Salience and Policy Instruments: Evidence from the Auto Market

Cristian Huse, Nikita Koptyug

Abstract: We take advantage of a unique institutional setting that allows consumers to separately value fuel and vehicle (or road) taxes. We estimate a structural model of vehicle choice using consumer-level revealed preference and find that consumers undervalue both policy instruments, but undervaluation of the latter 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 then 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 policy instruments more salient to consumers.

JEL Codes: D12, L62, Q40, Q41, Q48, Q50, Q58

Keywords: automobiles, energy efficiency, energy paradox, energy efficiency gap, fuel tax, standard, fuel economy, CO_2 emissions, vehicle tax, road tax

ELECTRICITY PRODUCTION AND TRANSPORTATION are major sources of greenhouse gas emissions (GHGs) in the United States and the European Union, at over 50%

Cristian Huse is at the University of Oldenburg, Department of Economics, Applied Microeconomics, Ammerlaender Heerstr. 114–118, 26129 Oldenburg, Germany (cristian.huse@uol.de; http://cristianhuse.webs.com). Nikita Koptyug is at the Research Institute of Industrial Economics (IFN), Grevgatan 34, Box 55665, 102 15 Stockholm, Sweden (nikita.koptyug@ifn.se). We thank seminar participants at AUROE, CEMFI, the Economics of Low-Carbon Markets Workshop, ETH, IFN, Institut d'Economia de Barcelona, KAPSARC, the International Industrial Organization Conference (IIOC), the Toulouse Conference on the Economics of Energy and Climate, and the World Congress for Environmental and Resource Economists (WCERE) for

Dataverse data: https://doi.org/10.7910/DVN/2VJRQZ

Received July 21, 2019; Accepted August 13, 2021; Published online February 7, 2022.

Journal of the Association of Environmental and Resource Economists, volume 9, number 2, March 2022.

https://doi.org/10.1086/716878

and 33%, respectively (Eurostat 2021; US EPA 2021). Driving energy consumption in both sectors are 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 with their use are a nontrivial share of a typical household's expenditure raises the question of whether consumers take into account future operating costs when they purchase such products.

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 (expected) lifetime operating costs is 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 with regulation and expected consumer behavior.

This paper examines 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 secondhand vehicle upon its purchase—our null hypothesis is that consumers value them correctly. Under standard assumptions in applied welfare analysis (Phaneuf and Requate 2017), 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 CO_2 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 on consumer decisions. Despite recent empirical findings casting doubt on this assumption, many central theoretical results in public economics rely on consumers being full optimizers 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 in the market we study. Thus,

feedback. We are also indebted to Reyer Gerlagh, Shanjun Li, Shaun McRae, Sebastian Schwenen, Ulrich Wagner, and two referees for insightful comments. The usual disclaimer applies.

if valuations differ, and these differences can be attributed to differences in salience, policy makers would improve consumer welfare and the effectiveness of policy instruments by increasing their salience.

1. OVERVIEW

1.1. Empirical Strategy

To examine whether consumers correctly value the lifetime operating costs of durable products, we start from the primitives of a consumer choice model of rational behavior for individual consumers. We construct a data set of individual retail-level transactions in the used car market to examine our hypotheses of interest. Our revealed preference data come from used car auctions; we follow the activity of individual consumers over time, which enables us to construct their choice sets.

We specify and estimate a structural econometric model of vehicle choice where a consumer chooses the vehicle that maximizes their indirect utility, taking into account both product and consumer characteristics. In particular, we rely on a random coefficients logit model that accounts for consumer heterogeneity. We allow for heterogeneity at the individual level via demographics, heterogeneity in both prices and different components of operating costs through random coefficients, and also interact product and consumer characteristics. This is important because it will allow us to depart from the representative consumer framework frequently used in economic theory and most of the empirics in the related literature. After all, if automobiles are highly differentiated products, this is partly due to 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 coexist and thus can be separately examined. The first component is a fuel tax that 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, the efficacy of the fuel tax as a policy instrument crucially relies on its correct valuation by consumers.

The second component of interest is a vehicle tax, a policy instrument often compared to—and combined with—the fuel tax, especially in European countries. The vehicle tax (on emissions, fuel economy) is the instrument of choice for a number of policy makers 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 as compared to a fuel tax.

Different sources of variation in the data help identify the parameters of interest. In the case of the vehicle tax parameter, which depends on CO_2 emissions (which maps onto fuel economy, conditional on fuel), identification comes from variation across fuel types—since vehicles operating on different fuels are taxed differently—across variants of a model, conditional on fuel and vintage, and across vintages of a variant, conditional on variant and fuel. In the case of the fuel cost parameter, identification additionally comes from differences in utilization across fuels and vintages and from the interaction of fuel prices—which vary across vehicle types and over time—and fuel economy.

1.2. Main Findings

Under the baseline assumptions that consumers' fuel price forecasts follow a random walk and that their real discount rate is 5% per year (both of which are prevalent in the literature), we find the average valuation of total operating costs (fuel costs plus vehicle tax) to be 0.51 and reject the null hypothesis of correct valuation at the 1% significance level, which is consistent with its undervaluation by consumers. Once operating costs are decomposed into lifetime fuel costs and vehicle tax components, we document 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% 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.[1]

Given the difference in the valuations of the two policy instruments, we investigate potential 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 US$100 or "round" 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 that are functions of fuel prices and exploits the time series behavior of such variables.

We introduce salient components of vehicle tax by 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, salience of vehicle tax is induced by when it is due and temporary tax exemptions decrease the salience of vehicle tax.

Given the evidence consistent with salience, we examine two potential channels through which fuel costs might become more salient, namely, news and internet search activity. We document a number of cases in which 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

1. We also estimate versions of our main specification accounting for endogeneity of vehicle prices and fuel costs, obtaining largely robust results and finding no evidence of endogeneity of those variables; see the appendix (apps. A–F are available online).

of behavioral explanations in influencing consumer valuations of policy instruments and in closing the energy efficiency gap.

1.3. Contribution and Related Literature

To the best of our knowledge, we are the first to take advantage of a unique institutional setting that allows the joint evaluation of two key policy instruments, namely, the fuel tax and the vehicle tax. 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 et al. 2007). However, we document that the undervaluation of other instruments 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 data set focusing on retail consumers. 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 vehicle market. With the above in mind, we formulate a structural econometric model that we take to data.

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 relative prices of used cars fully adjust to changes in the relative net present value of fuel prices. 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 et al. (2016). Consistent with the recent literature looking at the transport sector, these papers use the time series variation in fuel prices interacted with fuel economy to estimate the valuation of energy costs.[2] 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) and Leard et al. (2017) document modest and moderate undervaluation of energy costs, respectively (their preferred estimates imply that consumers pay \$0.76 and \$0.54 for every \$1 of discounted fuel cost savings, respectively). Sallee et al. (2016) provide 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

2. For instance, Li et al. (2009) pursue a similar strategy to study vehicle scrappage whereas Klier and Linn (2010) study the effect on new vehicle sales.

that vehicles with different expected remaining lifetimes are affected differently by shocks to fuel prices when it comes to their lifetime fuel costs.

Methodologically, our paper more closely aligns to the structural branch of this literature, which goes back to Hausman (1979) and Dubin and McFadden (1984). 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 Verboven (2002), who uses aggregate product-level data to quantify the valuation of energy costs and the pricing behavior of car makers, finding mild undervaluation by consumers. More recently, Grigolon et al. (2018) quantify the valuation of energy costs using market-level data for a panel of European markets. Since not accounting for heterogeneity in willingness to pay (WTP) for fuel costs biases the valuation of energy costs (Bento et al. 2012), Grigolon et al. (2018) carefully control for consumer heterogeneity, especially in regard to mileage, finding that consumers modestly undervalue energy costs. Our contribution, as detailed below, is to bring to the literature aspects of salience that affect the valuation of policy instruments, something enabled by the detailed data set at the micro level we construct, and to show that there are some cases in which increasing the salience of a policy instrument will lead to its correct valuation.

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 (e.g., Conlisk 1996) of a vehicle. Since fuel costs are substantially higher than vehicle tax, they are more salient (Gerlagh et al. 2016) and thus receive more attention from 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.[3] For instance, Chetty et al. (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 in the store have larger effects on demand. Finkelstein (2009) finds that policy makers take advantage of the reduced salience induced by electronic toll collection systems. More recently, Li et al. (2014) find evidence that consumers respond more strongly to changes in

3. 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.

(more salient) fuel taxes than (less salient) net fuel prices. Moreover, 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. Recent evidence is mixed; while Allcott and Knittel (2019) find no effect of providing additional information and/or making consumers more attentive to fuel economy on the fuel economy of vehicles purchased after the intervention, Leard (2018) finds heterogeneity in willingness to pay for fuel cost savings between attentive and inattentive consumers.

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 et al. (2017) for a recent survey. While this literature typically offers a number of potential explanations for the gap, our results identify and provide quantitative support for the importance of behavioral explanations, of which salience and consumer inattention are particular cases, in closing the energy gap.

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 which they impose 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 overutilization 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 policy makers aim to address is CO_2 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. Thus, it can achieve the first-best outcome 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.[4]

4. For other externalities, different instruments would be first-best. For instance, road pricing would address the congestion externality, and a per-mile tax based on vehicle type might be first-best to address the accident externality. In the case of local pollutants, the design of policies becomes more complex (Fullerton and West 2002, 2010; Knittel and Sandler 2012).

Standards can equivalently be imposed on CO_2 emissions, fuel consumption, or fuel economy. 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.

The inefficiency of standards occurs due to the rebound effect.[5] That is, standards incentivize the purchase of more efficient cars, but the resulting savings generate an increase in vehicle utilization. Intuitively, this can be shown as follows. Assume that the only transport externality is CO_2 emissions, which are denoted by $\mathcal{E}_{ij}$ for driver i and vehicle j, and are measured in grams of CO_2 (gCO_2.) A fuel tax incides on both vehicle utilization θ_{ij} (measured in VKT, vehicle-kilometers-year, say) and emissions e_j whereas a standard 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.[6]

A vehicle tax is another important policy instrument to tackle emissions-related externalities. As happens with standards, the vehicle tax incides only on vehicle emissions, and the marginal cost of driving decreases in fuel economy. However, while standards incide on new vehicles, vehicle taxes incide on all circulating fleet and are due on a yearly basis. As pointed out at least since Verboven (2002),[7] vehicle taxes play an important role when it comes to technology adoption (gasoline vs. diesel vehicles) in the European car market.[8]

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.[9] The Swedish fuel tax comprises an energy tax and a carbon tax, which combined make one of the highest fuel taxes (thus fuel prices) worldwide. For reference, at US$5.93 per gallon in 2016:Q2, gasoline prices

5. Additional externalities, such as congestion, will favor the fuel tax.

6. However, the empirical evidence in the economics literature since at least Goldberg (1998) is that consumers respond little to increasing fuel prices (Small and Van Dender 2007; Hughes et al. 2008; Li et al. 2009).

7. More recently, Klier and Linn (2015) have also shown the role of vehicle taxes in the reduction of CO_2 emissions—to a great extent via an increase in the market share commanded by diesel vehicles—in three European markets.

8. Standards apply to new vehicles and might or might not be passed on to consumers by car makers. On the other hand, the vehicle tax is charged (typically yearly) on both new and used vehicles. Letting τ_s be the implicit tax rate, e a vehicle's CO_2 emissions, and $\varphi > 0$ the binding standard, the tax imposed is $\tau_s(e - \varphi)$, paid once. The same formula holds for an explicit vehicle tax τ_{vt}, but paid yearly—see app. A for details and formulas. To the extent that the explicit and recurrent 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.

9. The EU CO_2 emission standards introduced in 2009 affect only a minority of the used vehicles in our sample.

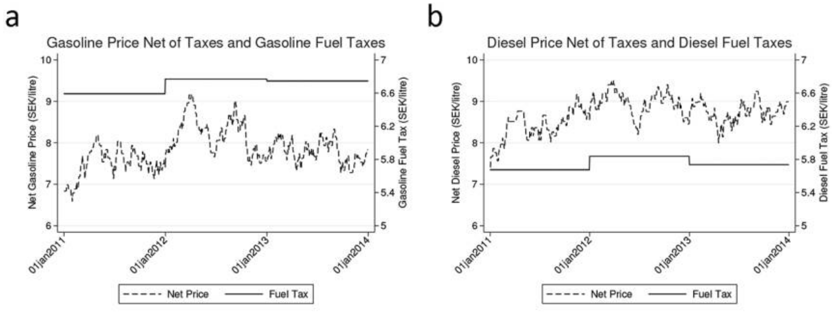

Figure 1. Fuel price components. This figure displays the time series of net gasoline (*a*) and diesel (*b*) prices (exclusive of taxes) together with the time series of the corresponding fuel taxes.

in Sweden are the ninth highest worldwide. In contrast, at US$2.57 per gallon, gasoline prices are substantially lower in the United States.[10]

Any changes to the fuel tax typically happen yearly, with the gasoline tax historically being higher than the diesel tax.[11] 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.

The vehicle tax is calculated based on the CO_2 emissions and the fuel type of a vehicle since the 2006 Swedish vehicle tax reform. Vehicle taxes on a diesel vehicle are approximately 2.5 times the value of a tax on a gasoline vehicle with identical CO_2 emissions (see app. A for details and formulas).[12] Crucially, and to our advantage, since the vehicle tax is an explicit tax effectively charged in monetary units (in contrast to a standard), we are able to value it in the same way we value the fuel tax.

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

10. See, e.g., the Bloomberg fuel price comparison engine: http://www.bloomberg.com/graphics/gas-prices/#20162:Sweden:USD:g.

11. 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 (Miravete et al. 2015).

12. Vehicle tax is paid according to the Road Traffic Tax Act (2006, 227), which came into force on May 1, 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, according to which the vehicle tax was charged based on a vehicle's weight. Exemptions are detailed in chap. 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.

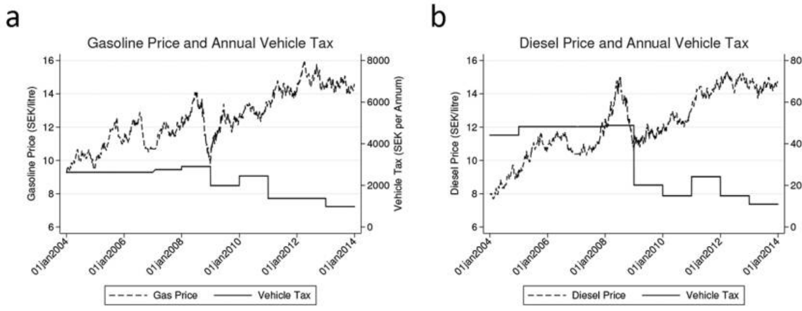

Figure 2. Fuel prices and vehicle taxes. This figure displays the time series of gasoline and diesel prices together with the time series of the annual vehicle tax of a baseline model Volvo V70 running on gasoline (*a*) and diesel (*b*). Source: New Car Guide.

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.[13]

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 2 for an illustration using baseline gasoline and diesel versions of the best-selling Volvo V70. What is more, despite substantial decreases in CO_2 emissions over time, they were not enough to counteract the penalties imposed on diesel vehicles.

Empirically, we will use information on fuel costs to draw implications for the fuel tax. 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.

3. DATA

We combine a number of data sets. Our unique transaction data from a major Scandinavian online auction platform 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 is less prone to concerns regarding the use of wholesale data to make inferences about consumer behavior (Busse, Lacetera, et al. 2013) and allows us to control for consumer heterogeneity at the micro level.[14]

13. Miravete et al. (2015) document that the share commanded by diesels in Europe increased from as low as 10% in 1990 to up to 70% by 2000 in some segments. Klier and Linn (2015) argue that vehicle taxes and demand for fuel economy combined explain diesel adoption in the continent. See also http://www.eea.europa.eu/data-and-maps/figures/dieselisation-in-the-eea.

14. 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

Our car inspection data enable 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 provide more reliable information in this important dimension. 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.

3.1. Transaction Data

The transaction data used in this study are auction data from a large Swedish platform collected from October 12, 2011, to June 13, 2013.[15] 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.[16] The history of bids is publicly available, although only the five latest bids are shown, together with the user name of the

---

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 resell 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 beg for the use of retail prices.

15. This platform is the leading player in Scandinavia, where online penetration and online shopping are among the highest worldwide. 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. While anecdotal, there are no reasons to believe that selection issues facing our sample are any more severe than those widely used in the recent literature, such as the dealer auction data provided by Manheim Consulting for the US market, for instance.

16. This means that should a bidder decide to back out after having won the auction, she or 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.

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.[17] 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, and size as well as the inspection period and tax details for each vehicle.[18] 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, brakes, gearbox, and interior. A given vehicle is available at one of 15 locations in Sweden and anyone interested may go and test-drive the given vehicle on prespecified 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 prespecified 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.

Given this information, the choice set for a given individual comprises the vehicle actually purchased by the individual in addition to all objects available on the market in weeks when the consumer was active, that is, placed at least one bid. (We include vehicles into the choice set by weeks because the auction house lists vehicles on a weekly basis. We do not observe whether an individual consumer was logged onto the platform.)

3.2. Car Inspection Data

We use car inspection data from 2007 to 2010, which include failure rates and average odometer readings disaggregated by model, age, fuel type, and year of publication. Vehicles are inspected at the ages of 3 and 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.

3.3. 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.

17. 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.

18. The Swedish Transport Agency also provides details about the month of payment for vehicle taxes. See the appendix for details.

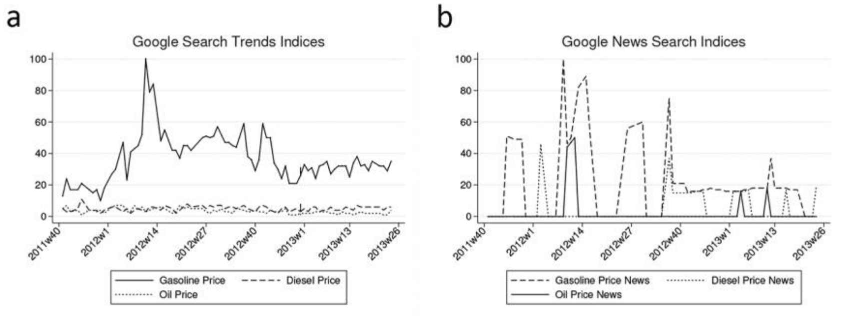

Figure 3. Search and news indices. This figure displays Google Search Trend Indices (*a*) and Google News Search Indices (*b*) at the weekly frequency for the sample period, October 2011 to June 2013.

3.4. Consumer Characteristics

We use median per capita income collected by the Swedish Statistics Bureau (SCB) for the municipalities where a transaction takes place as a proxy for consumer income.

3.5. 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 its fuel stations. We collect information on aggregate fuel volumes delivered within Sweden from the Swedish Petroleum and Bio-fuels Institute.

3.6. 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 3 displays their time series behavior.

3.7. Combining Data Sets

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 fuel and car age, and we use the car survival data to estimate and predict average car hazard rates as a function of total VKT traveled.[19] Table 1 displays an extract of a choice set. Table 2 reports summary

19. We consider only vehicles fueled by diesel and gasoline because for flexible-fuel vehicles, the leading alternative fuel vehicle then, fuel choice becomes another choice dimension that can significantly complicate the analysis (Salvo and Huse 2013; Huse and Lucinda 2014) and the

Table 1. Extract of Sample Choice Set

Brand	Model	Fuel	Price	Fuel Cost	Taxes	Weight	Age	Odometer	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 auctions where a randomly chosen user ("-RS6-") placed at least one bid. A full choice set comprises all vehicles auctioned in the weeks a given consumer was active. Prices, fuel cost, and taxes are in SEK, weight in kg, age in years, odometers in km; DV(.) denotes a dummy variable.

statistics on bidders, vehicle characteristics, and driving characteristics for the final sample of 11,262, unique registration numbers.

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.

4.1. 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 can choose one (and only one) among the J products on their choice set, namely, the one maximizing their indirect utility. The econometrician does observe individual choices, product characteristics, and consumer characteristics.[20] The indirect utility of consumer i $(i = 1, \dots, N)$ from purchasing product j $(j = 1, \dots, J)$ at period t is given by

$$\begin{aligned} U_{ijt} = {} & \phi_j^{\text{model}} + \phi_j^{\text{fuel}} + \phi_j^{\text{transm.}} + \phi_j^{\text{awd}} + \phi_j^{\text{turbo}} + \phi_t + x_j'\alpha + y_{it}'\gamma \\ & + \chi_{ijt}'\beta_i + \varepsilon_{ijt}, \end{aligned} \tag{1}$$

expectations of future fuel prices. Hybrid electric vehicles (e.g., the Toyota Prius) and Gasoline-CNG hybrids command negligible market shares and are also discarded.

20. 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.

Table 2. Descriptive Statistics

	Mean	SD	Min	Q1	Med	Q3	Max
	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
	B. Vehicle Characteristics						
DV (diesel)	.63	.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 (years)	4.65	2.71	.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	.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. NPV = net present value.

where ϕ_j^k, $k =$ model, fuel, transmission, all-wheeldrive(awd), turbo denote the fixed effect corresponding to characteristic k (model, fuel, transmission, all-wheel drive, turbo engine) for product j, and ϕ_t are time fixed effects. The vector x_j comprises characteristics such as vehicle ratings (for engine, body work, gearbox, interior, brakes), vehicle age, odometer reading, engine size, and total vehicle weight.[21] The vector y'_{it} comprises consumer characteristics such as median income within a region; the vector $\chi'_{ij} = (c_{ij}, \tau_j, p_j)$ comprises the monetary variables, namely, the discounted lifetime fuel costs, discounted lifetime vehicle taxes, and product price, respectively.[22] Finally, ε_{ij} is a stochastic term with a type 1 extreme value (T1EV) distribution.[23]

4.2. Heterogeneity

Given the prominence of heterogeneity in the vehicle market, we allow for different channels through which heterogeneity can affect choices.[24] On the product front, we allow for different fixed effects (e.g., brand-model, fuel), product characteristics (e.g., engine size, odometer readings), 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 across vintages, in addition to the current state of the product.

On the consumer front, we allow heterogeneity to affect choices via consumer characteristics, heterogeneity in utilization (which feed into fuel costs; details below), and consumer-type random coefficients. The consumer characteristic we use is the median income within a region, which measures the purchasing power of a consumer. Utilization 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 utilization typically decreases with vehicle age and diesel vehicles consistently are driven more than their gasoline counterparts. Finally, we endow the parameters associated with the monetary variables with random coefficients; see below.

---

21. We include such ratings (attributed to each vehicle by an independent mechanic; see sec. 3) to control for heterogeneity across vehicles of same observed characteristics (e.g., model-fuel-vintage-engine size) but different utilization.

22. The subscripts i in $\chi'_{ijt} = (c_{ijt}, \tau_{jt}, p_j)$ are due to the interaction between consumer (utilization, measured in vehicle-kilometers traveled) and product characteristics (details below). Thus, our specification comprises product characteristics, consumer characteristics, and interactions thereof.

23. We do not consider the existence of an outside option in our analysis, since our focus is on how the variation in the data influences which product to purchase, not on the timing of the decision to purchase. This is consistent with the purchase of used vehicles with replacement motives and has been used in similar models for other durables, e.g., Houde (2018) for refrigerators.

24. As shown in Bento et al. (2012), one needs to account for heterogeneity in the valuation of fuel costs at the risk of biasing the valuation of energy efficiency downward.

The model is estimated using maximum simulated likelihood (see app. B for details). In our application, we decompose the parameters associated with the monetary variables as $\beta_{ik} = \exp(\beta_k^* + \sigma_k \nu_{ik})$, where β_k^* is common across individuals, ν_{ik} is a consumer-specific random determinant of the taste for characteristic k, which we assume to follow independent draws of the standard Normal distribution, $\nu_{ik} \sim \mathcal{N}(0, 1)$, and σ_k measures the impact of ν on characteristic k, thus making the individual coefficients lognormally distributed.[25]

4.3. Details and Discussion

The discounted lifetime operating costs of vehicle services comprise fuel costs (c_{ij}) and taxes (τ_j), which are given by

$$c_{ij} = \sum_{t=0}^{T} \delta^t s_{jt} \pi_{jt} \phi_j \theta_{ijt}, \tag{2}$$

$$\tau_j = \sum_{t=0}^{T} \delta^t s_{jt} \tau_{jt}, \tag{3}$$

with δ being the discount factor, T the lifetime of a vehicle, and s_{jt} a vehicle's survival probability (which we estimate from scrappage data; see app. A). In our application, we follow the literature by setting vehicle lifetime to 25 years (Allcott and Wozny 2014) and the real discount rate to 5% per year (see app. A for details and app. C for robustness).[26] In the case of the fuel costs, $\pi_{jt}\phi_j\theta_{ijt}$ combines fuel price (π_{jt}, in SEK/liter), vehicle fuel economy (ϕ_j, in liters/100 km), and vehicle-kilometers/year (VKT, in θ_{ijt} km/year) whereas for the tax term, τ_{jt} is the annual vehicle tax. We assume consumers'

25. 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 is well defined and lognormally distributed. See the appendix for details.

26. We assume that the discount factor $\delta = 1/(1 + r)$ is constant across consumers. We set the discount rate r to 5% 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 US Environmental Protection Agency (2014). Nominal interest rates for car loans 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%, 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% reported 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%. Our strategy is also consistent with the literature, e.g., Verboven (2002) and Busse, Knittel, and Zettelmeyer (2013) use borrowing rates, Allcott and Wozny (2014) use the average of borrowing costs and stock market returns (calculated to be 6%), and Sallee et al. (2016) assume a discount rate of 5%.

expectations of real future prices to follow a random walk (or "no change" forecast), consistent with most of the recent literature (Busse, Knittel, and Zettelmeyer 2013; Sallee et al. 2016; Grigolon et al. 2018). This assumption is consistent with empirical evidence of consumer beliefs (Anderson et al. 2013) and has been shown to perform at least as well as alternative forecasts based on futures markets or expert surveys (Alquist and Kilian 2010; Alquist et al. 2013) used in, for example, Allcott and Wozny (2014) preferred results.

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 et al. (2018) for a recent case. This is consistent with empirical findings of a small and statistically insignificant rebound effect typically found in the literature.[27] We believe that this is 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.

While flexible, 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. In particular, there is an interdependence between the new and used car markets, something we abstract from by not including an outside good. However, as documented in Schiraldi (2011), the substitution between those segments is small. Our estimation approach thus represents a pragmatic modeling approximation to actual choice behavior in the industry, being consistent with much of the literature.

4.4. Identification

Our estimation strategy uses data on all the products a consumer could have bid for while active on the market to identify preference parameters instead of using information on auction outcomes only, for example, price and characteristics of the product sold. This is only possible by following the activity history of consumers in the sample over time.

The parameters of interest, β_i, are consistently estimated provided $E[\chi_{ij}\varepsilon_{ij}|I_{ij}] = 0$, where I_{ij} subsumes product and consumer characteristics. 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

27. For perspective, Small and Van Dender (2007) estimate a short-run rebound effect of 0.044 using yearly date whereas Klier and Linn (2010) estimate an elasticity close to 0.12 using monthly data. More recently, West et al. (2017) find no evidence of a rebound effect using micro data and accounting for attribute-based adjustment.

characteristics unobserved by the econometrician and related to, say, product reputation that may be correlated with the cost components.[28] 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, especially previous utilization and the current state of a product. If endogeneity is still a concern despite the above, we note that the resulting attenuation bias will render an overstated valuation parameter υ, as also seen in the appendix focusing on endogeneity, thus making it less likely to reject the null of correct valuation of energy efficiency.

The variation identifying the preference parameters differs between vehicle taxes and fuel costs. First, note that the quantities of interest are lifetime (present discounted values of) fuel costs and vehicle taxes, rather than the yearly quantities themselves, as per equations (2)–(3). Focusing on the yearly vehicle tax, the first source of variation is the difference in taxation between gasoline and diesel vehicles (see app. sec. A.2 for formulas). Moreover, conditional on fuel, vehicle taxes differ both cross-sectionally, due to the difference in CO_2 emissions of different variants of a model at a given point in time, and over time, due differences in CO_2 emissions of different vintages of the same variant; see figure 2*b* for an example.[29] As for fuel costs, variation across fuels and CO_2 emissions are interacted with time series variation in expected fuel prices. We illustrate the identifying variation in the data in figure 4. Figure 4*a* and figure 4*b* show how the present discounted value of vehicle tax and fuel costs, respectively, vary as a function of vehicle age and CO_2 emissions (and across fuels). Finally, figure 4*c* illustrates the covariation between the present discounted value of vehicle tax and fuel costs as a function of CO_2 emissions, fuel, and vehicle age.[30]

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

29. CO_2 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 fuels, and also across variants of a model 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, variants of the gasoline and diesel 2009 Audi A4 have CO_2 emissions in the range 149–89 gCO_2/km and 119–236 gCO_2/km, respectively, whereas the figures for their 2007 counterparts are 171–226 gCO_2/km and 140–200 gCO_2/km, respectively. For the full sample, the mean (median) ratio of maximum to minimum CO_2 emissions at the brand-model-fuel level is 1.41 (1.32).

30. We have also regressed the present discounted value of vehicle tax on (i) model fixed effects, (ii) model plus fuel fixed effects, (iii) model-fuel fixed effects, and (iv) model-fuel fixed effects plus the present discounted value of fuel costs. The R-squared of the last specification was a moderate 0.57, suggesting variation in the lifetime vehicle tax above and beyond the one of lifetime fuel costs.

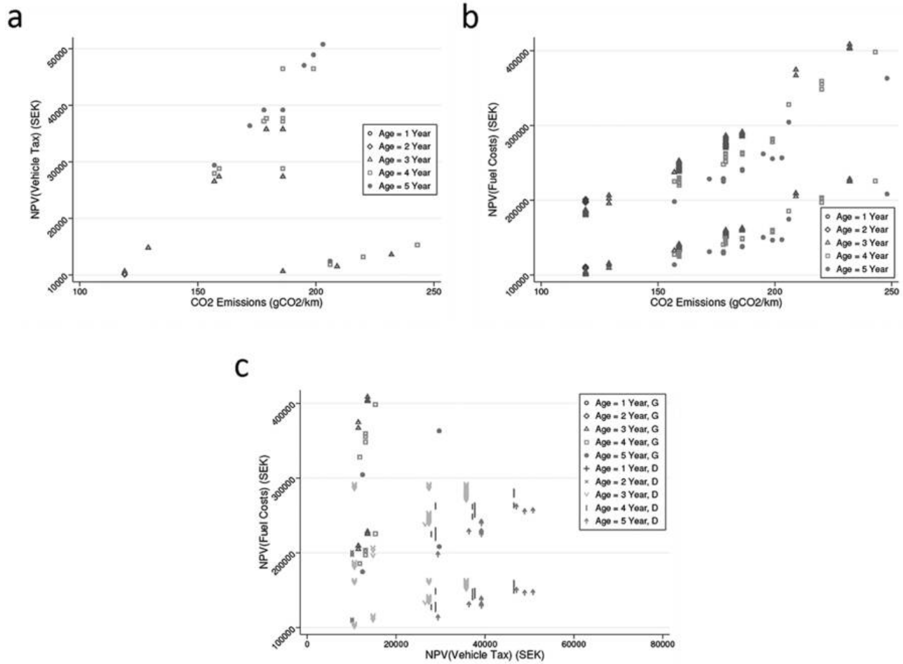

Figure 4. Identifying variation in cost components. This figure displays, for the model Volvo V70, the present discounted value of vehicle tax based on different CO_2 emissions and ages (*a*), the present discounted value of fuel costs based on different CO_2 emissions and ages (*b*), and the covariation between the present discounted value of fuel costs and the present discounted value of vehicle tax of different CO_2 emissions and ages (*c*). All data are from Q2:2013. NPV = net present value. Color version available as an online enhancement.

4.5. Testable Implications

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, $v_c := (\partial U_{ij}/\partial c_j)/(\partial U_{ij}/\partial p_j)$ and $v_\tau := (\partial U_{ij}/\partial \tau_j)/(\partial U_{ij}/\partial p_j)$.

Under the null hypothesis of full information and rationality, consumers correctly—and equally—value lifetime fuel and vehicle tax costs. Letting subscripts p, c, and τ denote the coefficients for product prices, fuel costs, and vehicle taxes, respectively, and assuming a logit model to fix ideas, the first testable implication—regarding the correct valuation of lifetime fuel costs—reads $v_c := \beta_c/\beta_p = 1$ whereas the second testable implication—regarding the correct valuation of lifetime vehicle taxes—reads $v_\tau := \beta_\tau/\beta_p = 1$. Finally, the correct valuation of both fuel costs and vehicle taxes can be expressed as $v_c = v_\tau = 1$.

5. VALUATION OF OPERATING COSTS

We estimate demand specifications comprising four components, namely, product characteristics, consumer characteristics, interaction terms between consumer and product characteristics, and the specification of consumer heterogeneity. Among product characteristics we consider price, engine size, vehicle age, odometer reading, and weight; we also include fixed effects for model, fuel type, transmission type, all-wheel drive, turbocharged engines, in addition to time (year-month) fixed effects. Finally, we include product ratings.

Next we include the median income at the municipality where a purchase occurred, which helps us control for consumer characteristics. We also include the cost components, a vehicle's (expected) lifetime fuel costs, and lifetime vehicle taxes, with the former containing an interaction between consumer and product characteristics.[31]

Finally, we endow price and cost components with random coefficients. This allows different consumers to have different sensitivities (thus preferences) to different characteristics, being a standard way to account for unobserved consumer heterogeneity in logit models. Standard errors are robust and clustered at the consumer (bidder) level.

5.1. Demand and Valuation Estimates

We report the main estimates of our specifications in table 3 (p-values are reported in parentheses). Specifications 1–3 are conditional logit models whereas specifications 4–6 are random coefficients logit (RCL) models; random draws are lognormally distributed,

31. The cost variables enter as lifetime rather than one-period quantities; therefore the acronym NPV in the tables, denoting net present value.

Table 3. Demand and Valuation Estimates

Variables	(1) Mean	(2) Mean	(3) Mean	(4) Mean	(4) SD	(5) Mean	(5) SD	(6) Mean	(6) SD
Price	−6.46e-06***	−6.23e-06***	−6.25e-06***	−11.95***	.01	−11.98***	.01	−11.98***	.01
	(.00)	(.00)	(.00)	(.00)	(.86)	(.00)	(.86)	(.00)	(.84)
NPV (fuel + tax)	−3.26e-06***			−12.63***	.02				
	(.00)			(.00)	(.28)				
NPV (fuel costs)		−3.90e-06***	−3.86e-06***			−12.45***	.01	−12.49***	.01
		(.00)	(.00)			(.00)	(.24)	(.00)	(.27)
NPV (vehicle tax)			−2.82e-07					−15.47***	1.74*
			(.93)					(.00)	(.07)
Average valuations:									
υ_{tc} (fuel + tax)	.51***				.51***				
υ_c (fuel costs)		.63**	.62*				.63**		.60**
υ_τ (vehicle tax)			.06*						.14***
υ_τ/υ_c			.08						.23**
Observations	397,943	397,943	397,943	397,943		397,943		397,943	
Log-likelihood	−10,859.6	−10,859.2	−10,859.2	−10,859.6		−10,859.2		−10,858.9	

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. *p*-values are reported in parentheses. Standard errors are robust and clustered at the consumer level. 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. The parameters υ_{tc}, υ_c, and υ_τ denote the valuation of lifetime operating costs, fuel costs, and vehicle tax, respectively. The null hypothesis for all valuation parameters is $H_0 : \upsilon = 1$. Mean and SD parameters are denoted by β_k^* and σ_k, respectively.

* Statistically significant at the 10% level.

** Statistically significant at the 5% level.

*** Statistically significant at the 1% level.

and both price and cost components are endowed with a random coefficient.[32] Our baseline specification is 6.

Specifications 1–3 differ with respect to the cost terms. While specification 1 aggregates both cost components into total operating costs, specification 2 considers only lifetime fuel costs (thus purposely omitting the vehicle tax term) whereas 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% significance level and fairly similar. As expected, the cost component in specification 2 loads higher than in specification 1. Finally, in contrast with a significant (at the 1% level) parameter for lifetime fuel costs in specification 3, the lifetime vehicle tax parameter is not significant. The associated valuation parameters are 0.51 for specification 1 and 0.63 for specification 2 (null of correct valuation rejected at the 1% and 5% significance levels, respectively), and 0.62 for fuel costs and 0.06 for vehicle tax in specification 3 (null of correct valuation rejected at the 10% significance level for both).

To better account for heterogeneity, we also estimate a set of RCL models. Specifications 4 and 5 are the RCL counterparts of specifications 1 and 2, respectively. Since the heterogeneity (σ, standard deviation) parameters in these specifications are small and not statistically significant, the valuation parameters do not change when compared to their logit counterparts and the associated log-likelihoods barely increase. As before, the null of correct valuation of operating costs is comfortably rejected for both 4 and 5.

We now proceed as for 3 and decompose total lifetime operating costs into fuel costs and vehicle tax components, but allowing for heterogeneity in the form of random coefficients; see specification 6. Now, heterogeneity becomes statistically significant (at the 10% level) for lifetime vehicle tax. That is, there is evidence of consumer heterogeneity in the responses to lifetime vehicle taxes. This can be attributed to the fact that newer (less than 5 years) and older vehicles are treated differently, with the former benefiting from a tax exemption (details below).

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 lifetime vehicle 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% 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

32. Although in our baseline the draws are uncorrelated, we have also experimented with more general specifications of the matrix of random coefficients. Our choice was driven by (i) convergence issues and (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., nonzero 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.

the 5% significance level (see v_T/v_c in table 3). That is, on top of undervaluing both components, consumers' undervaluation of lifetime vehicle tax is more severe than that of lifetime fuel costs. Moreover, while consumer reaction is relatively homogeneous in the case of lifetime fuel costs, it is heterogeneous in the case of lifetime vehicle tax.

The role of heterogeneity in valuations can be illustrated by examining the distribution of valuations for the baseline; see figure 5. Valuations are substantially more heterogeneous for the vehicle tax (fig. 5*b*) than for fuel costs (fig. 5*a*).

The above findings have important policy implications. First, we document that consumers 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 policy makers and researchers—is also undervalued by consumers. Moreover, 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.

5.2. Potential Endogeneity and Robustness Checks

One concern with the main results presented in table 3 is that variables such as vehicle prices or lifetime fuel costs are endogenous, potentially leading to biased and inconsistent

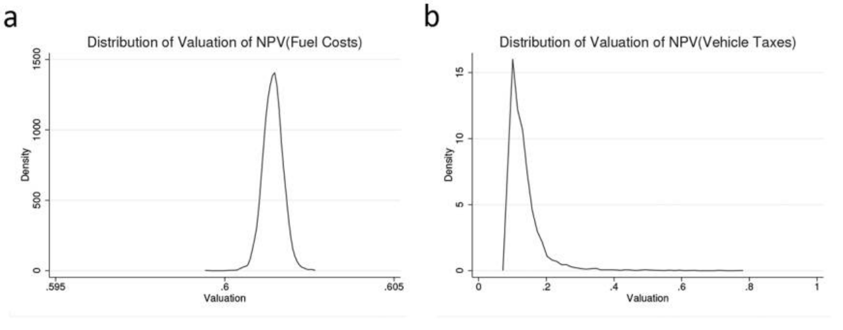

Figure 5. Distribution of valuations of lifetime fuel costs and vehicle tax. This figure displays the distribution of valuations of lifetime fuel costs (*a*) and vehicle tax (*b*) as per the baseline specification, specification 6 in table 3. The underlying model is a random coefficients logit (RCL) with lognormally distributed parameters on price and 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.

estimates. To address these concerns, we report in the appendix a two-stage control function approach as in Petrin and Train (2010). We interpret the resulting estimates as supporting the view that our main results are not due to the endogeneity of car prices or lifetime fuel costs.

We also conduct a number of robustness checks. As detailed in the appendix, the estimates of undervaluation are largely robust, with variation similar to the ones in the literature.

6. POTENTIAL EXPLANATIONS

We have documented that (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. Following the literature, for example, 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.[33] However, in our case vehicle taxes are more persistent than fuel prices yet more severely undervalued by consumers; thus undervaluation is unlikely to be driven by persistence, see appendix D for details.

To better appreciate how undervaluation relates to salience, consider the model of inattentive agents put forth in Conlisk (1996), according to which computing a payoff function is costly for decision makers in the sense of requiring deliberation due to its complexity. In such case, the decision maker applies a costly deliberation technology so that he maximizes the expected payoff net of the cost of deliberating. Translated to our setting, the correct computation of fuel costs is both more costly (due to its complexity) and more rewarding to consumers on the market for a used vehicle (table 2 reports that the ratio of mean lifetime fuel costs over mean lifetime vehicle tax is about seven). Since it is unlikely 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. 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 what follows, we empirically examine the role of salience on valuations. We do so by, first, splitting a cost component into salient and nonsalient components and then comparing their valuations by reestimating our baseline specification. In the case of vehicle tax, we exploit the within-year variation induced by different due dates of the

33. 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.

vehicle tax to show how increasing the salience of vehicle tax increases its valuation by consumers. For fuel costs, we define salient and nonsalient versions of lifetime fuel costs where the salience is switched on whenever a proxy variable is activated. For instance, a salient component of fuel costs consists of the actual fuel costs interacted with an indicator function of oil prices being at least US$100. The testable implications are twofold. First, the valuation of the salient component of cost exceeds that of its nonsalient counterpart. Second, 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.

We also employ salience proxies related to news activity and online search activity. While these mechanisms are likely different from the consumers's standpoint, they should point in the same direction: news activity variables (e.g., an indicator for "news about gasoline/gasoline prices is above its long-run average") can be seen as more passive 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. On the other hand, online search activity (e.g., an indicator for "search for keyword gasoline/gasoline prices is above its long-run average") suggest that consumers (perhaps after being exposed to news) take an active role to look for (additional) information on fuel prices.

6.1. Salience in Vehicle Tax

We rely on the institutional background and examine how consumers react to lifetime vehicle taxes as they vary in salience. 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. Concretely, we interact the indicator of nonexempt vehicle tax status with indicators of tax being due in the next quarter and in the next month, respectively.

Table 4 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, with all estimates close to the baseline estimate of 0.60. The estimates also share with the baseline the fact that the null of correct valuation is rejected at the 5% significance level. Perhaps unsurprisingly, the nonsalient valuation estimates of lifetime vehicle tax are low, with the null of their correct valuation being rejected at the 1% significance level. Specification 1 distinguishes between newer and older cars, that is the salient component switches on for vehicles older than 5 years. As expected, the valuation of the salient component is higher than the valuation of the nonsalient one (0.05 vs. 0.01), but the null of equality between them cannot be rejected against a two-sided alternative for standard significance levels.

Table 4. Salience of Vehicle Tax

Variables	(1)[a] Mean	(1)[a] SD	(2)[a] Mean	(2)[a] SD	(3) Mean	(3) SD
Price	−11.98***	.01	−11.95***	.01	−11.96***	.01
	(.00)	(.68)	(.00)	(.59)	(.00)	(.85)
NPV (fuel costs)	−12.46***	.01	−12.47***	.01	−12.49***	.01
	(.00)	(.65)	(.00)	(.69)	(.00)	(.28)
NPV (vehicle tax)	−19.01***	1.79***	−20.73***	3.18*	−16.11***	1.90**
	(.00)	(.00)	(.00)	(.06)	(.00)	(.01)
NPV (vehicle tax, age $\geq$ 5)	−15.14	.65				
	(.23)	(.93)				
NPV (vehicle tax, age $\geq$ 5, due$_{q+1}$)			−12.04***	.31		
			(.00)	(.64)		
NPV (vehicle tax, age $\geq$ 5, due$_{m+1}$)					−11.46***	.12
					(.00)	(.60)
Average valuations:						
υ_c (fuel costs)		.62**		.59**		.59**
υ_τ (vehicle tax, nonsalient)		.01***		.02***		.10***
$\upsilon_\tau^{\text{Salient}}$ (vehicle tax, salient)		.05**		.96		1.67
υ_τ/υ_c		.01***		.03***		.16***
$\upsilon_\tau^{\text{Salient}}/\upsilon_\tau$		13.43		132.00***		20.60***
Observations	397,943		397,943		397,943	
Log-likelihood	−10,859.2		−10,856.7		−10,856.1	

Note. Estimated model is a random coefficients logit (RCL) with lognormally distributed parameters on price and all net present value (NPV) terms. p-values are reported in parentheses. Standard errors are robust and clustered at the consumer level. 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. The parameters υ_c and υ_τ denote the valuation of lifetime fuel costs and vehicle tax, respectively. The parameter $\upsilon_\tau^{\text{Salient}}$ denotes the salient component of lifetime vehicle tax. The null hypothesis for all valuation parameters is $H_0 : \upsilon = 1$.

[a] Specification 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. Mean and SD parameters are denoted by β_k^* and σ_k, respectively.

* Statistically significant at the 10% level.

** Statistically significant at the 5% level.

*** Statistically significant at the 1% level.

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 either case, the null of correct valuation is rejected at the 1% significance level.

Focusing now on the valuation of the salient component of vehicle tax, we cannot reject the null of correct valuation for either specification 2 or specification 3—the point estimates are 0.96 and 1.67, respectively. While the former estimate implies that consumers essentially correctly value the vehicle tax, the latter may be suggestive of hyperbolic discounting (Frederick et al. 2002). Importantly, we reject the null of equality between the standard and salient components of vehicle tax at the 1% 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 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. Perhaps more importantly than the undervaluation of nonsalient 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 policy maker.

6.2. Salience in Fuel Costs

When examining salience in fuel costs, we selected variables based on the literature, with results reported in table 5. In specification 1, the salient component ($\upsilon_c^{\text{Salient}}$) is switched on whenever oil prices reach at least US$100.[34] Specification 2 postulates that the salient component switches on whenever fuel prices are within 10 cents of an integer number, for example 15±0.10 SEK/liter. Specification 3 switches on separate salient components for price increases and decreases for the months with the 5% largest increases and 5% largest decreases with respect to the previous month. 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, Knittel, and Zettelmeyer (2013).

The estimates for specification 1 show that fuel costs are more highly valued if oil prices are salient, with a valuation of 0.85 for $\upsilon_c^{\text{Salient}}$. While the increased valuation of the salient component of fuel costs is perhaps not surprising, it is interesting to note that while nonsalient fuel costs do not have a statistically significant standard deviation parameter (σ), suggesting that consumers respond to such changes in a homogeneous

34. Our choice of price threshold comes from the fact that 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* states that "oil prices [are] 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 November 9, 2007; see http://www.nytimes.com/2007/11/09/business/worldbusiness/09oil.html?_r=2. Similarly, 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 January 31, 2011; see http://www.bbc.com/news/business-12328745.

way just as in our baseline results, the heterogeneity parameter in the salient component of fuel costs is statistically significant, a finding consistent with heterogeneous consumer reactions whenever oil prices cross the US$100 threshold. Importantly, the null of correct valuation for the salient component of fuel costs is no longer rejected. When comparing valuations, while we do not reject the null of equal valuation of the salient and nonsalient components of fuel costs, we do reject the null of equal valuation of the nonsalient component of fuel costs and that of vehicle tax at the 10% significance level.

The estimates for specification 2 suggest a similar yet less extreme pattern than that of specification 1. That is, the nonsalient components of fuel costs and vehicle tax have valuations similar to those of the baseline specification and the null of correct valuation for both components is 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 substantially larger than the latter (0.73 vs. 0.53, 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% significance level. When comparing the salient components with their nonsalient counterpart, the findings are intuitive in that one does reject (at the 10% significance level) the null of equality between the salient price increase, component and the nonsalient component, but not the null of equality between the salient price decrease, component and the nonsalient component. All in all, 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 nonsalient counterpart (0.64 vs. 0.60). However, the null of equality of these two components cannot be rejected. Taken together, we interpret the findings above as suggestive 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 nonsalient component of fuel costs and the valuation of vehicle taxes remains largely unaffected.

6.3. Salience Induced by News and Web Searches

We investigate the role of salience using data from web searches and news related to oil, gasoline, and diesel prices as proxies thereof. Table 6 reports estimates related to news activity with the component being defined as salient if the activity is above its mean value for the sample.[35] The nonsalient components of fuel cost valuations are largely similar to the baseline estimate of 0.60, being in the range 0.54–0.70. On the other hand,

---

35. We have also estimated the same specification defining the component as salient if it is above the 75% percentile activity, obtaining similar results.

Table 5. Salience of Fuel Costs

	(1)[a]		(2)		(3)		(4)	
Variables	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Price	−11.98***	.01	−11.98***	.01	−11.98***	.01	−11.98***	.01
	(.00)	(.43)	(.00)	(.87)	(.00)	(.87)	(.00)	(.92)
NPV (fuel costs)	−12.59***	.10	−12.50***	.02	−12.49***	.03	−12.50***	.08
	(.00)	(.64)	(.00)	(.15)	(.00)	(.11)	(.00)	(.15)
NPV (vehicle tax)	−17.00**	1.82	−15.22**	1.23	−15.15**	1.21	−15.26**	1.23
	(.01)	(.26)	(.01)	(.46)	(.02)	(.46)	(.02)	(.47)
NPV (fuel costs, oil ≥ \$100)	−12.35***	.64***						
	(.00)	(.00)						
NPV (fuel costs, “round” prices)			−12.32***	.01				
			(.00)	(.57)				
NPV (fuel costs, large price inc.)					−12.29***	.02		
					(.00)	(.65)		
NPV (fuel costs, large price dec.)					−12.63***	.03		
					(.00)	(.19)		
NPV (fuel costs, price ≥ 75 percentile)							−12.43***	.13
							(.00)	(.80)

Average valuations:				
υ_c (fuel costs, nonsalient)	.54***	.60**	.61**	.60**
$\upsilon_c^{\text{Salient}}$ (fuel costs, salient)	.85	.71	.73/.53***	.64*
υ_τ (vehicle tax)	.03***	.08***	.09***	.08***
$\upsilon_c^{\text{Salient}}/\upsilon_c$	1.56	1.19***	1.22*/.85	1.07
υ_τ/υ_c	.06***	.14	.12	.13
Observations	397,943	397,943	397,943	397,943
Log-likelihood	−10,856.4	−10,855.4	−10,858.3	−10,858.8

Note. Estimated model is a random coefficients logit (RCL) with lognormally distributed parameters on price and all net present value (NPV) terms. p-values are reported in parentheses. Standard errors are robust and clustered at the consumer level. 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. The parameters υ_c and υ_τ denote the valuation of lifetime fuel costs and vehicle tax, respectively. The parameter $\upsilon_c^{\text{Salient}}$ denotes the salient component of lifetime fuel costs. The null hypothesis for all valuation parameters is $H_0 : \upsilon = 1$. 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 $\upsilon_c^{\text{Salient}}$ and $\upsilon_c^{\text{Salient}}/\upsilon_c$ are for periods of large increases and right-most numbers for periods with large decreases.

[a] Specification 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. Mean and SD parameters are denoted by β_k^* and σ_k, respectively.

* Statistically significant at the 10% level.

** Statistically significant at the 5% level.

*** Statistically significant at the 1% level.

Table 6. Effect of News Activity

	(1)		(2)		(3)	
Variables	Mean	SD	Mean	SD	Mean	SD
Price	−11.98***	.01	−11.96***	.01	−11.98***	.01
	(.00)	(.96)	(.00)	(.95)	(.00)	(.90)
NPV (fuel costs)	−12.61***	.16	−12.32***	.02	−12.49***	.03
	(.00)	(.40)	(.00)	(.56)	(.00)	(.17)
NPV (vehicle tax)	−15.31**	.90	−17.36	1.50	−15.33**	1.26
	(.00)	(.63)	(.24)	(.53)	(.02)	(.49)
NPV (fuel costs, gas news > mean)	−12.28***	.72***				
	(.00)	(.00)				
NPV (fuel costs, diesel news > mean)			−13.00***	1.00***		
			(.00)	(.00)		
NPV (fuel costs, oil news > mean)					−12.45***	.85**
					(.00)	(.03)
Average valuations:						
υ_c (fuel costs, nonsalient)		.54***		.70		.60**
$\upsilon_c^{\text{Salient}}$ (fuel costs, salient)		.96		.58**		.89
υ_τ (vehicle tax)		.05***		.01***		.08***
$\upsilon_c^{\text{Salient}}/\upsilon_c$		1.80		.84		1.49
υ_τ/υ_c		.10		.02***		.13
Observations	397,943		397,943		397,943	
Log-likelihood	−10,848.7		−10,853.8		−10,857.1	

Note. Estimated model is a random coefficients logit (RCL) with lognormally distributed parameters on price and all net present value (NPV) terms. p-values are reported in parentheses. Standard errors are robust and clustered at the consumer level. 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. The parameters υ_c and υ_τ denote the valuation of lifetime fuel costs and vehicle tax, respectively. The parameter $\upsilon_c^{\text{Salient}}$ denotes the salient component of lifetime fuel costs. The null hypothesis for all valuation parameters is $H_0 : \upsilon = 1$. Mean and SD parameters are denoted by β_k^* and σ_k, respectively.

* Statistically significant at the 10% level.

** Statistically significant at the 5% level.

*** Statistically significant at the 1% level.

the valuation estimates of the salient components of fuel costs vary somewhat; specifications 1 and 3, which focus on gasoline and oil news activity, respectively, have estimates of 0.96 and 0.89. In either case, the null of correct valuation is not rejected. On the other hand, specification 2, which focuses on diesel news, has a salient component which is actually lower than its nonsalient counterpart (0.58 vs. 0.70), with the null of correct valuation being rejected at the 1% significance level. 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 (at least in the short run), as documented in the appendix.

We note that the only statistically significant heterogeneity parameters are those associated with the salient components, which suggests substantial heterogeneity in the responses to increased salience, potentially due to the existence of attentive and inattentive consumers. All in all, table 6 documents how news related to gasoline and oil prices is associated with an increased valuation of fuel costs, for which the null of correct valuation cannot be rejected any longer.

Table 7 reports estimates where salience is associated with web search activity; the salient component in specifications 1–3 switches on if the search activity associated with the corresponding keyword is above its sample average. We first note that being in the range 0.60–0.62, the valuation estimates of nonsalient lifetime fuel costs are quite similar to the baseline of 0.60. Moreover, the null hypothesis of correct valuation of this component is rejected at the 5% significance level in all cases. The valuations of lifetime vehicle tax are robust across columns, and the null of correct valuation is rejected at least at the 5% 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 that of the salient component in this specification. The associated valuation estimate is 0.71, with the null of correct valuation not being rejected at standard significance levels. Thus, the valuation of fuel costs not only increases when this component becomes more salient, as measured by an increase in searches, but one cannot reject the null of correct valuation anymore. Moreover, the responses to this increased salience are heterogeneous across consumers.

Specifications 2 and 3 focus on diesel and oil search intensities, respectively. With estimates 0.63 and 0.65, the valuation of the salient component of fuel costs does not increase as much as for specification 1, yet the associated significance level of the rejection of the null of correct valuation is now 10% (compared to 5% for the nonsalient component). While the heterogeneity parameter associated with the salient component is not statistically significant in the case of diesel searches, it is significant at the 5% significance level for oil searches, suggesting heterogeneous reactions to the increase in its salience.

Taken together, the above findings for news and searches suggest a clear pattern whereby making gasoline more salient increases the valuation of fuel costs by consumers to the point that the null of their correct valuation cannot be rejected any longer. The same cannot be said for diesel, which we take as at least being partly driven by its lower short-run price elasticity of demand. The results for oil are somewhat in between, with an effect analogous to that of gasoline for news and a less pronounced one

Table 7. Effect of Search Activity

Variables	(1) Mean	(1) SD	(2)[a] Mean	(2)[a] SD	(3) Mean	(3) SD
Price	−11.97***	.01	−11.97***	.01	−11.98***	.01
	(.00)	(.88)	(.00)	(.68)	(.00)	(.99)
NPV (fuel costs)	−12.48***	.08	−12.56***	.47**	−12.47***	.04
	(.00)	(.39)	(.00)	(.04)	(.00)	(.85)
NPV (vehicle tax)	−15.81***	1.45	−15.45**	1.08	−15.54***	1.35
	(.01)	(.30)	(.02)	(.53)	(.01)	(.33)
NPV (fuel costs, gas search > mean)	−12.53***	.65***				
	(.00)	(.00)				
NPV (fuel costs, diesel search > mean)			−12.45***	.10		
			(.00)	(.25)		
NPV (fuel costs, oil search > mean)					−12.49***	.40**
					(.00)	(.04)
Average valuations:						
υ_c (fuel costs, nonsalient)	.60**		.62**		.61**	
$\upsilon_c^{\text{Salient}}$ (fuel costs, salient)	.71		.63*		.65*	
υ_τ (vehicle tax)	.06***		.08**		.07***	
$\upsilon_c^{\text{Salient}}/\upsilon_c$	1.18		1.02		1.05	
υ_τ/υ_c	.10**		.12		.12*	
Observations	397,943		397,943		397,943	
Log-likelihood	−10,857.3		−10,858.5		−10,858.7	

Note. Estimated model is random coefficients logit (RCL) with lognormally distributed parameters on price and all net present value (NPV) terms. p-values are reported in parentheses. Standard errors are robust and clustered at the consumer level. 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. The parameters υ_c and υ_τ denote the valuation of lifetime fuel costs and vehicle tax, respectively. The parameter $\upsilon_c^{\text{Salient}}$ denotes the salient component of lifetime fuel costs. The null hypothesis for all valuation parameters is $H_0 : \upsilon = 1$.

[a] Specification 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. Mean and SD parameters are denoted by β_k^* and σ_k, respectively.

* Statistically significant at the 10% level.

** Statistically significant at the 5% level.

*** Statistically significant at the 1% level.

for search intensity. In sum, news and search activity are helpful yet imperfect proxies for salience in fuel costs.

7. POLICY IMPLICATIONS

We have documented that consumers undervalue both lifetime fuel costs and lifetime vehicle taxes. Moreover, we have also documented 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 receives less attention (or deliberation) from the part of consumers. In fact, we have documented a number of cases where we cannot reject the null of correct valuation of the policy instruments once they become more salient.

The particular institutional setting we consider allows quantifying the valuation of the fuel tax and the vehicle tax, a standard-like policy instrument. Both policy instruments 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 the vehicle tax is even more undervalued than fuel costs and should be seen as complementary rather than a substitute to the fuel tax, thus departing from the textbook analytical framework typically used to study the efficiency of policy instruments.

Taken together, we see the results as supporting salience—more generally, behavioral effects—as the reason behind (the different levels of) undervaluation documented in the data. Intuitively, economic agents will devote attention to an attribute as long as the associated expected utility exceeds 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 that undervaluation 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, that is, internalize any externalities, policy makers should avoid any tax exemptions since they obfuscate important aspects of the tax system.

Second, allowing the automation of payments and/or making them infrequent should be avoided, due to unintended consequences already documented in Sexton (2015) for electricity bills; while such measures arguably increase the likelihood of taxes being paid, they also minimize any nonmonetary costs associated with them and make them less salient to consumers. Finally, publicizing information about taxes should be encouraged, so it becomes transparent to tax payers which taxes they are due to pay.

8. CONCLUSION

This paper examines the energy efficiency gap using unique revealed preference data from retail consumers in the used car market. We quantify consumer valuation of the fuel tax and of the vehicle tax. We document that consumers undervalue both the 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. We document that behavioral explanations, in particular salience, lie at the root of our findings. In fact, once we introduce a salient component of vehicle tax exploiting information of both tax liability and the timing of tax payment, we cannot reject the null hypothesis 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 the salience of policy instruments helps in closing the energy efficiency gap.

Our findings suggest that a fuel tax is unlikely to provide a first-best allocation. Rather, policy makers should make policy instruments more salient and rely on a menu of them in order to make consumers internalize any externalities they create.

REFERENCES

Allcott, Hunt, and Christopher Knittel. 2019. Are consumers poorly informed about fuel economy? Evidence from two experiments. *American Economic Journal: Economic Policy* 11 (1): 1–37.

Allcott, Hunt, and Nathan Wozny. 2014. Gasoline prices, fuel economy, and the energy paradox. *Review of Economics and Statistics* 96 (5): 779–95.

Alquist, Ron, and Lutz Kilian. 2010. What do we learn from the price of crude oil futures? *Journal of Applied Econometrics* 25 (4): 539–73.

Alquist, Ron, Lutz Kilian, and Robert J. Vigfusson. 2013. Forecasting the price of oil. In *Handbook of economic forecasting*, ed. G. Elliot and A. Timmerman, 427–507. Amsterdam: North-Holland.

Anderson, Soren T., Ryan Kellogg, and James M. Sallee. 2013. What do consumers believe about future gasoline prices? *Journal of Environmental Economics and Management* 66:383–403.

Atkinson, Anthony B., and Joseph E. Stiglitz. 1976. The design of tax structure: Direct versus indirect taxation. *Journal of Public Economics* 6 (1): 55–75.

Bento, Antonio M., Shanjun Li, and Kevin Roth. 2012. Is there an energy paradox in fuel economy? A note on the role of consumer heterogeneity and sorting bias. *Economics Letters* 115:44–48.

Busse, Meghan R., Christopher Knittel, and Florian Zettelmeyer. 2013. Are consumers myopic? Evidence from new and used car purchases. *American Economic Review* 103 (1): 220–56.

Busse, Meghan R., Nicola Lacetera, Devin G. Pope, Jorge Silva-Risso, and Justin Sydnor. 2013. Estimating the effect of salience in wholesale and retail car markets. *American Economic Review: Papers and Proceedings* 103 (3): 575–79.

Chetty, Raj, Adam Looney, and Kory Kroft. 2009. Salience and taxation: Theory and evidence. *American Economic Review* 99 (4): 1145–77.

Conlisk, John. 1996. Why bounded rationality? *Journal of Economic Literature* 34 (2): 669–700.

Dubin, James A., and Daniel L. McFadden. 1984. An econometric analysis of residential electric appliance holdings and consumption. *Econometrica* 52 (2): 345–62.

Eurostat. 2021. Greenhouse gas emission statistics: Air emissions accounts. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Greenhouse_gas_emission_statistics_-_air_emissions_accounts#Analysis_by_economic_activity.

Finkelstein, Amy. 2009. E-Ztax: Tax salience and tax rates. *Quarterly Journal of Economics* 124 (3): 969–1010.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue. 2002. Time discounting and time preference: A critical review. *Journal of Economic Literature* 40 (2): 351–401.

Fullerton, Don, and Sarah E. West. 2002. Can taxes on cars and gasoline mimic an unavailable tax on emissions? *Journal of Environmental Economics and Management* 42 (1): 135–57.

———. 2010. Tax and subsidy combinations for the control of car pollution. *B. E. Journal of Economic Analysis and Policy (Advances)* 10 (1).

Gerarden, Todd, Richard G. Newell, and Robert N. Stavins. 2017. Assessing the energy-efficiency gap. *Journal of Economic Literature* 55 (4): 1486–1525.

Gerlagh, Reyer, Inge van den Bijgaart, Hans Nijl, and Thomas Michielsen. 2016. Fiscal policy and CO_2 emissions of new passenger cars in the EU. *Environmental and Resource Economics* 69 (1): 103–34.

Goldberg, Pinelopi K. 1998. The effects of the corporate average fuel efficiency standards in the US. *Journal of Industrial Economics* 46:1–33.

Grigolon, Laura, Mathias Reynaert, and Frank Verboven. 2018. Consumer valuation of fuel costs and tax policy: Evidence from the European car market, *American Economic Journal: Economic Policy* 10 (3): 193–225.

Harberger, Arnold C. 1964. The measurement of waste. *American Economic Review* 54 (3): 58–76.

Hausman, Jerry. 1979. Individual discount rates and the purchase and utilization of energy-using durables. *Bell Journal of Economics* 10:33–54.

Houde, Sebastien. 2018. How consumers respond to product certification and the value of energy information. *RAND Journal of Economics* 49 (2): 453–77.

Hughes, James, Christopher Knittel, and Daniel Sperling. 2008. Evidence of a shift in the short-run price elasticity of gasoline demand. *Energy Journal* 29 (1): 113–34.

Huse, Cristian, and Claudio Lucinda. 2014. The market impact and the cost of environmental policy: Evidence from the Swedish green car rebate. *Economic Journal* 124: F393–F419.

Jaffe, Adam, and Robert N. Stavins. 1994. The energy paradox and the diffusion of conservation technology. *Resource and Energy Economics* 16:91–122.

Kahn, James A. 1986. Gasoline prices and the used automobile market: A rational expectations asset price approach. *Quarterly Journal of Economics* 101 (2): 323–39.

Klier, Thomas, and Joshua 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.

———. 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.

Knittel, Christopher, and Ryan Sandler. 2012. Carbon prices and automobile greenhouse gas emissions: The extensive and intensive margins. In *The design and implementation of U.S. climate policy*, ed. Don Fullerton and Catherine Wolfram, 287–99. Chicago: University of Chicago Press.

Koptyug, Nikita. 2015. Asymmetric information in auctions: Are resellers better appraisers? IFN Working paper no. 1110, Research Institute of Industrial Economics, Stockholm.

Leard, Benjamin. 2018. Consumer inattention and the demand for vehicle fuel cost savings. *Journal of Choice Modeling* 29:1–16.

Leard, Benjamin, Joshua Linn, and Christy Zhou. 2017. How much do consumers value fuel economy and performance? Evidence from technology adoption. RFF Working paper, Resources for the Future, Washington, DC.

Li, Shanjun, Joshua Linn, and Erich Muehlegger. 2014. Gasoline taxes and consumer behavior. *American Economic Journal: Economic Policy* 6 (4): 302–42.

Li, Shanjun, Christopher Timmins, and Roger von Haefen. 2009. How do gasoline prices affect fleet fuel economy? *American Economic Journal: Economic Policy* 1 (2): 113–37.

Miravete, Eugenio J., Maria J. Moral, and Jeff Thurk. 2015. Innovation, emissions policy, and competitive advantage in the diffusion of European diesel automobiles. CEPR Discussion paper no. 10783, Center for Economic and Policy Research, London.

Mirrlees, James A. 1971. An exploration in the theory of optimum income taxation. *Review of Economic Studies* 38 (2): 175–208.

Parry, Ian, Margaret Walls, and Winston Harrington. 2007. Automobile externalities and policies. *Journal of Economic Literature* 45:374–400.

Petrin, Amil, and Kenneth Train. 2010. A control function approach to endogeneity in consumer choice models. *Journal of Marketing Research* 47 (1): 3–13.

Phaneuf, Daniel J., and Till Requate. 2017. *A course in environmental economics: Theory, policy, and practice*. New York: Cambridge University Press.

Ramsey, Frank P. 1927. A contribution to the theory of taxation. *Economic Journal* 37 (145): 47–61.

Revelt David, and Kenneth Train. 1998. Mixed logit with repeated choices: Households' choices of appliance efficiency level. *Review of Economics and Statistics* 80:647–57.

Sallee, James M., Sarah E. West, and Wei 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.

Salvo, Alberto, and Cristian Huse. 2013. Build it, but will they come? Evidence from consumer choice between gasoline and sugarcane ethanol. *Journal of Environmental Economics and Management* 66:251–79.

Schiraldi, Paschale. 2011. Automobile replacement: A dynamic structural approach. *RAND Journal of Economics* 42:266–91.

Sexton, Steven E. 2015. Automatic bill payment and salience effects: Evidence from electricity consumption. *Review of Economics and Statistics* 97 (2): 229–41.

Small, Kenneth A., and Kurt Van Dender. 2007. Fuel efficiency and motor vehicle travel: The declining rebound effect. *Energy Journal* 28 (1): 25–51.

Train, Kenneth. 1998. Recreation demand models with taste differences over people. *Land Economics* 74:230–39.

US EPA (US Environmental Protection Agency). 2014. Guidelines for preparing economic analyses. https://www.epa.gov/environmental-economics/guidelines-preparing-economic-analyses.

———. 2021. Inventory of U.S. greenhouse gas emissions and sinks: 1990–2019. https://www.epa.gov/sites/production/files/2021-04/documents/us-ghg-inventory-2021-main-text.pdf.

Verboven, Frank. 2002. Quality-based price discrimination and tax incidence: Evidence from gasoline and diesel cars. *RAND Journal of Economics* 33 (2): 275–97.

West, Jeremy, Mark Hoekstra, Jonathan Meer, and Steven L. Puller. 2017. Vehicle miles (not) traveled: Fuel economy requirements, vehicle characteristics, and household driving? *Journal of Public Economics* 145:65–81.