rate is the baseline 5%, the survival probability is assumed to be 1, the car is assumed to have 7 years remaining until its terminal age. Vehicle taxes, as a function of CO_2 emissions, are calculated according to the fomulas detailed in Appendix A.2.3.

A Data

A.1 Sample and Choice Sets

From the data that we collect, we restrict our sample in a number of dimensions. First, we restrict our analysis only to users that have actively placed bids in at least two auctions and purchased at most one car. This helps ensure that we are studying the behavior of consumers rather than dealers. Each of these individuals' choice set is the set of all cars that were auctioned in the same weeks in which the individual placed an active bid. We include vehicles into the choice set by weeks because the auction house lists vehicles on a week-by-week basis without the bidder knowing what will be listed the following week. Further, we restrict attention to gasoline and diesel operated cars, excluding ethanol, gas, electric or hybrid cars.[52] Finally, we include only the 100 most popular car models (out of 393) in our sample, which amounts to 87 percent of all cars, in order to keep estimation of the random coefficients logit model with all the included fixed effects tractable. Finally, we obtain a sample size of 397,943 observations constituting 3,150 individual buyers and 12,961 unique auctions.

A.2 Construction of Cost Variables

A.2.1 Vehicle Prices

The auction house from which we have collected our data provides value estimates for every auctioned car. This value is made public at the beginning of each auction and therefore does not depend on the number of bidders or other auction-specific characteristics other than those of the cars themselves. The way the auction house calculates the value estimate is undisclosed but aims to give the buyers a fair and unbiased assessment of the car's value. The value estimates provided are not for the car in its condition at the time of sale but for the car in a hypothetically perfect condition. On average, the transaction price is 27 percent below the auction house estimate because the cars are not sold in perfect condition. As our main price measure, we will use the auction house estimate discounted by 27 percent. We believe this gives the a reasonable estimate of a product's expected price a potential buyer will have had while at the same time not being affected by auction-specific factors such as the number of bidders or the bidding process. This is similar in spirit to the approach of Allcott and Wozny (2014) who are able to use their vehicle auction data in order to predict price estimates for each car, essentially clearing our any auction-specific effects. We are not able to use exactly the same approach due to data limitations and rely instead on the auction house estimate.

A.2.2 Net Present Value of Fuel Costs

The net present value of fuel costs for individual i, car model j is calculated as

$$c_{ij} = \sum_{t=0}^{T} \delta^t s_{jt} \pi_{jt} \phi_j \theta_{ijt}$$

[52]One key reason for excluding these types of vehicles is to be more conservative and look at more simple consumer choice problems; we do not look at cars for which the buyer needs to account for possible fuel switching (see e.g. Huse and Lucinda, 2014) between regular and alternative fuels or all the vehicle tax exemptions that these types of cars enjoyed.

where δ is the annual discount rate, s is the survival probability, π is the fuel price (of the fuel used by vehicle j), ϕ is fuel consumption in liters of fuel per 100km driven and θ is the number of vehicle-kilometers traveled (VKT). To calculate a net present value for each car in an individual's choice set, we must make assumptions about all of these parameters. In our baseline specifications we will assume that the discount rate is 0.05 and we detail the assumptions for survival probabilities and terminal age, fuel prices and VKT below.

Vehicle Kilometers Traveled (VKT). Using data from mandatory vehicle inspections we estimate the logarithmized VKT as a function of car fuel type and age

$$\ln \theta_{jft} = \lambda_{0,f} + \lambda_{1,f}\alpha_{jt} + u_{fjt}$$

where θ_{jft} is the VKT in 10's of kilometers of car j of fuel type f in period t as a function of the car's age at time t, α_{jt}.

The data used to estimate the parameters is taken from the Swedish Automobile Inspection Association (SAIA) and contains average odometer readings disaggregated by model, age, fuel type and year of testing. For instance, the average odometer reading of a 3 year old, gasoline-operated Audi A4 in 2009 was 68,718 km whereas for a 5 year old version it was instead 85,048 km. As we observe only average odometer readings, and only for cars that are 3,5,7,8 or 10 years old, we extrapolate the average VKT as

$$\theta_{\alpha+1,\alpha+\triangle} \approx \frac{O(\alpha+\triangle) - O(\alpha)}{\triangle}$$

where $\triangle$ is the age gap between two consecutive odometer observations and $\theta_{\alpha,\alpha+\triangle}$ is the implied VKT for the car for the ages $\alpha+1$ to $\alpha+\triangle$. Using the previous Audi A4 example,

$$\theta_{4,5} \approx \frac{85,048 - 68,718}{5-3} = 8,165$$

that is, the approximate VKT for a 4 or 5 year old, gasoline-operated Audi A4 is extrapolated to be 8,165km per year. We extrapolate VKT for all models and fuel types up to the age of 10, the oldest age of cars that are reported in the SAIA data, and estimate the λ parameters using this data. The results are presented in the table below and capture the stylized fact in the data that diesel-operated cars drive substantially more than gasoline-operated cars.

With these parameter estimates, for each car in our sample we predict annual VKT as a function of the revealed fuel type of the buyer and the age of the car at purchase. As an additional robustness check, we calculate VKT as a function of the fuel type and effective age of the car in each period as well as a VKT based only on vehicle age.

Table A1: **Parameter Estimates of Mileage (VKT) Regressions**

Car Fuel Type	Gasoline	Diesel
$\hat{\lambda}_0$	7.41***	8.03***
$\hat{\lambda}_1$	-0.03***	-0.05***
N	666	199

Survival Probabilities and Terminal Age. In our baseline specifications, the probability that gasoline and diesel cars survive in any given period is assumed to be exponentially decreasing with a per-period survival probability of 98 percent. This assumption has a number of drawbacks; firstly it oversimplifies the effect of age or usage on the survival probability. In reality, the probability of survival is likely to be relatively high for cars below the age of 5 years and to decrease thereafter. Secondly, it assumes that diesel and gasoline cars have the same probability of survival, something that may not be the case due to mechanical differences in the two types of cars. Furthermore, as diesel cars are used more on average, the probability of being in an accident that leaves the car beyond repair would likely be higher. To try to capture these differences between gasoline and diesel-operated cars, we assume that all cars have a terminal age of 25 years.

Further, as an additional robustness check, and as in Sallee et al (2015), we assume survival probabilities to be a function of the total number of kilometers car j has traveled, as measured by the car's odometer reading O_{jt}, at time t;

$$s_{jt} = \frac{\exp\left(c_0 + c_1 O_{jt}\right)}{1 + \exp\left(c_0 + c_1 O_{jt}\right)} + \epsilon_{jt}$$

To estimate the parameters, we use data from 2011 collected from Bil Sweden on the total number of cars that were scrapped as compared to the total number of cars in the Swedish car fleet disaggregated by car age. For instance, there were 52,986 cars that are 18 years old and of these 11,665 were scrapped, implying a conditional survival probability of roughly 80 percent. We are able to calculate implied survival probabilities conditional on age for cars up to 21 years old with the data available. We then merge these conditional survival probabilities with the average odometer reading, conditional on car age, from the fleet statistics of Statistics Sweden. We estimate the parameters using non-linear least squares and use these to predict survival probabilities conditional on the predicted VKT. Finally, when calculating the net present value of fuel costs or taxes using these predicted survival probabilities, we set the terminal age of all cars to 25 years as in Sallee et al (2015).

Fuel Prices. To model future fuel prices, we assume that fuel prices follow a random walk;

$$\pi_{ft} = \pi_{ft-1} + \varepsilon_{ft}$$

with ε_{ft} being mean zero random shocks that are independently and identically distributed across periods. This assumption is justified by the work of Anderson, Kellogg and Sallee (2013) who find that consumer expectations are in line with fuel prices following a random walk. The assumption implies that the best estimate of the fuel price in any future period is this period's fuel price. In our baseline specifications we use monthly average fuel prices to calculate the net present value of fuel costs and use longer-term averages as robustness checks (see the Appendix on robustness).

A.2.3 Net Present Value of Vehicle Tax

As discussed in Section 2, Sweden changed the way in which vehicle taxes are calculated in 2006, going from a weight-based to a CO2 emissions-based taxation system. The vehicle tax is further differentiated by the type of fuel that the car operates on and the European exhaust emissions standards classification of the car.[53] The exact amount of vehicle tax payable (in SEK) in any given year is determined according to the pre-specified formulas

$$\tau_{Gt} = 360 + 22 \times (\epsilon - 111)$$

for gasoline-operated cars and

$$\tau_{Dt} = 2.33\tau_{Gt} + 250 \times \mathbf{1}[y \geq 2008] + 500 \times \mathbf{1}[y < 2008]$$

[53] The European exhaust emissions standard of each car sold in the European Union is set by the European Commission.

for diesel operated cars, with ϵ being the CO2 emissions of a vehicle measured in gCO2/km and $\mathbf{1}[y \geq 2008]$ an indicator for whether the car was registered for use after the 1st of January 2008. For perspective, 111 gCO2/km is equivalent to 50.7 mpg in the case of gasoline and 56.0 mpg in the case of diesel.[54]

From the above equations, one obtains that the vehicle tax can be written as $\tau_{ft} = \tau_{fn} + \tau_f(\epsilon - \sigma)$, $f = D, G$, where the first right-hand side term consists of a (fuel-dependent) per-vehicle tax, and the second right-hand side term consists of a (fuel-dependent) CO2 emissions standard with target σ. We subsume the per-vehicle tax component into total vehicle tax due to the lack of variation required to identify any associated parameters, as the per-vehicle tax varies across, but not within, fuels.

The month in which the tax is due is determined by the last digit of the vehicle's registration number. For vehicles whose tax owed is in excess of 3,600 SEK, the owner can pay the due amount in three installments. However, only six percent of the vehicles in our sample would qualify for payments in installments. The exact conversion table between the last digit of the vehicle's registration number and the month in which the vehicle tax payment is due is given below.

Table A2: **Timing of Vehicle Tax Payment**

Last Digit of Registration	Due Month	Additional Months
0	March	July and November
1	April	August and December
2	May	September and January
3	June	October and February
4	August	December and April
5	October	February and June
6	November	March and July
7	December	April and August
8	January	May and September
9	February	June and October

Note. Table relating last digit of the vehicle registration number (X, given format ABC 12X) to the month in which the annual vehicle tax payment is due. For vehicles with high annual vehicle tax payments ($>$ 3600 SEK/Year) payments can be made in installments which are due in the months listed in the column "Additional Months". Source: The Swedish Transport Agency.

For each vehicle j we calculate the net present value of vehicle taxes as

$$\tau_j = \sum_{t=0}^{T} \delta^t s_{jt} \tau_{jt}$$

where τ_{jt} is the annual vehicle tax, δ is the discount rate, T is the lifetime of a vehicle and s_{jt} is the vehicle survival probability. The starting period for vehicle tax payments depends on whether the car is exempt from paying taxes before it reaches the age of 5 years. Exemptions change over time and depend both on the fuel the car operates on as well as the emissions standard of the car. Diesel operated cars on the other hand were exempt during our sample period if their emissions standard was "2005 PM". In our sample, we are able to determine the exemption status of most, but not all cars, because the average age of the cars is less than 5 years old and for most cars we observe the

[54]The auction house provides the exact annual vehicle tax value for each car together with the remaining technical information. Therefore, the buyer does not have to make the calculations or know the underlying parameter values. The "diesel factor" was increased from 2.33 to 2.37 in 2015, thus out of the sample period.

exact emissions standard. Furthermore, exemption rules can change every year and determining the exact exemption status of a car can be very difficult. To simplify calculations, we assume that all cars in our sample are exempt for the first 5 years. After the first five years, the vehicle tax is constant and given by the formulas discussed previously.

A.2.4 Search and News Indices

The Internet is a way in which consumers can become informed about fuel and oil prices. Consumers can read about fuel and oil price shocks or price changes through online newspapers and can actively look for information about recommended daily fuel prices through search engines such as Google. As a proxy for how well-informed consumers are about fuel and oil prices at any given point in time, we collect search volume and news publication indices for our sample period from Google Trends. Although we do not argue that search indices necessarily cause consumers to be more informed, for instance because more informed consumers are the ones making follow-up searches, it is likely that fuel prices are more salient when search volumes are high and less salient when search volumes are low. We collect search index values for searches of the keywords “Oil price”, “Gas Price” and “Diesel Price” and news publication indices for the same keywords.

B SML Estimation

Define $\theta := (\beta^*, \sigma, \gamma, \alpha)$ and write $\theta_i = \theta^* + \sigma_\theta \nu_i$. Following the literature (see Train 2009), the estimation of a RC logit model is performed using Simulated Maximum Likelihood (SML), with the log-likelihood function given by

$$lnL(\theta) = \sum_{i=1}^{I} \sum_{j=1}^{J} d_{ij} ln(p_{ij})$$

where d_{ij} is the indicator that consumer i chose product j. In the particular case of lognormally distributed random draws, the choice probabilities take the form

$$p_{ij} := Prob(d_i = j) = \int \frac{exp\,(U_{ij}(\theta))}{\sum_{s=1}^{J} exp\,(U_{ij}(\theta))} \phi(\theta_i | \theta^*, \sigma_\theta)$$

where $\phi(.|.)$ is the density of a lognormal random variable with parameters $\theta^*$and σ_θ. Intuitively, a RC logit consists of a weighted average of logit models, where the weights are given by the distribution of random draws.

The parameter of interest expressing the valuation of energy efficiency is defined as $\upsilon_i := exp(\beta_c^* + \sigma_c \upsilon_{ci})/exp(\beta_p^* + \sigma_p \upsilon_{pi})$ due to the assumption of lognormal draws. The valuation parameter is itself lognormally distributed with mean $(\beta_c - \beta_p)$ and variance $(\sigma_c^2 + \sigma_p^2 + 2\sigma_{pc})$.[55]

In our baseline specifications we estimate the parameters using Stata's mixlogit command and use 50 random draws from a Halton Sequence with a burn rate of 15. To make sure our results are not driven by the number of draws or the burn rate we have rerun our main specification using 300 draws and a burn rate of 50 without any substantial difference in the estimated parameters.

[55]We specify heterogeneity to come from lognormal random draws to avoid ill-defined moments of the distribution of the valuation parameter υ. As pointed out in Daly et al (2012), one alternative parameterization would be to treat the price coefficient as having no heterogeneity and thus divide the numerator mixing distribution by a scalar, but we feel heterogeneity crucial to model the price sensitivity of consumers in a realistic way (in fact, it is shown below to be highly significant).

C Robustness: Alternative Specifications

We estimate a number of alternative specifications, some of which were already put forth in the literature. The results are displayed in Tables C1-C3. (For reference, recall that in our baseline specification the average valuations of fuel costs and vehicle tax are 0.60 and 0.14, respectively.)

Enlarged Choice Sets. In our baseline results we consider the top-100 best-selling vehicles on the auction platform which amounts to 87 percent of all cars auctioned. In this robustness check we consider the top-150 best-selling vehicles, which amounts to 93 percent of all auctions, see Specification (1) of Table C1. The sample increases by slightly over 5 percent as compared to our original sample and the demand estimates are very close to the original ones; the main change is the loss of statistical significance of the vehicle tax heterogeneity parameter. As for the valuations, while the valuation of fuel cost changes slightly, the one of vehicle taxes is more heavily affected, becoming even smaller than our baseline one (0.08 vs. 0.14). Despite any changes, the conclusions of all valuation tests are unchanged when compared to the baseline ones. Our interpretation of the results is that the inclusion of vehicles which are rarer on the Swedish market is to blame, with consumers looking for more specialized or rarer cars having little regard for the valuation of vehicle taxes.

Discount Rate. We change the discount rate from 5 percent to 10 percent and re-estimate our baseline model, see Specification (2) of Table C1. While the price mean and heterogeneity parameter estimates are robust, the same cannot be said about the cost parameters. As a result, the valuation of fuel costs increases from 0.60 in the baseline to 1.09, along the lines of what one would expect, and the null of correct valuation is not rejected anymore. However, the change in the vehicle tax parameter estimates is less dramatic, resulting in a valuation parameter which decreases from 0.14 to 0.06, with significance results unchanged.

The intuition of the above findings is as follows. In the case of fuel costs, increasing the discount rate results in lower lifetime fuel costs, making the valuation parameter increase mechanically. In the case of vehicle taxes we believe the decrease in the valuation to come from the fact that by increasing the discount rate we are effectively reducing the effect of older cars - the cars for which we show that vehicle taxes are more salient. By increasing the discount rate and decreasing the importance of older cars we are decreasing the valuation by increasing the importance of younger cars for which taxes are less salient.

Alternative VKT Estimates I. We use less detailed information of VKT where we disaggregate VKT by fuel and weight (proxy of vehicle class) but not vehicle age. That is, we abstract away from the declining VKT pattern observed during the lifetime of a vehicle, as is often done in the literature. The results are displayed in Specification (3) of Table C1. The parameter estimates are relatively close to those of the baseline: the are slight changes in the mean parameter estimates of both fuel cost and vehicle tax, and the heterogeneity parameter of vehicle tax now becomes insignificant. In terms of valuations, fuel costs are more undervalued while vehicle tax is largely unchanged (0.44 vs. 0.60 and 0.16 vs. 0.14, respectively), with the null of correct valuation being rejected at the 1 percent significance level in either case. This is due to an overall increase in total VKT and therefore an increase in average fuel costs, all else equal – 223,722 SEK for weight-based VKT and 181,886 SEK for our baseline VKT calculation (see Table 2).

Alternative VKT Estimates II. As an additional VKT robustness check we disaggregate VKT by fuel and vehicle age, but not weight. That is, we allow VKT to differ across cars of different fuel types and ages by estimating VKT parameters for both age and fuel as described in Appendix A2. The results are displayed in Specification (4) of Table C1 and are once again similar to those of the baseline model albeit with higher valuations, especially for the valuation of vehicle taxes (0.63 vs. 0.60 and 0.29 vs. 0.14, respectively).

Sampling Weights. We re-estimate our baseline specification after introducing sampling weights, see Specification (1) of Table C2. This is done in order to make our sample more representative of the population (Solon, Haider and Wooldridge 2013). To do so, the weights are such that the market share of each fuel in our data matches the market share of each fuel in the Swedish car fleet. The estimates from the resulting model are extremely close to those of the baseline, with the valuation of fuel costs and vehicle tax now being 0.58 and 0.12, with the null of correct valuation again being rejected at the 5 and 1 percent significance levels. The null of equal valuation of fuel costs and vehicle tax components is rejected at the 10 percent significance level. The fact that the sampling weights do not have a substantial impact on our results suggests that our results are not primarily driven by a larger proportion of diesel vehicles in our sample.

Sticky Fuel Prices. As consumers may well require some time to adjust to new information, we also estimate a specification with sticky fuel prices. That is, for a given month, we assume that consumers are basing their estimates of fuel costs on average fuel prices from the previous quarter. The parameter estimates from this alternative specification are quite similar to those of the baseline, as seen from Specification (2) of Table C2, resulting in valuation parameters that change only slightly: while the valuation of fuel costs increases marginally from 0.60 to 0.62, the valuation of vehicle taxes decreases from 0.14 to 0.02.

Additional Consumer Characteristics. We add region fixed-effects of the locations where the products are on sale, see Specification (3) of Table C2. In the former case, the valuations of 0.53 and 0.16 are very close to those in the baseline.

Income. One concern regarding our baseline specification comes from the fact that low- and high-income consumers likely have different price sensitivities; the use of a utility function linear in price would imply a larger price elasticity (in magnitude) for high-end vehicles, a counter-intuitive pattern. To address this concern, we estimate a logit model with a price-income interaction term, with income taken as the average income within a region where a consumer purchased his vehicle.

The results reported in Specification (4) of Table C2 document that the valuation estimates are largely consistent with our baseline, at 0.59 and 0.03 (as compared to 0.60 and 0.14), rejecting the null of correct valuation at standard significance levels.

Our interpretation of the above findings for the price-income effect is that income plays a limited role in the used car market; although Table 2 reports that our sample comprises vehicles priced in the range 5,000-547,000 SEK, most cars are priced between 69,000 SEK (25th percentile) and 131,000 SEK (75th percentile), so income does not likely play a role as important in the used car market as in the new car market. Another factor likely dampening the effect of income is the fact that it is possible to finance a used car without major problems at competitive rates, even directly from the auction platform.

Alternative and/or Additional Characteristics. We estimate a specification where we changed product characteristics; we have replaced vehicle weight with vehicle size (length $\times$ width) and additionally replaced engine size with the ratio of HP (engine power) and weight. We obtain robust parameter estimates resulting in valuations for 0.64 and 0.22 for fuel costs and vehicle taxes, respectively. Both valuation parameters are statistically significantly different from 1 at the 5 percent significance level as seen in Specification (1) of Table C3.

Vehicle Hazard Rates. We re-estimate our baseline model including odometer-based survival probabilities to more accurately account for scrapping, see Specification (2) of Table C3. Intuitively, this amounts to the assumption that car quality degrades with usage rather than a fixed fraction every period, as in our baseline specification. That is, the vehicle survival probability is declining in the total VKT but is independent of age and other vehicle characteristics.[56]

[56] Ideally, we would like to allow survival probabilities to be as general as possible but our data only permits estimating a survival probability that is heterogeneous only in the total VKT traveled by a given car.

While the parameter estimates for price, fuel costs and vehicle taxes are of similar magnitude as compared to those of the baseline specification, when it comes to the valuations, there is an increase in the undervaluation of fuel costs (from 0.60 to 0.55) and a slight increase in the valuation parameter of vehicle (from 0.14 to 0.22). For both valuations, we

reject the null of correct valuation – the valuation of fuel costs is statistically significant at the 1 percent significance level and that of vehicle taxes is statistically significant at the 5 percent significance level. The results suggest that assuming that the survival probability is exponentially declining, as in our baseline specification, is a reasonable approximation to more complicated functional forms that account explicitly for car usage.

Correlated Random Draws. We re-estimate our baseline specification allowing for correlated draws between price and each of the two cost components, see Specification (3) of Table C3. The main changes with respect to the baseline happen in the vehicle tax parameters. While this does not seem to affect the valuation of fuel costs (0.62 vs. 0.60 in the baseline), it substantially reduces the valuation of vehicle tax (0.02 vs. 0.14 in the baseline). We draw the following lessons from this exercise. First, heterogeneity seems to manifest itself only when it comes to the vehicle tax component, despite rich sources of variation in the other components. Second, even with a more general model (three additional parameters), the changes in valuations are minimal. These factors combined make us settle for the constrained version of the RCL as our baseline.

Discussion. We present a number of alternative specifications to our baseline specification. These changes range from replacing our baseline discount rate with an alternative one to changes in characteristics to the inclusion of additional fixed-effects and additional random coefficients, an issue especially demanding for our econometric methodology. Our first take-away from this exercise is that both fuel costs and vehicle taxes remain undervalued, which allows us to comfortably reject the null of correct valuation for both components.

The second lesson we learn from this exercise is that fuel costs remain more highly valued than vehicle taxes, despite any changes in mean valuations that may occur – and the rejection of the null hypotheses of equal valuations does hold in a number of robustness exercises.

Finally, the valuation of vehicle taxes seems more robust than that of vehicle taxes, thanks to the richer variation identifying the underlying coefficients of the former. In fact, identifying the vehicle tax parameter tends to be a challenge unless we distinguish between vehicles which are tax-exempt from those which are not. At the root of this problem lies the lack of variation of vehicle taxes over time following the 2006 tax reform.

Conditional on a 5 percent discount rate, our alternative specifications provide valuation estimates in the range 0.41-0.70 for fuel costs (baseline is 0.60) and in the range 0.03-0.22 for vehicle taxes (baseline is 0.14). The null of correct valuation is rejected in all cases for vehicle tax and in all but one case for the fuel cost component.

While we are cautious in not claiming outright robustness of our baseline specification, ours is not the only case where valuations change across specification. For instance, despite a substantially larger dataset in which a number of features are taken to be model-vintage averages rather than the raw characteristics as in our case, Allcott and Wozny (2014) provide an extensive set of robustness checks where there is substantial variation in their valuation estimates (0.58-0.87 (0.46-0.68) for their preferred discount rate following their baseline of 0.76 (0.55) for expectations of fuel prices following futures prices (martingales), see their Table 4). Similarly, the estimates of the valuation parameter in Sallee et al (2016) vary in the range 0.48-0.98, see their Table 5.

Table C1: **Alternative Demand and Valuation Estimates I**

	Enlarged Choice Sets		Disc. Rate $r = 10\%$		Weight-based VKT		Fuel-Age VKT	
Variables	(1)		(2)		(4)		(5)	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Price	-11.99***	0.04	-11.93***	0.01	-11.93***	0.01	-11.99***	0.01
	(0.00)	(0.21)	(0.00)	(0.76)	(0.00)	(0.88)	(0.00)	(0.91)
NPV(Fuel Costs)	-12.42***	0.01	-11.84***	0.01	-12.76***	0.04	-12.47***	0.03
	(0.00)	(0.92)	(0.00)	(0.22)	(0.00)	(0.56)	(0.00)	(0.64)
NPV(Vehicle Tax)	-14.58**	0.08	-17.46 ***	2.33**	-15.07***	1.61	-14.05***	1.29
	(0.03)	(0.97)	(0.00)	(0.02)	(0.00)	(0.11)	(0.00)	(0.23)
Average Valuations								
υ_c	0.66*		1.09		0.44***		0.63	
υ_τ	0.08*		0.06***		0.16***		0.29	
υ_τ/υ_c	0.04		0.06***		0.36		0.47	
Obs.	419,149		397,943		397,943		397,943	
Log-Likelihood	-11,217		-10,854		-10,866		-10,866	

Note. Estimated model is a Random Coefficients Logit (RCL) with lognormally distributed parameters on price and all Net Present Value (NPV) terms. The covariance matrix in RCL is constrained to be diagonal as per discussion in text. Auto fixed-effects in every specification include model, fuel type, transmission type, all-wheel drive, turbo charged and year-month fixed effects. Auto controls in every specification include ratings (engine, body, transmission, brakes and interior), engine size, vehicle age, odometer reading and weight. P-values are reported in parentheses. Standard errors are robust and clustered at the consumer level. Significance is given at the 1% level (***), 5% level (**) and 10% level (*). The parameters $\upsilon_c, \upsilon_\tau$ denote the valuation of lifetime fuel costs, and vehicle tax, respectively. The null hypotheses for all valuation parameters are $H_0 : \upsilon = 1$.

Table C2: **Alternative Demand and Valuation Estimates II**

	Fuel-Weighted		Previous Quarter Fuel Prices		Region FEs		Income#	
Variables	(1)		(2)		(3)		(4)	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Price	-11.92***	0.01	-11.97***	0.01	-11.88***	0.01	-8.61e-06***	-
	(0.00)	(0.66)	(0.00)	(0.84)	(0.00)	(0.84)	(0.00)	-
NPV(Fuel Costs)	-12.47***	0.02	-12.46***	0.01	-12.52***	0.01	-3.84e-06***	-
	(0.00)	(0.16)	(0.00)	(0.22)	(0.00)	(0.23)	(0.00)	-
NPV(Vehicle Tax)	-15.49***	1.76**	-15.64***	1.78*	-14.81***	1.48	-2.10e-07	-
	(0.00)	(0.03)	(0.00)	(0.06)	(0.00)	(0.18)	(0.95)	-
Price*Income							7.70e-09*	-
							(0.06)	-
Average Valuations								
υ_c	0.58**		0.62**		0.53**		0.59**	
υ_τ	0.13***		0.02***		0.16***		0.03*	
υ_τ/υ_c	0.23**		0.20**		0.30		0.06	
Obs.	397,943		397,943		397,943		397,943	
Log-Likelihood	-18,148		-10,860		-10,853		-10,859	

Note. Estimated model is a Random Coefficients Logit (RCL) with lognormally distributed parameters on price and all Net Present Value (NPV) terms. The covariance matrix in RCL is constrained to be diagonal as per discussion in text. Auto fixed-effects in every specification include model, fuel type, transmission type, all-wheel drive, turbo charged and year-month fixed effects. Auto controls in every specification include ratings (engine, body, transmission, brakes and interior), engine size, vehicle age, odometer reading and weight. P-values are reported in parentheses. Standard errors are robust and clustered at the consumer level. Significance is given at the 1% level (***), 5% level (**) and 10% level (*). The parameters $\upsilon_c, \upsilon_\tau$ denote the valuation of lifetime fuel costs, and vehicle tax, respectively. The null hypotheses for all valuation parameters are $H_0 : \upsilon = 1$. Symbol # denotes that valuations are calculated at the median income (fuel cost and vehicle tax valuations at the first and third quartiles are $0.58^{**}, 0.61^{**}$and $0.03^{***}, 0.03^{***}$).

Table C3: **Alternative Demand and Valuation Estimates III**

	Size		Odometer Survival#		Unconstr.			
Variables	(1)		(2)		(3)			
	Mean	SD	Mean	SD	Mean	Covariances		
Price	-11.96***	0.01	-11.95***	0.01	-11.98***	0.01		
	(0.00)	(0.86)	(0.00)	(0.85)	(0.00)	(0.94)		
NPV(Fuel Costs)	-12.41***	0.01	-12.55***	0.02	-12.47***	-0.01	0.01	
	(0.00)	(0.25)	(0.00)	(0.21)	(0.00)	(0.74)	(0.25)	
NPV(Vehicle Tax)	-14.56***	1.47	-14.33***	1.31	-30.16***	2.50***	-0.95***	5.78***
	(0.00)	(0.15)	(0.00)	(0.19)	(0.00)	(0.00)	(0.00)	(0.00)
Average Valuations								
υ_c	0.64**		0.55***		0.62***			
υ_τ	0.22**		0.22**		0.02***			
υ_τ/υ_c	0.35		0.40		-			
Obs.	397,943		397,943		397,943			
Log-Likelihood	-10,861		-10,857		-10,860			

Note. Model type is a Random Coefficients Logit (RCL) with lognormally distributed parameters on price and all Net Present Value (NPV) terms. Except for Specification 3, which displays an unconstrained covariance matrix, the covariance matrix in RCL is constrained to be diagonal as per discussion in text. Auto fixed-effects in every specification include model, fuel type, transmission type, all-wheel drive, turbo charged and year-month fixed effects. Auto controls in every specification include ratings (engine, body, transmission, brakes and interior), engine size, vehicle age, odometer reading and weight. P-values are reported in parentheses. Standard errors are robust and clustered at the consumer level. Significance is given at the 1% level (***), 5% level (**) and 10% level (*). The parameters $\upsilon_c, \upsilon_\tau$ denote the valuation of lifetime fuel costs, and vehicle tax, respectively. The null hypotheses for all valuation parameters are $H_0 : \upsilon = 1$. Symbol# denotes that the Specification contains a third order polynomial in odometer reading.

D Details on Salience

Table D1 documents the higher price sensitivity of gasoline sales as compared to diesel sales by means of a reduced-form demand equation. This is consistent with a higher valuation of the salient component in the case of gasoline as compared to diesel found in our results.

We estimate demand equations for diesel and gasoline using monthly data for the period January 2005-October 2014. The dependent variable is the logarithm of sales of either diesel or gasoline (measured in cubic meters) in Sweden. Price is measured in SEK/liter. Specifications (1) and (2) report OLS estimates with month fixed-effects to control for seasonality. The point estimates suggest that gasoline is substantially more responsive than diesel to prices (-0.39 vs. -0.25), with both estimates being statistically significant at the 1 percent significance level.

As fuel prices are likely endogenous with respect to fuel sales, Specifications (3) and (4) report IV estimates using Brent oil prices to instrument for diesel and gasoline prices.[57] As before, month fixed-effects control for seasonality. The resulting IV estimates are similar to their OLS counterparts in that the demand for gasoline remains more responsive than that of diesel to changes in the respective prices (-0.35 vs. -0.23), with both estimates significant at the 1 percent level.

Finally, Specifications (5) and (6) report IV estimates with hours worked added as a proxy for income. The estimates are similar to the previously reported ones and once again the demand for gasoline is more responsive to prices than that of diesel (-0.35 vs. -0.22), with both estimates being significant at the 1 percent level. The rebound effects implied by these estimates are -2.4e-4 and -3.5e-5 for gasoline and diesel, respectively, when evaluated at mean values (-3.5e-4 and -3.2e-5 when evaluated at median values).

Intuitively, the findings are consistent with evidence that diesel drivers drive more and are less responsive to prices than gasoline drivers.

[57] Oil prices are quoted in SEK/barrel; the results are virtually unchanged to the use of USD/barrel.

Table D1: **Fuel Demand Estimates – Reduced-form Evidence**

	OLS		IV		IV	
	Diesel	Gasoline	Diesel	Gasoline	Diesel	Gasoline
Variables	(1)	(2)	(3)	(4)	(5)	(6)
Price	-0.25***	-0.39***	-0.23***	-0.35***	-0.22***	-0.35***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Month FE's	Yes	Yes	Yes	Yes	Yes	Yes
Hours					Yes	Yes
Obs.	117	117	117	117	117	117
R-squared	0.50	0.80	0.50	0.80	0.50	0.80

Note. This table displays estimates of fuel demand for diesel and gasoline for the Swedish market (sample period is January 2005-October 2014). The dependent variable is the logarithm of monthly fuel sales (either diesel of gasoline), measured in cubic meters. Price is measured in SEK/liter; the income proxy is Hours, the number of hours actually worked per week for persons aged 15-74 (a good proxy especially under rigid labor contracts). Both diesel and gasoline prices are instrumented with Brent oil prices measured in SEK/barrel (using USD/barrel only marginally affects the results). Standard errors are robust, p-values of the estimates are reported in parentheses. Significance is given at the 1% level (***), 5% level (**) and 10% level (*).