

Munk-Nielsen, Anders

Doctoral Thesis

Car ownership, type choice and use

PhD Series, No. 178

Provided in Cooperation with:

University of Copenhagen, Department of Economics

Suggested Citation: Munk-Nielsen, Anders (2015) : Car ownership, type choice and use, PhD Series, No. 178, ISBN 978-87-93428-01-0, University of Copenhagen, Department of Economics, Copenhagen

This Version is available at:

<https://hdl.handle.net/10419/233442>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



PhD Dissertation
Anders Munk-Nielsen



Car Ownership, Type Choice and Use

Car Ownership, Type Choice and Use

PhD Dissertation

Anders Munk-Nielsen

Department of Economics,
University of Copenhagen

Submitted: August 31, 2015

Main Supervisor: Bertel Schjerning

Co-supervisor: Søren Leth-Petersen

This dissertation is part of the IRUC research project financed
by the Danish Council for Strategic Research (DSF).

Financial support is gratefully acknowledged.

Contents

	page
Foreword	ii
English Summary	iii
Danish Summary	v
1. The Tail Wagging the Dog: Commuting and the Fuel Price Response in Driving	1
2. Diesel Cars and Environmental Policy	85
3. A Dynamic Model of Car Ownership, Type Choice and Usage	145

Foreword

Most kids grow up wanting to be astronauts but when I was a kid, I dreamed of being a truck driver. Although I have clearly failed in this regard, I hope that I at least earn some credit in the young eyes of my past self for having spent my time as a PhD student studying cars. Although the field of cars in some sense is fitting in this life-cycle perspective, the choice of the topic was entirely due to my supervisor, Bertel Schjerning. His passion and skill at what he does inspired me so profoundly that when he said that he had a project for a PhD, I immediately jumped at the chance. With the exception of my choice not to study law, this is the best decision I have made academically. I am deeply grateful for having been a part of the wonderful team of co-authors under the IRUC research project. Every time we have gone on one of our week-long intensive work sessions, I have come back afterwards energized and inspired in spite of having worked around the clock.

I am also deeply indebted to Søren Leth-Petersen, my co-supervisor. I worked for Søren as a research assistant and since on a joint project so he has followed me from the very start in economics. Søren and Bertel have complemented each other perfectly as a team of supervisors, teaching me the best from two sides of modern econometrics. I also want to thank Kenneth Gillingham for hosting me at Yale University and to Jesse Burkhardt for welcoming me there and teaching me the joy of climbing.

Of course, no dissertation would be here if it were not for the coffee bus. Without the long discussions about everything and nothing, academic and otherwise, who knows what I would have spent my time doing. And a special thanks for waiting for my slow-brewing coffee. I am sure that you all secretly suspected that it was just a test to see how much you enjoyed my company. The office mates for the final part of the PhD also deserve a thanks: to Thais for dragging me off to crossfit and to Patrick and Jeppe for listening to us brag about it afterwards.

Finally, I want to thank my family and friends and especially my girlfriend, Sigrid, for supporting me when times were tough and the lemonade got sour. And I want to thank my grandfather, Flemming Ib Nielsen, who has always encouraged my curiosity about everything in life, all the way from ABC to PhD. I dedicate this dissertation to him.

Anders Munk-Nielsen
Copenhagen, August 2015

English Summary

This PhD dissertation consists of three self-contained chapters on car ownership, type choice and use. They complement each other by taking three different angles on household decisions about cars; the first chapter explores the car use decision in great detail, focusing on how households change their driving decisions in response to changes in fuel prices. The second chapter zooms out and considers how households choose which car to purchase in the new car market. When they make the type choice decision, they take into account their driving behavior. The final chapter zooms out even further and looks at the problem in a dynamic context; here, households make decisions about driving, replacement, scrappage and car type in one joint framework. When they trade in the market for used cars it gives rise to equilibrium prices, which in turn shape the movements in the car fleet over time, forming *waves* of car vintages that travel through the car age distribution over time.

One of the key findings in chapter two is that fuel taxes are more effective at reaching environmental goals than car taxes. This has also been found by other studies in different settings and using different techniques. The implication is a clear policy recommendation about taxing fuel rather than cars. Chapter one documents that there is rich heterogeneity in the fuel price elasticity and shows how this translates into heterogeneity in the distributional consequences of fuel taxes. In particular, we find substantial heterogeneity in the deadweight loss depending on household work distance. Finally, chapter three documents that the reactions to fuel taxes in the used car market may have complicated dynamic impacts. If for example the fuel taxes are increased and the revenue from this is used to lower registration taxes, then the implication will be accelerated scrappage of older cars. This novel model framework lends itself nicely to policy analysis, providing a tool for understanding the changes in the car fleet over time from a more long-run perspective and with much greater realism than existing models are able to.

Below I go into greater detail about each chapter.

1. The Tail Wagging the Dog: Commuting and the Fuel Price Response in Driving

with Kenneth Gillingham

This chapter explores how driving by households responds to changes in fuel prices. This responsiveness is key to the welfare consequences of gasoline price changes. This study uses rich data covering the entire population of vehicles and consumers in Denmark, to find a one-year mean price elasticity of driving of -0.30 . We uncover an important feature of driving demand: two small groups of much more responsive households than the bulk of the population that make up the lower and upper *tails* of the work distance distribution. The first group lives in urban areas and close to work. The second group tends to live outside of major urban areas and has the longest commutes. We provide evidence that the

response to gasoline prices for both groups is mediated by access to public transportation, which is nearly universally available in Denmark. Further, we illustrate the importance of accounting for heterogeneity in the fuel price elasticity for the deadweight loss and environmental implications of changing gasoline prices. We find that raising fuel taxes by 1 DKK/liter implies a deadweight loss of 0.56 DKK/liter.

2. Diesel Cars and Environmental Policy

In this chapter, I investigate how households choose between different types of cars in the new car market. The purpose is to measure the costs of environmental taxation of car ownership and usage in Denmark. Using full population Danish register data covering 1997–2006, I estimate a discrete-continuous model of car choice and usage that explicitly allows households to select cars based on expected usage conditional on observed and unobserved heterogeneity. I validate the model using a major Danish reform in 2007 which prompted a substantial shift in the characteristics of purchased cars unique to the Danish setting compared to the rest of Europe. Through counterfactual simulations, I find that both Danish reforms in 1997 and 2007 were cost-ineffective at reducing CO₂ emissions compared to a fuel tax. Moreover, I find that the diesel market share responds strongly to taxation but that environmental goals can be reached both with and without a large diesel share in the fleet.

3. A Dynamic Model of Vehicle Ownership, Type Choice and Use

with Kenneth Gillingham, Fedor Iskhakov, John Rust and Bertel Schjerning

The focus of this chapter is to understand the equilibrium at used car market. Towards this end, we develop an estimable structural microeconomic model of car choice and usage that features endogenous equilibrium prices on the used-car market. Households buy and sell cars in the market and car owners choose how much to drive their car in a finite-horizon model. Moreover, we explicitly model the choice between scrapping the car or selling it on the used-car market. We estimate the model using full-population Danish register data on car ownership, driving and demographics for the period 1996–2009, covering all Danish households and cars. Simulations show that the equilibrium prices are essential for producing realistic simulations of the car age distribution and scrappage patterns over the macro cycle. We illustrate the usefulness of the model for policy analysis with a counterfactual simulation that reduces new car prices but raises fuel taxes. The simulations show how equilibrium prices imply that the boom in new car sales come at the cost of accelerated scrappage of older cars. Furthermore, the model gives predictions on tax revenue, fuel use, emissions, the lifetime of vehicles as well as the composition of types and ages of cars in the future.

Resumé

Denne PhD afhandling består af tre separate kapitler, der alle omhandler husholdningers beslutninger vedrørende køb og brug af biler. De supplerer hinanden ved at lægge tre forskellige vinkler på husholdningernes beslutninger om biler: Det første kapitel undersøger, hvordan husholdningernes kørsel påvirkes af brændstofpriserne. Det andet kapitel zoomer ud og ser på, hvordan husholdningerne vælger bil under hensyntagen til deres kørselsbehov. Det sidste kapitel zoomer helt ud og ser på ligevægtsdannelsen på brugtbilmarkedet. Her skal husholdningerne vælge, hvor meget de vil køre, hvilken bil de vil have, hvornår bilen skal udskiftes, samt om den skal skrottes eller sælges på brugtbilmarkedet. Når husholdningerne interagerer på brugtbilmarkedet opstår ligevægtspriserne, som skaber de aggregerede bevægelser i bilparken over tid, som kommer til udtryk som bølger bilparkens alder over tid.

Et af de vigtigste resultater i kapitel to er, at brændstofafgifter er et mere effektivt beskatningsværktøj end bilafgifter. Dette er også blevet fundet i studier i andre kontekster og ved brug af andre metoder. Implikationen er en klar politisk anbefaling om beskatning af brændstof snarere end af biler. Kapitel et dokumenterer heterogeniteten i, hvor prisfølsomme husholdningernes kørsel er. Denne heterogenitet i prisfølsomheden udmønter sig i heterogenitet i de fordelingsmæssige konsekvenser af brændstofafgifter. Især husholdningernes afstand til arbejde er afgørende for størrelsen af dødvægtstabet ved brændstofafgifter. Sluttelig viser kapitel tre, hvorledes reaktionerne i brugtbilmarkedet på ændringer i skatter og afgifter kan have afgørende betydning for den samlede effekt. Hvis eksempelvis der implementeres en kombination af højere brændstofafgifter og en lavere registreringsskat, så vil konsekvenserne være, at værdien af gamle biler ændres med fremskyndet skrotning til følge. Modellen fra kapitel tre er velegnet til policy-analyser med lidt længere sigt for øje, idet den fremskriver bilparkens aldersfordeling med langt større realisme, end hvad eksisterende modeller er i stand til.

I det følgende går jeg i dyden med hver af de tre kapitler.

1. The Tail Wagging the Dog: Commuting and the Fuel Price Response in Driving

med Kenneth Gillingham

Dette kapitel undersøger, hvordan husholdningernes kørsel reagerer på ændringer i brændstofpriserne. Denne prisfølsomhed er nøglen til velfærdskonsekvenser ved ændringer i brændstofafgifterne. Kapitlet bruger detaljerede data, der dækker hele befolkningen af forbrugere i Danmark samt alle køretøjer, og hovedresultatet er, at priselasticiteten for kørslen på et etårigt sigt er på -0.30 . Husholdningernes evne til at ændre deres kørsel demonstreres at hænge sammen med adgangen til offentlig transport en rolle i at mediere reduktionen i kørslen. Sluttelig illustreres det, hvordan heterogeniteten i priselasticiteten

udmønter sig i heterogenitet i dødvægtstabet ved brændstofafgifter. Dødvægtstabet ved at hæve brændstofafgifterne med 1 kr. pr. liter er beregnet til at udgøre 0,56 kr. pr. liter.

2. Diesel Cars and Environmental Policy

Dette kapitel undersøger, hvordan husholdninger vælger mellem forskellige typer af biler under hensyntagen til deres efterfølgende kørselsbehov. Formålet er at måle omkostningerne ved forskellige typer af miljømæssig beskatning af bilsektoren. Ved brug af danske registerdata for 1997–2006 estimerer jeg en model for det simultane valg af biltype og efterfølgende kørsel. Modellen tillader for, at husholdningerne vælger deres bil endogen på baggrund af både observerbare og uobserverbare karakteristika. Modellen valideres ved brug af en større dansk reform i 2007, som havde en betydelig effekt på husholdningernes typevalg. Af kontrafaktiske simulationer fremgår det, at både beskatningsreformen i 1997 og 2007 var mindre omkostningseffektive end en brændstofafgift ville have været. Et gennemgående resultat er desuden, at dieselandelen af salget af nye biler reagerer særdeles kraftigt på den relative beskatning af diesel- og benzinbiler.

3. A Dynamic Model of Vehicle Ownership, Type Choice and Use

Det tredje kapitel er fokuseret på at forstå ligevægtsprisdannelsen på brugtbilsmarkedet. For at gøre dette udvikles en ny strukturel, mikroøkonometrisk model for bilvalg og brug, med endogene priser. Modellen indeholder desuden en eksplicit modellering af valget mellem at sælge eller skrotte en bil. Modellen er estimeret på danske registerdata for 1996–2009, som dækker samtlige biler i den danske bilpark. Simulationer illustrer vigtigheden af at have endogene ligevægtspriser på brugte biler. Uden dette kan modellen ikke producere realistiske simulationer af aldersfordelingen af biler samt af skrotning. En simpel policy-analyse illustrerer, hvordan skatteændringer der påvirker forholdet mellem faste og variable omkostninger fører til en kraftig fremskyndelse af skrotning af ældre biler. Endvidere giver modellen forudsigelser om skatteindtægter, brændstofforbrug, emissioner, levetid af køretøjer samt sammensætningen af typer og aldre af biler i fremtiden.

Chapter 1

The Tail Wagging the Dog: Commuting and the Fuel Price Response in Driving

The Tail Wagging the Dog: Commuting and the Fuel Price Response in Driving*

Kenneth Gillingham, *Yale University and NBER*
Anders Munk-Nielsen, *University of Copenhagen*

December 3, 2015

Abstract

The consumer price responsiveness of driving demand is key to the welfare consequences of gasoline price changes. This study uses rich data covering the entire population of vehicles and consumers in Denmark, to find a one-year mean price elasticity of driving of -0.30 . We uncover an important feature of driving demand: two small groups of much more responsive households than the bulk of the population that make up the lower and upper *tails* of the work distance distribution. The first group lives in urban areas and close to work. The second group tends to live outside of major urban areas and has the longest commutes. We provide evidence that the response to gasoline prices for both groups is mediated by access to public transportation, which is nearly universally available in Denmark. Further, we illustrate the importance of accounting for heterogeneity in the fuel price elasticity for the deadweight loss and environmental implications of changing gasoline prices. We find that raising fuel taxes by 1 DKK/liter implies a deadweight loss of 0.56 DKK/liter.

Keywords: urban transportation, heterogeneity, environmental taxes, deadweight loss.

JEL classification codes: R2, R4, Q2, Q5.

*The authors wish to thank Bertel Schjerning, John Rust, Fedor Iskhakov and Søren Leth-Petersen for helpful comments and feedback. The authors would also like to acknowledge funding from the IRUC research project, financed by the Danish Council for Independent Research.

1 Introduction

Oil prices have historically been highly variable, with Brent crude spot prices ranging over the past decade from \$139/barrel at the peak in June 2008 to below \$50/barrel in January 2015. These large oil price gyrations lead to corresponding changes in gasoline prices, influencing transport decisions, congestion, and environmental outcomes. Understanding the price elasticity of driving, which underpins the price elasticity of gasoline consumption, is therefore of considerable policy interest. Not only is it valuable for anticipating responses to future swings in oil prices, it is also useful for measuring the macroeconomic effects of oil price fluctuations (e.g., Edelstein and Kilian, 2009) and providing insight into the role of speculators during oil price shocks (Hamilton, 2009; Kilian and Murphy, 2014). Furthermore, it forms the basis for measuring the welfare consequences of changes in gasoline prices.

This study estimates the price elasticity of driving and provides new insight into the underlying determinants of this elasticity. Using vehicle-level odometer readings matched to individual-level location and demographic information from the Danish registers, we uncover two small groups of households who are much more responsive to changing gasoline prices than most of the population. These households are in the *tails* of the work distance distribution; one group has very short commutes and the other has the longest commutes. Our mean elasticity estimate of -0.30 is considerably influenced by these two groups of tail households, each of which have an elasticity estimate of -0.6. These findings can be rationalized with a model of switching costs incurred when switching from driving to other modes of transport, such as public transport. Danes have almost universal access to public transport and we posit that our results hold in similar settings around the world.

This research contributes to three strands of literature with major policy importance. First, it provides a new point estimate for the gasoline price elasticity of driving, which is a dominant component in the modeling of gasoline demand. There is a vast literature aiming to estimate the price elasticity of gasoline demand (e.g., for some recent studies see Coglianese et al., 2015; Davis and Kilian, 2011; Hughes, Knittel, and Sperling, 2008; Li, Linn, and Muehlegger, 2014; Hymel and Small, 2015; Small and van Dender, 2007), largely using aggregate data at the regional or national level.¹ More recently, a handful of studies have estimated the elasticity of vehicle-miles-traveled with respect to the price of gasoline using disaggregated micro-level data, either from surveys or inspection odometer reading data (Linn, 2013; Bento et al., 2009; Knittel and Sandler, 2013; Gillingham, 2013, 2014; Munk-Nielsen, 2015; Levin, Lewis, and Wolak, 2014). Most of these recent short-run elasticity estimates are for drivers in the United States and are in the range of -0.15 to -0.35. Interestingly, similar benchmark estimates for Europe tend to show a more

¹Review articles cover dozens of studies going back decades, most using aggregate data. For example, see Dahl and Sterner (1991), Espey (1998), Graham and Glaister (2004), and Brons et al. (2008).

elastic response. For example, Frondel and Vance (2013) estimate a short-run driving elasticity with respect to the gasoline price of -0.45 in Germany.² Our study not only helps to reconcile these differing estimates across countries, but it also sheds light on the mechanisms underlying the differences. In particular, by identifying the tails in the distribution of consumer response and the reason for these tails, we can posit that there are groups of more-responsive households in Europe that simply do not exist in the United States.

Identifying the composition of the tail households contributes to a second vein of literature on the complex relationship between urban form, gasoline prices, and consumer decisions about how much to drive. There is growing evidence that urban form and the spatial structure of labor force demand affect travel choices and commuting behavior (Bento et al., 2005a; Grazi, van den Bergh, and van Ommeren, 2008; Brownstone and Golob, 2010). Since at least McFadden (1974), it has been long-recognized that access to public transport is an important mediator of travel choices, with clear environmental implications (e.g., Glaeser and Kahn, 2010). Denmark provides a very useful empirical setting for exploring these issues, for access to public transport is near-universal, yet there is considerable variation in commute distances and the degree of access to appealing substitutes to driving. Our findings are informative for the development of models of household location choice and access to public transport by revealing the detailed spatial relationship between location and driving.

The third strand of the literature to which we contribute is the analysis of environmental tax reforms on the light duty vehicle fleet. Several recent papers focus on vehicle registration tax reforms using discrete vehicle choice models (e.g. D’Haultf  uille, Givord, and Boutin, 2014; Adamou, Clerides, and Zachariadis, 2013; Huse and Lucinda, 2013). Without modeling the endogenous choice of driving, these papers can only calculate a rough estimate of the environmental implications of such policies. Other work incorporates the driving decision, for example in a discrete-continuous framework (Jacobsen, 2013; Gillingham, 2013; Munk-Nielsen, 2015; Grigolon, Reynaert, and Verboven, 2015) in order to evaluate environmental policies focused on vehicles. However, the computational complexity of such models prevents a sufficiently detailed modeling of the driving decision to fully capture the heterogeneity of the response. Our study provides a comprehensive picture of the consumer response on the intensive margin, which can help inform choice of salient features to include in discrete-continuous models of both vehicle choice and utilization intended to examine policies affecting both the intensive and extensive margins.

In this paper, we underpin our empirical analysis with a simple theoretical model that predicts the existence the tail households. A key feature in this model is the presence

²-0.45 is the fixed effects estimate, which we believe is better identified than other estimates in the paper, which are closer to -0.6.

of switching costs incurred when changing transport modes. Consider households with very high work distances. When fuel prices increase, these households stand to gain more from switching to public transportation and will therefore respond more strongly than households who do not commute as far. Households with very short commutes face a different decision problem, one in which nearly all driving demand is for non-commute trips such as shopping or leisure travel. For those households, the price sensitivity of the marginal kilometer in a shopping or leisure trips determines their overall elasticity. If leisure trips are more price sensitive than work trips, as would be expected, then this group of households will exhibit a greater price responsiveness. We present empirical evidence consistent with these explanations using both a quantile regression framework and a standard linear framework with a rich set of interactions to explore the determinants of greater price responsiveness.

We illustrate what our results mean for policy through a counterfactual analysis of a price increase of 1 DKK/liter for both gasoline and diesel fuel (i.e., just over \$0.50/gal). Decomposing the total response in driving, we find that the most-responsive 5% of drivers are responsible for 14.4% of the total reduction in driving. Moreover, we calculate the deadweight loss from this increase in fuel prices as 0.56 DDK/l, and show that this deadweight loss is highly heterogeneous. In fact, the deadweight loss for both highly responsive tails of households is more than four times greater than for the less responsive households in the middle of the work distance distribution. These results underscore the differing distributional consequences of changes in fuel taxes.

The remainder of this paper is organized as follows. The next section lays out our simple theoretical model to provide a framework for the economics underlying our results. Section 3 describes the Danish register data we are using and provides descriptive evidence on the primary features of the data relevant to estimating the driving responsiveness. Section 4 describes our empirical strategy, while section 5 presents the results and a set of robustness checks. Section 6 estimates the policy counterfactuals, while section 8 concludes.

2 A Simple Model of Travel Decisions

This section develops a simple model of the travel decision of a car-owning agent in order to build intuition. To focus on the economics relevant to our setting, we abstract from decisions about where to live, what employment to accept, and how much to travel for non-work trips. Instead, we focus on modeling the key features of how work distance can influence the price responsiveness of driving. Our model is well-suited for a setting where the decision-maker has access to public transport. Such a setting is relevant to nearly all of Denmark, as well as much of Europe and many other areas in the world. For example, in 2014, 87 percent of Danes live within one kilometers (km) of a public transport stop

and nearly all the remainder are served by on-call buses (“telebusser”).³ We model a static setting for a given finite amount of time, such as one week.

To simplify our setting, we hold the total number of km traveled by the agent fixed at T . The agent can travel by personal vehicle or by other modes of transport, including public transport, biking or walking. Let the km traveled by personal vehicle be denoted by v , so the remaining km traveled is $T - v$. Consider two types of travel. The first type is repeated travel that occurs several times a week, such as for a commute to work. The second is discretionary, shopping, or leisure travel. Let $d^w \in [0, 1]$ be the decision of how much to drive for commuting trips. $d^w = 1$ if all of commuting is accomplished by driving and $d^w = 0$ if all of commuting is done by other modes of transport. Similarly, let $d^l \in [0, 1]$ be the same decision for non-commuting (leisure) trips.

As driving is a more flexible form of transport, let $g^w(d^w)$ be the additional utility from commuting to work by driving rather than other forms of transport. Similarly, let $g^l(d^l)$ be the utility from driving for non-work trips. Assume $\frac{\partial g^l(d^w)}{\partial (d^w)} > 0$ and $\frac{\partial g^l(d^l)}{\partial (d^l)} > 0$. However, there is an important difference between the commuting trips and other trips that motivates our specification of these functions. While trips for shopping or leisure involve travel to a diverse set of locations, commute trips are from the same origin to the same destination and usually at the same time of day. Thus, for a given set of commute trips in a given time period, we would expect that the marginal utility from commuting by personal car rather than other forms of transport is constant, regardless of the amount of driving. This allows us to define $g^w(d^w) \equiv \gamma^w d^w$, where γ^w is a constant. In contrast, there is inherent heterogeneity in the ability to bike, walk, or take public transport for non-commute trips. For some shopping or leisure trips, public transport or biking are very attractive modes of travel; for others, they are highly unappealing due to the distance or destination. Thus, one would expect some curvature of $g^l(d^l)$, i.e., the marginal utility of driving will vary with the fraction of non-work trips driven: $\frac{\partial^2 g^l(d^l)}{\partial (d^l)^2} \neq 0$ (and we might expect that g^l is concave, but it is not necessary to assume this).

Denote the km traveled for work by w and the km traveled for non-commute trips by l . Consider an agent who maximizes utility subject to a budget constraint:

$$\begin{aligned} \max_{d^w \in [0,1], d^l \in [0,1]} & u(x) + g^w(d^w) + g^l(d^l) \\ \text{s.t.} & y \geq p^v v + p^b(T - v) + x, \end{aligned}$$

where x is the outside good (whose price is normalized to 1), y is total income, p^v is the price per km of driving, and p^b is the price per km of the non-driving mode.

³See <http://passagerpulsen.taenk.dk/file/68/download?token=fy19yEeh> (Accessed June 16, 2015).

Inserting the assumed form of g^w , the Lagrangian for this problem can be written as

$$\max_{d^w \in [0,1], d^l \in [0,1]} u(x) + \gamma^w d^w + g^l(d^l) + \lambda [y - (p^v - p^b)v - p^b T - x],$$

where λ is the shadow price or marginal utility of income.

We can now solve for d^w and d^l . Assuming standard regularity conditions and using $v = d^w w + d^l l$, the optimal leisure travel decision can be characterized by the following first-order condition:

$$\frac{\partial g(d^l)}{\partial d^l} = \lambda(p^v - p^b)l.$$

This condition is entirely standard; the household will choose the fraction of non-commute driving, $d^l \in [0, 1]$, so that the marginal utility of an additional kilometer traveled by car is equal to the marginal cost (converted in terms of utility). In other words, since shopping and leisure trips are heterogenous, the household will shift the least inconvenient trips to public transport, walking or biking when fuel prices increase. Of course, corner solutions at 0 and 1 are possible, if the marginal cost is sufficiently high or low. Otherwise, $\frac{\partial^2 g^l(d^l)}{\partial (d^l)^2} \neq 0$ and the monotonicity of $g^l(d^l)$ assures an interior solution.

The setting is different for commuting, since $\frac{\partial g^w(d^w)}{\partial d^w} = \gamma^w$. Given this, as long as we do not have exact indifference (i.e., $\gamma^w = \lambda(p^v - p^b)w$), a utility-maximizing household would never choose an interior solution. Instead, we obtain the following “bang-bang” solution for the choice of mode for commute travel:

$$d^w = \begin{cases} 1 & \text{if } \gamma^w \geq \lambda(p^v - p^b)w \\ 0 & \text{else.} \end{cases} \quad (2.1)$$

If the marginal utility from driving is greater than marginal cost (converted into utility units), then $d^w = 1$ and all commute trips are done by driving. Otherwise, all commute trips are done by other forms of transport, such as public transport, cycling, or walking. Even though the model is static, we can think of γ^w intuitively as a *switching cost* that prevents a change in commute driving unless there is a sufficiently large change in the marginal cost. It can be thought of as the marginal utility of driving over other forms of transport, which includes such factors as the effort in planning transport trips or the psychological cost of changing habits.

This framework has important implications for our empirical setting. We are interested in the fuel price sensitivity of driving and the heterogeneity in this sensitivity. That is, we are interested in $\frac{\partial v}{\partial p^v}$, holding w and l fixed. Consider the comparative statics with change in gasoline prices at the optimal values of d^w and d^l . From the implicit function

theorem we know that for leisure driving,

$$\frac{\partial d^l}{\partial p^v} = \frac{\lambda l}{\frac{\partial^2 g^l}{\partial (d^l)^2}} \quad (2.2)$$

For commute driving, the discontinuity in the optimal mode choice implies a discontinuity in the response so that the derivative is zero (almost) everywhere. We thus consider a change in gasoline prices leading to a change from p_0^v to p_1^v . Consumers will switch from driving to other modes of transport at the threshold $p^v = p^b + \frac{\gamma^w}{\lambda w}$. So the change in driving with the given change in gasoline prices is

$$\Delta d^w = \begin{cases} 1 & \text{if } p_1^v < p^b + \frac{\gamma^w}{\lambda w} < p_0^v, \\ -1 & \text{if } p_1^v > p^b + \frac{\gamma^w}{\lambda w} > p_0^v, \\ 0 & \text{otherwise.} \end{cases}$$

This expression highlights when switching might occur with a fuel prices rise. For example, in order for there to be a switch away from driving for commutes, the increase in the marginal cost of driving must be sufficient to overcome the marginal cost of the other option p^b plus the marginal utility of driving above other sources, scaled by the distance of the commute and put in monetary terms.

Now the response in total driving to the change in gasoline prices is given by:

$$\Delta v = \Delta d^w w + \Delta d^l l. \quad (2.3)$$

Thus, for households with very long commutes (i.e., a large w), a fuel price change sufficiently large to induce a switch would imply a much greater decrease in driving. This can be restated as our first testable implication:

Proposition 1. *Households with a longer work distances are expected to be more responsive to changes in gasoline prices over a period of time with sufficiently large gasoline price variation.*

Equation (2.3) also shows that for the households with the shortest work distances, Δv becomes determined entirely the change in leisure driving. According to equation (2.2), the price sensitivity for these households ultimately boils down to the curvature of the utility of driving for leisure travel ($g^l(\cdot)$). The underlying fundamentals determining the shape of $g^l(\cdot)$ are factors such as the availability of appealing substitutes to driving, the closeness of amenities, and the types of leisure activities that households with low work distances engage in. Our model imposes no a priori restriction on the curvature of $g^l(\cdot)$, so this is an empirical question. However, it is common to assume that leisure travel is more discretionary and thus may be more responsive to changes in gasoline price. Our second testable implication summarizes:

Proposition 2. *For households with very short work distances, the fuel price sensitivity of total driving approaches the fuel price sensitivity for leisure trips. To the extent that leisure trips are more discretionary, we would predict more responsiveness for households with very short work distances.*

This simple model lays a theoretical foundation for analyzing the heterogeneity of fuel price elasticity of driving. It is intentionally a simple static model to build intuition. We would expect to see the same switching behavior in a dynamic setting, whereby households could “invest” in switching if the discounted savings from doing so outweigh the switching cost.⁴ Since the savings are just the work distance times the fuel price differential, this implies that households with longer work distances will switch for smaller changes in the fuel price. Such a mechanism can be explained in our model by allowing γ^w to be heterogeneous in the population and increasing in w .

Several other simplifications may also be relaxed without changing the basic intuition. Allowing leisure travel demand to change endogenously in the model would be straightforward, but tedious. It would change the model quantitatively, but not qualitatively. When a household switches work trips from driving to public transportation, it will then use some of the savings towards increased leisure driving, slightly dampening the primary effect we show above. Endogenizing work distance would add further complexity, but would be possible with a model of household location choice. Although unquestionably important for longer-term policy, adding a location choice would again not change our primary findings, even if it would provide another channel that could dampen the effects we treat in the model above.

3 Data

To introduce the reader to the dataset, we first describe the data sources, then briefly describe the creation of the final dataset and finally present some descriptive evidence on both the mean and distribution of driving over the sample period.

3.1 Data Sources

This paper draws upon data from the Danish registers on the population of both households and vehicles in Denmark from 1998 to 2008. There are three main sources. The first is the vehicle license plate register, which contains the vehicle identification number, gross vehicle weight rating, fuel type, date of registration, owner identification number,

⁴The intuition is similar to the intuition in an (S, s) -model of portfolio choice; for small changes in the fuel prices, most households will stick with their baseline mode choice and avoid paying the switching cost. For larger changes, however, they will be forced to re-optimize.

and whether the vehicle type is a personal car or a van.⁵ The second is the inspection database. Starting on July 1, 1998, all vehicles in Denmark are required undertake a mandatory safety and emission inspection at periodic intervals after the first registration of the vehicle. In Denmark, the first inspection is roughly four years out, and subsequent inspections are every other year.⁶ Only a small number of used vehicles are imported into Denmark, in part because they pay the same vehicle registration fee and value-added tax that are assessed on new vehicles. The fee and tax schedule are primarily based on the value of the vehicle for both new and imported vehicles. Vehicles that are imported also undergo a registration and inspection at the time of importation. The inspection database contains odometer readings, which can be used to determine the amount of driving between two inspections.

The third primary data source is the household register, which contains detailed demographic data at the calendar year-level. These data include the number of members of the household, ages and sex of these members, municipality of the household, income of the household members, the assets of the household, and a measure of work distance used to calculate the tax deduction for work travel. This measure of work distance is the product of the reported number of days that work travel occurred and the reported work distance. Since the address of the work place is known to the tax authorities, this number is subject to auditing. The individual is only eligible for a deduction if the distance is greater than 12 km but there is no minimum requirement on the number of days worked. The work distance measure will therefore be equal to zero if the individual lives closer than 12 km from the work place, or if the individual does not work. More details and statistics are in Appendix A.3. For 2000 through 2008, we have data on the actual work distance for 79.61% of the households measured using a shortest-path algorithm and provided by Statistics Denmark.⁷ We find that these two measures are actually quite similar and can compare results using this measure to the results using the tax deduction measure as a robustness check (Appendix A.4.4).

In addition to the register data, we also bring in daily price data for octane 95 gasoline and diesel fuel from the Danish Oil Industry Association.⁸ Similarly, we also bring in daily West Texas Intermediate crude oil price data (www.eia.gov/dnav/pet/hist/LeafHandler.ashx?n=PET, Accessed June 15, 2015). Finally, we use data from Journey Planner (www.journeyplanner.dk, Accessed April 19, 2013) on all bus and train stops in Denmark in 2013.

We also have access to some additional car characteristics, namely fuel efficiency in

⁵Company cars are not in our database and are not linked to a person but rather to the firm. However, individuals with access to a company car must pay a tax for this, and we observe that (3.7% of our households have at least one member paying this tax).

⁶This is a very similar schedule to inspections in states in the United States, such as California. Details about the driving period lengths are in Appendix A.2.1.

⁷Statistics Denmark has access to the actual addresses of individuals. This information, however, is anonymized in our dataset so we cannot perform any operations based on GIS information.

⁸See www.eof.dk, Accessed June 17, 2015.

km/l and the manufacturer suggested retail price (MSRP). This data comes from a dataset from the Danish Automobile Dealer Association (DAF). However, these variables are not available for car vintages older than 1997 so we have not included them in the preferred specification but instead use them for robustness checks.

3.2 Development of the Final Dataset

We combine the data from the various sources to create a final dataset where the unit of observation is a vehicle driving period between two inspections. So if a driver has a first inspection of her vehicle on June 1, 2004 and the next inspection on June 6, 2006, the driving period will be the 735 days between these two tests. We use the difference in odometer readings between these two inspections to calculate the total kilometers driven and the kilometers driven per day over the driving period. Similarly, we calculate the average gasoline, diesel, and oil price over the same driving period. If a car changes owners during a driving period, we include an observation for both households that have contributed to the driving. We also create a variable measuring the fraction of the driving period that each household owns the car.

To match our calendar year demographic data with driving periods, we construct a weighted average of the values of the demographic variables over the years covered by the driving period. For example, if a driving period covers half of 2001, all of 2002, and half of 2003, the demographic variables values would be given a weight of 0.25 for 2001, 0.5 for 2002, and 0.25 for 2003. We merge in the public transport stops by the count at the municipality level. For a detailed description of the variables used, see appendix A.3.

In the data, we have 7,254,893 driving periods that can be matched to the owner of the vehicle. Of these, we delete 377,708 driving periods based on driving period less than two months long and 277,294 periods that are missing demographic variables. Finally, in order to use household fixed effects, we delete the 744,267 driving periods pertaining to households that are only observed with a single driving period. Thus, the final dataset consists of 5,855,446 driving period observations covering nearly all driving periods by Danish drivers over the period from 1998 to 2011. Table 3.1 presents summary statistics for the final dataset. Appendix A provides further details on the data cleaning process.

3.3 Descriptive Evidence

There has been considerable variation in both gasoline and diesel prices in Denmark over the 1998 to 2008 period. Figure 3.1 shows average gasoline and diesel prices over time in our dataset. The x-axis denotes the time of the inspection at the beginning of the driving period.⁹ Figure 3.1 also plots the average vehicle kilometers traveled (VKT) over

⁹Appendix A contains a similar graph of daily prices of gasoline and diesel fuel that is not averaged over driving periods.


Table 3.1: Summary Statistics

	Mean	Std.
Real gross income (DKK)	574055.9	(627921.3)
Real gross income (couples)	646638.1	(628011.1)
Real gross income (singles)	320975.4	(558097.7)
Couple (D)	0.777	(0.416)
Age (oldest member)	49.83	(14.39)
Work Distance (km)	12.19	(19.73)
Work Distance > 12 km (D)	0.500	(0.50)
Work Distance (actual, km) ^a	23.44	(35.55)
# of kids	0.761	(1.023)
Urban (D)	0.159	(0.362)
Company car (D)	0.0338	(0.181)
Self employed (D)	0.0981	(0.297)
# periods observed	4.650	(2.723)
Bus stops per km ²	15.86	(18.42)
VKT (km/day)	46.59	(40.17)
Weight (ton)	1671.4	(330.5)
Diesel (D)	0.143	(0.350)
Van (D)	0.0783	(0.269)
Percent owned of period	0.794	(0.299)
Driving period length (years)	2.388	(0.890)
Car age ultimo (years)	6.972	(5.166)
# cars owned	0.338	(0.602)
# vans owned	0.0537	(0.244)
# motorcycles owned	0.0530	(0.266)
# mopeds owned	0.0270	(0.159)
# trailers owned	0.303	(0.599)
Observations	5855446	

A “D” denotes a dummy variable.

^a: Only available for 76.17% of the sample.

Figure 3.1: Vehicle Kilometers Traveled and Fuel Price



the driving period, illustrating a negative relationship between fuel prices and driving. The unconditional distribution of VKT is shown in Figure A.8.

The rich Danish register data allows us to explore the relationship between fuel prices and driving in greater detail. Figure 3.2 divides the sample into ten groups based on the percentiles of driving in each year. The vehicles in each group may change over time, as we recalculate the percentiles in each year. The figure shows a fascinating pattern. For most groups, there appears to be very little change in driving over time, even as fuel prices change significantly. However, the 1 percent of drivers who drive the most show a noticeable decrease in VKT during driving periods that begin between 2003 and 2005, just as gasoline prices are rising. This provides the first descriptive evidence of the existence of one tail of more responsive drivers.

Who are these drivers who drive the most? Table 3.2 shows the demographics and other characteristics of drivers based simply on the amount driven. Not surprisingly, higher VKT drivers have a higher income, have more vehicles, and live further away from their workplaces. They also tend to drive heavier and younger cars and diesel cars more frequently. Otherwise, these higher VKT drivers are similar to the average driver in most other characteristics. For details about the individual variables, see appendix A.3.

To visualize where the high-driving households live, we supplement the summary statistics of Table 3.2 with a map showing the spatial distribution of VKT. Figure 3.3 shows a map of Denmark where each municipality is colored according to the average VKT of the households having their private address in that municipality. The figure shows that the high-VKT households tend to be in or very close to the major urban areas, in particular Aarhus, Odense as well as the Copenhagen metropolitan region as

Table 3.2: Average Characteristics by VKT


	VKT<100	VKT $\geq$ 100
Real gross income (DKK)	571327.5	645390.1
Real gross income (couples)	644262.1	705237.5
Real gross income (singles)	319519.5	369007.6
Couple (D)	0.7754	0.8220
Age (oldest member)	50.0	45.1
Work Distance (km)	0.4374	0.6269
Work Distance > 12 km (D)	11.6	26.9
Work Distance (actual, km) ^a	22.7	39.3
# of kids	0.7515	1.0040
Urban (D)	0.1597	0.1365
Company car (D)	0.0337	0.0376
Self employed (D)	0.0959	0.1561
# of periods observed	4.6236	5.3430
Bus stops per km ²	15.9	14.2
VKT (km/day)	42.6	151.4
Weight (ton)	1665.0	1836.4
Diesel (D)	0.1292	0.5037
Van (D)	0.0765	0.1258
Percent owned of period	0.7957	0.7357
Driving period length (years)	2.3813	2.5730
Car age ultimo (years)	7.0321	5.4133
# cars owned ^b	0.3276	0.5958
# vans owned	0.0535	0.0599
# motorcycles owned	0.0528	0.0588
# mopeds owned	0.0270	0.0270
# trailers owned	0.3018	0.3220
Observations	5639738	215708

A “D” denotes a dummy variable.

^a: Only observed for 76.17% of our sample

^b: Number of cars in excess of the current.

Figure 3.2: VKT Percentiles Over Time




well as the region to the north of Copenhagen. When interpreting this it is important to remember that the dataset conditions on car ownership; the correct interpretation of the map is that if a household chooses to have a car in the city, they tend to use it more intensively than a corresponding household living in a rural area. Of course, we would expect that car ownership rates are much lower in the cities. In the appendix, Figure A.12 shows the average work distance by municipality, which shows that municipalities with high work distances tend to coincide with high driving.

For a driver who lives further from work to be able to reduce driving, they must have access to public transport. In many countries, such as the United States, such access tends to be very limited. Figure A.10 illustrates the prevalence of public transport access throughout Denmark. As shown in the figure, there are bus or train stops nearly everywhere in Denmark. Moreover, there is on-call public transport available in rural municipalities where the stops are sparser. This pervasiveness of public transport may be essential for allowing switching behavior for those in the tail of the distribution of responsiveness.

In Appendix A.4.2, we show the relationship between driving and four demographic variables: work distance, car age, income and household age. The work distance relationship is monotonically increasing, from 40 km/day for work distances below 12 km to 70 km/day for households with work distances of 90 km. If we assume that households drive to work each day, then the graph indicates that leisure driving for the households with the shortest driving in the sample is about 80–90% of total driving. For the households with the highest work distances, they drive on average about 20% less than they would have to in order to commute to work on every working day by car. This indicates that they

Figure 3.3: Average VKT by Municipality



must either not be going to their work place every day or be using alternative methods of transportation at least on some occasions.

4 Methodology

In this section, we will first go through the various econometric models to be estimated in section 5. Next, we discuss the variation in the data that identifies the model. Finally, we derive explicit formulas for welfare calculations, which will be used in section 6.

4.1 Econometrics

A primary goal of this paper is to investigate the fuel price elasticity and to explore the heterogeneity in that parameter. Following a vast literature on estimating fuel price elasticities, we use a linear log-log specification. This specification not only provides for a ready interpretation of the coefficient of interest, but we find that it also fits the data very well.

Consider the demand for driving for household i during a driving period t . Note that t differs from a typical panel data time dimension in that a household may have two cars driving at the same time, which will be represented by two separate observations. In the most general form, which nests all of the specifications used in this paper, we model the

demand for driving with the following random parameters model:

$$\log \text{VKT}_{it} = \gamma_{it} \log p_{it} + \mathbf{x}_{it} \boldsymbol{\beta}_{it} + \sum_{t=1998}^{2011} \sum_{f=\text{gas,diesel}} \delta_{fy} \omega(i, t, y) + \eta_i + \varepsilon_{it}. \quad (4.1)$$

VKT_{it} is the average daily driving in km and p_{it} is the average daily fuel price over the driving period (gasoline or diesel depending on the car type) and $\mathbf{x}_{it}$ denotes a vector of controls. In $\mathbf{x}_{it}$ we include controls for whether and by how much the driving period overlaps with other driving periods by the same household.¹⁰ The coefficient, γ_{it} , is our primary coefficient of interest, which is the fuel price elasticity for household i in driving period t . A key focus of this paper will be to explore the heterogeneity in this coefficient, so we think of it as a parameter that varies across observations. We employ fixed effects at the household level, η_i .

The variable $\omega(i, t, y)$ captures time controls. In our primary specification, $\omega(i, t, y)$ denotes the fraction of driving period t that falls within the year $y \in \{1998, \dots, 2011\}$. So if a driving period starts on July 1st 2001 and ends on June 30th 2003, $\omega(i, t, y)$ will be 0.25 for $y = 2001, 2003$ and 0.5 for $y = 2002$. The coefficients δ_{fy} will therefore have a similar interpretation to a model with fuel type specific year fixed effects even though ω are continuous variables. The reason why we cannot include traditional year fixed effects is that a driving period is not exclusively in one year but rather covers two to five years. Since the weights sum to unity, we omit year 2003 as the reference year. As robustness, we also try alternative specifications, including having only one set of time controls ($\delta_{fy} = \delta_y$) and simplifying to a linear time trend ($\omega(i, t, y) = y, \delta_{fy} = \delta$).

We use three different approaches to parameter heterogeneity: a standard linear fixed effects model, a panel quantile regression and a linear fixed effects model with interactions. The linear fixed effects version of (4.1) takes the form

$$\log \text{VKT}_{it} = \gamma_0 \log p_{it} + \mathbf{x}_{it} \boldsymbol{\beta} + \sum_{t=1998}^{2011} \delta_{fy} \omega(i, t, y) + \eta_i + \varepsilon_{it}. \quad (4.2)$$

This model is useful for estimating the mean elasticity for comparison to other studies. We also focus most of our robustness checks around this model. The dependent variable corresponds to the total driving by car, which we denoted v . Note that we do not observe the mode choices for the work and leisure trips, d^w, d^l , but only the composite outcome v . The work distance, w , is included in the regressors, $\mathbf{x}$, and is key for identifying the switching behavior.

¹⁰If the car changes owner mid-way through the driving period, the driving period is included as an observation by both households and we add a control for the percent of the driving period each household owns the car. We also add controls for ownership of other vehicles that do not admit driving observations such as motorcycles, mopeds, trailers, etc. See Appendix A.2.

Our second model is the conditional quantile model

$$\begin{aligned} \log \text{VKT}_{it}(\tau) &= \gamma(\tau) \log p_{it} + \mathbf{x}_{it} \boldsymbol{\beta}(\tau) + \sum_{t=1998}^{2011} \delta_{fy}(\tau) \omega(i, t, y) + \eta_i + e_{it}(\tau), \\ Q_{e_{it}(\tau)}[\tau | \mathbf{x}_{it}, p_{it}, \omega(i, t, y)] &= 0, \eta_i \perp\!\!\!\perp e_{it}(\tau) | \mathbf{x}_{it}, p_{it}, \omega(i, t, y). \end{aligned} \quad (4.3)$$

This model is motivated by the shift in the upper quantiles of the VKT distribution shown in figure 3.2. We are interested in assessing whether this shift in the upper quantiles can be attributed to increasing fuel prices, which we can explore by examining whether $\gamma(\tau)$ increases as $\tau \rightarrow 1$.

We estimate the parameters using the panel quantile estimator of Canay (2011), which proceeds in two steps: first, we use a standard panel regression to obtain the within-estimate of $\hat{\eta}_i$. Second, we construct the regressand $\tilde{y}_{it} := y_{it} - \hat{\eta}_{it}$. Third, we run the pooled quantile regression of $\tilde{y}_{it}$ on all our regressors. We are aware of other approaches to dealing with person-specific unobservables but the computational advantage of the Canay (2011) approach is appealing given the size of our dataset.¹¹

Finally, we consider a linear model with interactions between a subset of our controls, $\mathbf{x}_{it}^1$, and the log of the fuel price to allow the fuel price elasticity to vary with demographics in a linear fashion, $\gamma_{it} = \gamma_0 + \boldsymbol{\gamma} \mathbf{x}_{it}^1$. The subset includes all demographic variables and car-related variables but excludes the time and period controls because that would suck up all the identifying variation. Thus, the third specification we estimate is the following model, which we estimate using a standard linear fixed effects estimator,

$$\log \text{VKT}_{it} = \left(\gamma_0 + \boldsymbol{\gamma} \mathbf{x}_{it}^1 \right) \log p_{it} + \mathbf{x}_{it} \boldsymbol{\beta} + \sum_{t=1998}^{2011} \delta_{fy} \omega(i, t, y) + \eta_i + \eta_{m_{it}} + \varepsilon_{it}. \quad (4.4)$$

This model allows us to explore the underlying factors behind the heterogeneity in γ_{it} and relate them to observables. The virtue of this approach is the simplicity since this model can be estimated using OLS. However, the model places no restrictions on the values that γ_{it} so in particular we may get positive values. It is theoretically possible that some people respond to rising fuel prices by increasing their driving; if driving is a complement to an activity that is negatively correlated with fuel prices, then the driving will inherit this correlation. Other authors have encountered upwards-sloping demand curves when studying driving or fuel demand. For example, Blundell, Horowitz, and Parey (2012)

¹¹For example, Abrevaya and Dahl (2008) use a Chamberlain-style random effects estimator, projecting η_i on covariates from all periods or the time-averages. This essentially amounts to adding more regressors and the running a pooled quantile regression. Koenker (2004) on the other hand takes a high-dimensional penalization approach, treating the η_i s as N additional parameters to be estimated, and penalizing the sum of absolute values of the fixed effects in the spirit of the LASSO estimator. Using clever computational tricks, he makes the approach feasible in CPU terms but with the size of our dataset, we would run into problems with RAM usage.

consider fuel demand and formulate a non-parametric estimator that imposes negative elasticities over the entire region of the fuel price. They argue that the upwards sloping regions are due to small-sample error.¹² In our setting, it is hard to argue that any positive-elasticity regions of the demand curve would be due to small sample problems and we tend to believe that the regions are actually due to omitted time-varying variables. In our estimation, we do not impose positivity.

4.2 Identification

The central relationship of interest in this paper is between the fuel price and VKT. Since there is no geographical variation in the fuel price in our data, we rely on time series variation in fuel prices. As shown in Figure A.7, the primary variation in our empirical setting comes from the oil price. Since the price of oil is determined on the world market and Denmark is a small market, it follows that the price setting process for fuel prices in Denmark can reasonably be considered exogenous. It is unlikely that Denmark-specific demand shocks affect fuel prices. However, there may still be common demand shocks across countries that affect driving demand across countries and thus might affect prices. We include time controls and run robustness checks to confirm that this is not appreciably biasing our results.

In addition, we employ household fixed effects, so identification is coming from variations from the mean household driving under different fuel prices. When considering a variable that primarily varies over time, one should always be wary of unobserved, time-varying factors that affect driving. One example could be if the dropping prices of flight tickets have caused more households to use flights instead of their cars when going on holidays. By employing time controls, we address any trends over time that influence all households in Denmark. The precise timing of the driving period is emphasized, leveraging for example the difference between a driving period beginning in January and December in the same year. This approach relies on a large dataset for sufficient precision and a long-enough time horizon; studies that use only a short time horizon will often be unable to disentangle a change in driving into the response to fuel prices and secular shifts in driving.

In a sense, the task of estimating the fuel price elasticity requires assuming that the general trend in driving not associated with fuel prices follows a smooth path over time. In the extreme case where we allow the time path to be completely nonparametric, the fuel price coefficient becomes unidentified. This is a general problem when estimating the fuel price elasticity without cross sectional variation.¹³

¹²An alternative strategy would be to modify (4.4) so that $\gamma_{it} = \varphi(\gamma \mathbf{x}_{it})$, where $\varphi : \mathbb{R} \mapsto (-\infty; 0]$. This would require a normalization, for example on the constant in γ . Then the model could be estimated using for example the method of sieves.

¹³With cross sectional variation, one might worry as to what generates the spatial variation in fuel

4.3 Deriving the Deadweight Loss

In section 6, we analyze the welfare implications of a fuel tax increase leading to a change in the fuel price from p_0 to p_1 . In particular we consider the welfare consequences of such a policy intervention. Assuming that the Danish market is sufficiently small that the world market price is unaffected by such a tax, there will be no supply side response to the policy. Therefore, we focus on the deadweight loss to consumers by analyzing the change in consumer surplus based on changes from the current fuel tax, and consider this the total change in deadweight loss. Note that our calculations do not require that the fuel tax is initially set at the socially optimal level. The deadweight loss derived below is the same as the standard “triangle area” in a linear model. However, we find that a log-log model fits the data very well (Figure 7.1) so we develop a deadweight loss formula that is consistent with that model.

We start by deriving the deadweight loss for the single agent model. It is straightforward to extend this formula to take into account parameter heterogeneity. This framework can then be applied to compute the deadweight loss for each of the estimated models and it can be evaluated either for the average household or we can evaluate it individually and average over households. This will allow us to assess the importance of taking into account heterogeneity.

The preferred specification estimated in this paper is a log-log regression, i.e. of log VKT on log fuel prices and controls. By exponentiating both sides, this is algebraically equivalent with the following Cobb-Douglas demand for VKT,

$$\text{VKT} = \varepsilon p^\gamma \prod_{k=1}^K z_k^{\theta_k},$$

where p is the fuel price and z contains all the other observables, $z = (x, \omega, \eta)$, and $\theta = (\beta, \delta, 1)$. All the parameters here can be taken directly from the log-log estimates. Suppose that the world-market supply of fuel is inelastic with respect to the Danish price of fuel. Then the deadweight loss from a tax increase causing the price to rise from p_0 to p_1 is given by the “triangle” between the VKT demand curve and the inelastic supply

prices. If fuel stations change prices depending on regional differences in the fuel price elasticity, then the price-setting behavior poses an endogeneity problem. Indeed, our results suggest that there are clear patterns in the spatial distribution of fuel price elasticities, with fuel stations in oligopolistic competition.

curve at p_1 over the interval $[p_0; p_1]$, i.e.,

$$\begin{aligned} \text{DWL}(p_0, p_1, \gamma) &= \int_{p_0}^{p_1} \text{VKT}(p) - \text{VKT}(p_1) dp \\ &= \varepsilon \prod_{k=1}^K z_k^{\theta_k} \left[\frac{1}{1+\gamma} p^{1+\gamma} \right]_{p_0}^{p_1} - \varepsilon p_1^\gamma \prod_{k=1}^K z_k^{\theta_k} (p_1 - p_0) \\ &= \text{VKT}(1) \left[\frac{1}{1+\gamma} (p_1^{1+\gamma} - p_0^{1+\gamma}) - p_1^\gamma (p_1 - p_0) \right], \end{aligned}$$

since $\text{VKT}(1) = \varepsilon \prod_{k=1}^K z_k^{\theta_k}$. The deadweight loss will be positive because $\text{VKT}(p) > \text{VKT}(p_1)$ for all $p \in [p_0; p_1)$. However, $\text{DWL}(p_0, p_1, \gamma)$ is not necessarily monotonic on $\gamma \in (-1, 0)$. This non-monotonicity stems from the fact that the Cobb-Douglas form implies that for a fixed γ , a higher baseline VKT implies a higher deadweight loss. This is not true for a linear demand curve.¹⁴

Now, let us suppose that agents are heterogeneous both in terms of observables, $p_{it}, \mathbf{z}_{it}$, and parameters, indicated by subscripts (i, t) . The demand equation is then

$$\text{VKT}_{it}(p) = \varepsilon_{it} p^{\gamma_{it}} \prod_{k=1}^K z_{itk}^{\theta_{itk}},$$

where once again $z_{it} = (x_{it}, \omega(y, i, t), \eta_i)$ and $\theta_{it} = (\beta_{it}, \delta, 1)$. Given this demand equation, we can compute the standard deadweight loss as the integral

$$\text{DWL}_i(p_0, p_1) = \varepsilon_{it} \prod_{k=1}^K z_{itk}^{\theta_{itk}} \left[\frac{1}{1+\gamma_{it}} (p_0^{1+\gamma_{it}} - p_1^{1+\gamma_{it}}) - p_1^{\gamma_{it}} (p_1 - p_0) \right]$$

Inserting estimated parameters and evaluating the expression at the average error, $\varepsilon_{it} = 1$, we can compute this expression for all observations and compute the average deadweight loss in our sample.¹⁵

We will now derive an expression for the deadweight loss for a quantile model. Of course, we can simply take the quantile regression estimates for a given quantile $u \in [0; 1]$ and insert them in equation (4.5), setting $\gamma_{it} = \hat{\gamma}(u)$. Doing so will teach us something about what the deadweight loss looks like for different conditional quantiles of driving.

¹⁴If we had assumed a linear VKT demand, the deadweight loss would instead be given by

$$\int_{p_0}^{p_1} \gamma p + \mathbf{x}\beta + \varepsilon - (\gamma p_1 + \mathbf{x}\beta + \varepsilon) dp = \gamma \left[p_1 p_0 - \frac{1}{2} (p_1^2 + p_0^2) \right],$$

which is monotonic in γ . Note that the γ here should come from estimating a linear regression of VKT on fuel prices, both in levels. Our strict emphasis on the log-log model is motivated by semi-parametric demand curve, which we present in Figure 7.1; this illustrates the appropriateness of the log-log form.

¹⁵If we knew the distribution of ε , we could integrate it out for each household. However, we do not pursue this added precision. It will only be a concern if there is heteroscedasticity that might lead to a very skewed distribution so that the averaged expression would become different from the expression evaluated at the average of ε .

However, to ensure comparability with the results from the interacted model, we will instead now outline a procedure for obtaining a single measure of the deadweight loss for each household, taking into account the full range of quantile estimates.

To do this, we will consider an alternative formulation of the quantile model in equation (4.3); Koenker (2005, ch. 2.6) argues that instead of writing the equation for a given quantile, u , we may think of the model equivalently as each observation (i, t) randomly drawing a uniform quantile, u_{it} , and then being assigned parameters according to the quantile regression function, $\gamma_{it} = \gamma(u_{it})$ and $\beta_{it} = \beta(u_{it})$. With this in mind, we can write the model as

$$\log \text{VKT}_{it} = \gamma(u_{it}) \log p_{it} + \mathbf{x}_{it} \beta(u_{it}) + \eta_i, \quad u_{it} \sim \text{Uniform}(0, 1).$$

This random coefficient formulation is an equivalent way of thinking about the quantile regression model, which makes it easier to think about u_{it} for the purpose of calculating the deadweight loss;¹⁶ if we observed u_{it} , it would just be a matter of plugging it into (4.5).

Since u_{it} is not observed but has a known density (which we have already assumed for consistency of the quantile regression), we can integrate it out. This is in the spirit of Melly (2005) and Machado and Mata (2005). Thus, we replace the unobserved, latent deadweight loss with the *Integrated Deadweight Loss* (IDWL) given by

$$\begin{aligned} \text{IDWL}_i(p_0, p_1, 1) \\ = \int_0^1 \left(\prod_{k=1}^K z_{itk}^{\theta_k(u)} \right) \left[\frac{1}{1 + \gamma(u)} \left(p_0^{1+\gamma(u)} - p_1^{1+\gamma(u)} \right) - p_1^{\gamma(u)} (p_1 - p_0) s \right] du. \end{aligned} \quad (4.5)$$

This integral can be computed in a number of different ways; Machado and Mata (2005) use a simulation approach in a somewhat similar setting and Melly (2005) uses a grid. Given the computational requirements for estimating the model at even a single quantile (approximately 10 hours on a 64-core machine with 1 TB of RAM), we are forced to use a grid and let the computer time dictate the fineness of the mesh.¹⁷ We have worked with 21 grid points, comprising of $\{ .01, .05, .10, .15, \dots, .95, .99 \}$. Thus, we approximate the

¹⁶Note that the correlation between the random parameters is one since the only source of randomness is u_{it} , which is scalar.

¹⁷Portnoy (1991) shows that with a finite sample, the estimated quantile regression function, $u \mapsto (\hat{\gamma}(u), \hat{\beta}(u))$, will only change at a finite number of points on the interval $[0, 1]$ and that this number is $O(N \log N)$. Melly (2005) notes that for his estimator of the conditional distribution based on the quantile predictions, a mesh size on the order of $O(N^{-.5-\epsilon})$ will ensure that the asymptotics still hold. For computational reasons, we are unable to scale up accordingly.

IDWL as

$$\begin{aligned} & \text{IDWL}_i(p_0, p_1, 1) \\ & \cong \sum_{q=1}^{21} w_q \left(\prod_{k=1}^K z_{itk}^{\hat{\theta}_k^{u_q}} \right) \left[\frac{1}{\hat{\gamma}^{u_q} + 1} \left(p_0^{\hat{\gamma}^{u_q} + 1} - p_1^{\hat{\gamma}^{u_q} + 1} \right) - p_1^{\gamma^{(u_q)}} (p_1 - p_0) \right], \end{aligned} \quad (4.6)$$

where $\hat{\gamma}^{u_q}$ and $\hat{\beta}^{u_q}$ are the estimated coefficients from the u_q 'th quantile regression, $u_q \in \{.01, .05, .1, \dots, .90, .95, .99\}$, and where w_q is equal to the length of the corresponding intervals.¹⁸

5 Results

In this section, we present our econometric results. First, we present the estimates of the mean elasticity using the standard linear fixed effects estimator. Then we show how the fuel price elasticity varies over the conditional quantiles of VKT in a quantile model, forming an inverse U-shape. Finally, we analyze the underlying determinants of the remarkable heterogeneity in fuel price elasticity using a set of interactions with various demographic variables in a linear fixed effects model.

5.1 The Mean Elasticity

Table 5.1 shows results from estimating linear fixed effects model where we assume no heterogeneity in γ and β (equation (4.2)). The parameters omitted here are included for reference in the appendix in table B.1. The first specification, column (1), shows a fuel price elasticity of -0.866 only controls for car characteristics, seasonality (month controls) and controls for the period.¹⁹ When we add year controls and demographics in column (2), the elasticity drops to -0.298. This indicates the importance of controlling for precise individual-level demographics as well as using time controls. In columns (3) and (4), we add household fixed effects. Without the time controls, the elasticity is -0.515 but including them reduces the elasticity to -0.304, which is our preferred estimate for the elasticity. It may not be surprising that the elasticity becomes closer to zero when we control for general time trends in driving; as discussed in Section 4.2, our primary variation in the fuel price is time-series variation. In Appendix Table C.5 we show that the elasticity is highly robust to the functional form of the time controls.²⁰ The ability to simultaneously control for household fixed effects and time controls is a unique advantage

¹⁸So $w_q = .01$ for $q = 1, 21$, $w_q = .04$ for $q = 2, 20$ and $w_q = .05$ otherwise.

¹⁹For details about the individual variables and a full list of variables, the reader is referred to appendix A.3, tables A.1 and A.2.

²⁰Even in a specification with just a linear time trend in the starting year of the period, the elasticity is -0.313. In other words, it is important to control for a shift in average driving but the elasticity is not sensitive to the precise specification.

of our data, combining full population data with a long, 10 year time horizon.

It is worth noting that the difference in the fuel price elasticity is relatively small when we compare the specifications with and without household fixed effects in columns (2) and (4) (-0.298 vs. -0.304). We take this as an indication that our rich set of controls are capturing the most important determinants of the fuel price elasticity. In particular the controls for work distance, company cars and income (including transfers) appear to capture key components of driving demand as evidenced by an R^2 of 0.340 in column (2), which is quite high for micro data studies in a specification without household fixed effects.

When we inspect the estimates in the results in Table 5.1 we see that the effect of work distance is positive for both males, females and singles, indicating that households with longer commutes tend to drive more. Recall that the work distance information is censored by construction at 12km for households working 220 days a year so we have included a dummy for work distance being observed as well as a linear term.²¹

The dummy for urban residency is positive in the specification without household fixed effects. This seems at odds with the map of driving in Figure 3.3, which showed a slightly lower average VKT in major urban areas. However, recall that this is conditional on all other observables and Figure A.12 shows that the average work distance in urban regions is very small. So the coefficients indicate that *conditional* on how low their work distance tends to be in urban regions, those households tend to drive more. When we add household fixed effects in column (4), the urban dummy changes sign to be negative. In this specification, the identifying variation comes from households that have lived both in urban and rural areas and it shows that they drive less (conditional on work distance etc.) when they live in the urban areas. This difference can also be interpreted as the specification without household fixed effects suffering from selection bias; the types of households that choose to own a car in spite of living in the urban regions may have high driving demand for unobserved regions, since car ownership in urban regions in Denmark is generally much lower than in the rest of the country. For example, many Danish households use bikes for short trips and in particular in cities. In a survey of Danish households, DTU Transport (2013) find that 20% of all trips to work are done by bike. For trips shorter than 5km, close to 80% are carried out by walking or biking.

The bus stops variable becomes insignificant in the preferred specification in column (4). However, this variable has no time-series variation due to data availability and only varies by municipality, so it is purely identified from households moving between municipalities. When there are more kids present in the household, driving tends to be higher. Interestingly, this has the reverse sign before controlling for household fixed

²¹Details on the variable are in Appendix A.3. In Appendix A.4.4, we examine the validity of our primary work distance variable by comparing it with a measure from another source available for a subset of the data.

Table 5.1: Log-log Model of VKT

	OLS		Household FE	
	(1)	(2)	(3)	(4)
	No demo	Base	FE	Main
$\log p^{\text{fuel}}$	-0.866*** (0.00509)	-0.298*** (0.0143)	-0.515*** (0.00722)	-0.304*** (0.0154)
<i>Work Distance (WD) controls</i>				
WD, male		0.00234*** (0.0000207)	0.00251*** (0.0000339)	0.00242*** (0.0000336)
WD non-zero, male		0.0670*** (0.000745)	0.0299*** (0.00108)	0.0329*** (0.00107)
WD, female		0.00304*** (0.0000284)	0.00315*** (0.0000446)	0.00303*** (0.0000443)
WD non-zero, female		0.0581*** (0.000801)	0.0216*** (0.00111)	0.0257*** (0.00111)
WD, single		0.00420*** (0.0000484)	0.00428*** (0.0000837)	0.00419*** (0.0000835)
WD non-zero, single		0.165*** (0.00164)	0.0669*** (0.00243)	0.0724*** (0.00243)
<i>Age controls</i>				
Age, male		0.0176*** (0.000248)	0.00330 (0.00806)	0.0212** (0.00813)
Age, female		0.0145*** (0.000249)	0.0103 (0.00805)	0.0468*** (0.00813)
Age, single		0.0171*** (0.000211)	0.00460*** (0.000853)	0.0598*** (0.000939)
Age squared, male		-0.000222*** (0.00000248)	-0.0000945*** (0.0000112)	-0.0000930*** (0.0000112)
Age squared, female		-0.000203*** (0.00000261)	-0.000197*** (0.0000115)	-0.000195*** (0.0000115)
Age squared, single		-0.000305*** (0.00000212)	-0.000207*** (0.00000767)	-0.000206*** (0.00000767)
<i>Other demographic controls</i>				
log gross inc (couple)		-0.0318*** (0.000619)	-0.0276*** (0.00163)	-0.0242*** (0.00162)
log gross inc (single)		0.0249*** (0.000456)	0.0221*** (0.00288)	0.0199*** (0.00287)
Urban (dummy)		0.00357*** (0.000959)	-0.0135*** (0.00284)	-0.0249*** (0.00284)
# of kids		0.00607*** (0.000258)	-0.0177*** (0.000652)	-0.0168*** (0.000650)
Company car		-0.186*** (0.00135)	-0.103*** (0.00217)	-0.0977*** (0.00216)
Self employed		0.0275*** (0.000822)	0.00642*** (0.00136)	0.000712 (0.00136)
Bus stops per km ²		-0.00117*** (0.0000188)	-0.000313*** (0.0000544)	0.0000419 (0.0000548)
Year controls	No	Yes	No	Yes
% of each month	Yes	Yes	Yes	Yes
Car	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Household FE	No	No	Yes	Yes
N	5855446	5855446	5855446	5855446
R^2	0.198	0.340	0.175	0.180

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

effects, indicating that the households who never have children might differ along some unobserved dimension.

Another novel feature of our dataset is that we are able to control for the availability of a company car due to this being reported to the tax authorities.²² In our sample, 3.8% of the households have access to a company car. Note that the car doing the driving is not a company car but privately owned; the variable merely indicates that one of the household members has *access* to using a company car. Therefore, it makes sense that the dummy has a negative sign (-0.098). This implies that those households drive almost 10% less than other households in their privately owned car, *ceteris paribus*. The self employment dummy is insignificant, but this may be due to little time-series variation within a household so that it is captured by the household fixed effect.

Regarding the car characteristics, shown in table B.1, we see that households tend to drive longer if they have younger and heavier cars, where the gross weight accounts for car quality. Moreover, diesel cars are driven much further; this is in line with the reasoning that such cars are more expensive to buy but cheaper to use. Hence, households with a higher driving demand will self-select into this group.

5.2 Heterogeneous Elasticity: Quantile Results


We estimate the panel quantile regression model (4.3) using the Canay (2011) 2-step fixed effects estimator as described in Section 4.1.²³ The results are shown visually in Figure 5.1 and numerically in Tables B.2 and B.3. The figure shows a clear inverted U-shape in the fuel price elasticity; the lower and the higher conditional quantiles of VKT have distinctly higher fuel price elasticities (numerically) than the middle region. This is consistent with the predictions of the theoretical model; households with very long commutes stand to gain a lot from switching from commuting to work to using public transportation. On the other hand, households with low driving demand might be more responsive if their driving is made up of a diverse set of trips with good substitutes available.

In Table B.3, we present the coefficients on the demographic variables for the quantiles 1, 50 and 99. They indicate that the coefficients are stable over the distribution of VKT for many demographic variables, including age and children. Conversely, the controls for urban residency, income, company car presence, self employment and bus stops change. The large negative company car coefficient (-0.312) for the 1st percentile of VKT indicates that the households with very little driving in the private car might be doing much of their driving using the company car. Similarly, in the bottom percentile of driving, self employment is negative (-0.082) but in the top it is positive (0.070). If the self employment

²²For a more thorough description of the company car tax rules, see Table A.2.

²³We do not correct the standard errors for the two-stage estimation even though Canay (2011) does provide an analytic expression for the asymptotic variance, taking into account the effect of the first-stage estimation of η_i . Canay also conjectures that a bootstrap procedure will be valid but the computational burden proved to be too great for our application (an estimation for a single quantile takes 10 hours).

Figure 5.1: Elasticity by Conditional Quantile



and company car availability in the sample increases, these coefficients imply that we would expect to observe more cars that are driven very little and more cars that are driven very much, i.e. that the distribution would get heavier tails.


5.3 Heterogenous Elasticity: Interactions

Although the quantile results allow us to study the elasticity in detail and illustrate the higher elasticity of the “tail households” in the tails of the driving distribution, they are not informative about *who* falls into the tails. To corroborate our intuitive understanding of the high-elasticity households, we instead turn to the linear specification with interactions of equation (4.4). We allow the fuel price elasticity to vary with household demographics and car characteristics, but not with the year, month or period controls.²⁴ Table 5.2 the interactions involving demographic variables.

The most interesting interaction is with work distance where we find the underlying determinants of the inverted U-shape seen in Figure 5.1. For easier overview, Figure 5.2 shows the result graphically; we have first calculated the individual-level predicted elasticities as $\hat{\gamma}_{it} = \hat{\gamma}_0 + \mathbf{x}_{it}^1 \hat{\boldsymbol{\gamma}}$. Next, we divide the work distance (defined as the maximum within household for couples) into 20 quantiles. Since work distance is not observed between 0 and 12 km, 11 of these occur at a work distance of zero (see Figure A.5). Finally, we compute the average elasticity within each of these bins and plot them in Figure 5.2. The figure shows an inverted U-shape over the work distance; for the shortest work distances, the fuel price elasticity is relatively high (-0.30) but then drops (numerically) for the households with work distance just over 12 km (-0.05), after which it increases mono-

²⁴Allowing the fuel price elasticity to vary over time will remove the remaining variation that we are using for identification. To identify such a model, we would either rely on functional form assumptions or we would need to have cross-sectional variation in fuel prices.

Figure 5.2: Fuel Price Elasticity and Work Distance



tonically (numerically) as work distance increases (up to about -0.45 for work distances just over 70 km). Both ends of the inverted U-shape in Figure 5.2 can be understood from the point of view of the model from section 2; for very high work distances, smaller increases in fuel prices are required in order for a switch to public transportation to recover the switching costs. For very low work distances, driving will mostly be made up of leisure driving, which consists of a more diverse set of trips where some will likely be easily substitutable for biking or public transportation.

When we turn to the individual coefficients on work distance interactions by male, female and singles, we find that there are highly significant differences in patterns by gender and single status; the primary driver of the elasticity pattern is males with high work distances. For females and singles, the effect is not really there for the high work distances. On the other hand, it turns out that married males by far have the highest work distances (see Appendix A.4.4 for details). This indicates that either females tend to have jobs that are available almost everywhere or couples tend to locate close to the female's work place and let the male take the longer commute. One would expect singles to be able to relocate more easily but we still see some singles with fairly high work distances in our sample, which appear to be less responsive (in light of the work distance interactions for singles in Table 5.2).

Turning to the other estimates, the interaction between the public transit measure (bus or train stops per km^2) implies that households living in a region with one standard deviation more stops available (18.4 stops/ km^2) will have a 25.2% larger absolute elasticity at the mean. This indicates that the access to public transportation is allowing households to more easily switch away from driving by car. This has important ramifications for policy, since it indicates that provision of public transportation may help improve the effectiveness of fuel taxes in reducing driving by car. Moreover, note that the inclusion

of an urban dummy interaction effect shows that the coefficient is not driven by a simple comparison of urban and rural regions, but is also driven by differences outside of the urban areas. The coefficient is identified by within-household comparisons for households that have lived in both municipalities with many and few bus and train stops. Note, however, that we have not time-series variation in the variable.²⁵ Even if we are cautious about making strong causal claims, we are still contributing to a very sparse body of empirical evidence on public transit provision based on revealed preference data.

The rest of the coefficients are intuitive; households with kids are less responsive and younger households are more responsive. Column (1) indicates that high-income households are more responsive, but once we allow the elasticity to vary with the car characteristics as well the interaction becomes insignificant, indicating that income effects are captured fully by the car choice. Interestingly, self employment and company car availability is associated with a much higher responsiveness; this can be seen as an indication that such households have alternative modes of commuting available than their private car. Table B.4 furthermore shows that multi-car households are more responsive, consistent with within-household switching, potentially towards the more cost-effective car as suggested by De Borger, Mulalic, and Rouwendal (2013). Also, drivers of diesel cars are much more responsive than gasoline car drivers. Diesel cars are typically more expensive to buy up front but cost less to drive, so this makes sense because households with large driving demand will have self-selected into the diesel segment. Drivers of vans tend to be less price sensitive. This may be due to vans typically being used in relation to work activities although most are diesel-driven, which moves the elasticity in the other direction.²⁶

Based on the estimates from the interacted linear model, we can compute the predicted elasticity for each individual as $\hat{\gamma}_{it} = \hat{\gamma}_0 + \hat{\gamma}\mathbf{x}_{it}^1$. As reported in table 5.2, the average elasticity is -0.238 although the distribution has a long right tail and even has observations with positive predicted elasticities. We note that the average predicted elasticity from the interaction results is -0.238 but the distribution has a long right tail, consistent with the theoretical model. To better understand the interplay of all the interactions in one, we have split the sample based on whether γ_{it} falls above or below the 5th percentile of elasticities. The characteristics of the two groups are shown in table B.5.

We see that the high-elasticity households tend to have substantially higher work distances, higher incomes (note, however, that the income interaction was insignificant), have fewer kids, are more often self employed or have access to a company car and they tend to drive younger cars and diesel cars and own more cars. We also note that they have as good or maybe slightly better access to public transportation.

²⁵One might be worried about endogeneous provision of public transit by local governments. In that sense, we are not using variation concerning how provision changed in the period, which anecdotally is not very much. The variable is measured in 2013.

²⁶67.7% of vans are diesel-driven against only 9.8% of personal cars.

Table 5.2: Heterogeneous Elasticity: Log-log Model with Interactions

	(1)	(2)	(3)
Mean elasticity	-0.253	-0.288	-0.238
$\log p^{\text{fuel}}$	-0.879*** (0.240)	-3.847*** (0.0862)	-4.698*** (0.283)
<i>Work Distance (WD) interactions</i>			
WD, male $\times \log p^{\text{fuel}}$	-0.0104*** (0.000337)		-0.00870*** (0.000352)
WD, female $\times \log p^{\text{fuel}}$	0.00304*** (0.000419)		0.00451*** (0.000423)
WD, single $\times \log p^{\text{fuel}}$	0.00204** (0.000715)		0.00342*** (0.000715)
WD squared, male $\times \log p^{\text{fuel}}$	-0.00000602*** (0.000000225)		-0.00000606*** (0.000000224)
WD squared, female $\times \log p^{\text{fuel}}$	-0.00000773*** (0.00000101)		-0.00000779*** (0.00000101)
WD squared, single $\times \log p^{\text{fuel}}$	-0.0000117*** (0.000000597)		-0.0000117*** (0.000000598)
WD non-zero, male=1 $\times \log p^{\text{fuel}}$	0.103*** (0.0129)		0.116*** (0.0141)
WD non-zero, female=1 $\times \log p^{\text{fuel}}$	0.185*** (0.0127)		0.191*** (0.0128)
WD non-zero, single=1 $\times \log p^{\text{fuel}}$	0.325*** (0.0266)		0.326*** (0.0267)
<i>Other demographic interactions</i>			
Couple=1 $\times \log p^{\text{fuel}}$	1.044*** (0.280)		0.464 (0.334)
$\log$ gross inc (couple) $\times \log p^{\text{fuel}}$	-0.0855*** (0.0118)		-0.00352 (0.0139)
$\log$ gross inc (single) $\times \log p^{\text{fuel}}$	0.0847*** (0.0221)		0.0403 (0.0265)
Urban (dummy) $\times \log p^{\text{fuel}}$	-0.0143 (0.0183)		-0.0342 (0.0184)
Age (oldest member) $\times \log p^{\text{fuel}}$	0.0246*** (0.00137)		0.0229*** (0.00146)
Age squared (oldest) $\times \log p^{\text{fuel}}$	-0.000249*** (0.0000133)		-0.000227*** (0.0000138)
# of kids $\times \log p^{\text{fuel}}$	0.187*** (0.00451)		0.183*** (0.00491)
Bus stops per km ² $\times \log p^{\text{fuel}}$	-0.00406*** (0.000349)		-0.00327*** (0.000360)
Self employed=1 $\times \log p^{\text{fuel}}$	-0.176*** (0.0135)		-0.175*** (0.0142)
Company car=1 $\times \log p^{\text{fuel}}$	-0.408*** (0.0234)		-0.396*** (0.0239)
Car x p	No	Yes	Yes
Year controls	Yes	Yes	Yes
Month controls	Yes	Yes	Yes
Car	Yes	Yes	Yes
Period	Yes	Yes	Yes
Demographics	Yes	Yes	Yes
Household FE	Yes	Yes	Yes
N	5855446	5855446	5855446
R^2	0.182	0.182	0.185

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

5.4 Geographical Analysis

We turn now to a geographical analysis of the results. Figure 5.3 illustrates the spatial location of the most responsive households by showing the 10th percentile of the distribution of the fuel price elasticity within each municipality. Recall that the lowest 10% are the most responsive. The map shows that the most responsive households tend to live in the rural regions in the outskirts of Sealand. Taken together with the map showing work distance in Figure A.12, this is consistent with those households commuting to Copenhagen from the rural areas. Looking at the map of bus and train stops in Figure A.10, we see that there are major public transportation lines connecting even these regions to central Copenhagen.

Two additional features of Figure 5.3 are worth noting; firstly, the major urban areas appear to have higher elasticities (numerically). This is in part captured by the urban dummy interacted with fuel prices (-0.0342). Secondly, the region to the north of Copenhagen also has some of the most elastic households. These areas tend to have wealthy, high-educated households working nice jobs and the public transportation in to Copenhagen is very good. Taken together, these two groups tend to have both low work distances (cf. Figure A.12) and low driving (cf. Figure 3.3), so they are likely to be a major part of the lower tail we have identified both in the quantile regression and the interaction results. The spatial configuration is remarkably similar to that of the average elasticity, which is shown in Figure B.1.


Figure 5.3 also shows that the elasticity pattern is not just driven by the urban dummy; the small municipalities in the suburban regions surrounding Copenhagen appear to show a continuously increasing elasticity (numerically) as the distance to the central business district of Copenhagen increases. This gradual change in the elasticity is driven by the demographic composition of the household living in these regions and from the map of work distance in Figure A.12 indicates that the pattern is driven to a large extent by the work distance.

6 Counterfactual Simulations and Welfare

In this section, we analyze the implications of our empirical findings for the welfare effects of an illustrative increase in fuel prices by 1 DKK/l for both gasoline and diesel. The average gasoline price over the 1998 to 2008 period is 9.01 2005 DKK, so this represents a substantial price increase.²⁷ Such an increase may be due to a fuel tax policy or exogenous swings in oil prices. While there is no evidence in Denmark, there is some evidence that consumers in the United States respond more to gasoline price swings than to changes in

²⁷This maps to an increase in gasoline prices of \$0.57 per gallon based on the June 18, 2015 exchange rate of 6.54 DKK per dollar.

Figure 5.3: The 10th Percentile of Elasticities by Municipality



gasoline taxes (Li, Linn, and Muehlegger, 2014). To the extent that these differ, then this counterfactual can best be thought of as an analysis of an exogenous fuel price increase rather than a tax policy. For ease of exposition, we will describe this price change as a policy change in the remainder of this section. We begin with a discussion of the overall implications of this policy for tax revenue and emissions. Then we explore how the heterogeneity in our elasticity influences the response to the fuel tax across the quantiles of driving, which has important implications for the distributional effects of the policy. Finally, we compute the deadweight loss using the different methods outlined in Section 4.3 and examine the distributional consequences of the policy based on the household work distance.

With a fuel price elasticity of driving of -0.30 , the proposed increase of 1 DKK/l translates into a 3.33% reduction in driving and it implies that total tax revenue goes up by 13.23%.²⁸ In terms of emissions, if households do not respond to the price on the extensive margin, i.e., by changing their cars, and if all vehicles respond in the same way, then our -0.30 estimate implies an elasticity of carbon dioxide or local air pollutant emissions with respect to the fuel price of -0.30 . Munk-Nielsen (2015) estimates a discrete-continuous model of car choice and driving and finds that when fuel prices increase by 1%,

²⁸At 9 DKK/l, the increase of 1 DKK/l is 11.1%, which at an elasticity of -0.30 translates to a change in driving of 3.33%. Over the sample period, taxes make up 64.87% of the gasoline price, corresponding to 5.84 DKK/l at 9 DKK/l. An increase in 1 DKK/l thus corresponds to an increase of 17.13% in taxes, giving a total relative change in taxes of $(1 + 0.1713) \times (1 - 0.0333) = 13.23\%$.

the fuel efficiency of newly purchased cars only increases by 0.1%. Hence, the -0.30 may only be slightly higher than the true short-run to medium-run change in emissions from changing fuel prices. Fully analyzing the effect is other important vehicle externalities, such as congestion and accidents (Mayeres and Proost, 2013), is outside the scope of this paper, but our results can shed some light on how these external costs would change. For example, since the high work-distance tail households appear to live outside of the city, one might posit that congestion—which mostly occurs around major urban areas—may not be as strongly affected.

We see how the heterogeneity in the fuel price responsiveness can map into differing distributional consequences of a change in fuel prices by computing individual-level responses in VKT. Table B.6 in Appendix B.2 uses the quantile regression estimates to explore how much of the total change in VKT can be attributed to each of the conditional quantiles. The results indicate that the top 5% quantile accounts for 14.4% of the sum of the predicted responses in driving for the population. By comparison, even though the lower tail also has a high elasticity, the bottom 5% only accounts for 4.77% of the total predicted response. In this sense, the top tail of drivers with long work distances is clearly more important for the welfare and environmental implications of the reform. Intuitively, one might think that the log-log functional form would be biased towards a result like this given that it implies a constant elasticity. However, the quantile regression allows the price parameter to vary freely over the quantiles of VKT. Even if the true relationship were less than proportional, the quantile regression could fit the data well and provide useful estimates.

Next, we turn to the deadweight loss for both the linear model and integrated quantile regressions. The formulas for the deadweight loss in each case are in Section 4.3. Recall that these formulas calculate the deadweight loss as the loss in consumer surplus (not accounting for any external costs). For each of the models, we can either evaluate the deadweight loss at the sample average of the observable characteristics or at the observed characteristics and then take the average across observations. Since the deadweight loss is a non-linear function, the two will yield different results. Table 6.1 shows the results.

Rows (1) and (2) show the results for the preferred log-log model from Section 5.1. Row (1) evaluates the deadweight loss for the preferred elasticity of -0.304, which gives a deadweight loss of 0.592 DKK/l for the fuel tax of 1 DKK/l. This is the standard, single agent deadweight loss calculation from the log-log model. The number may seem quite high but recall that gasoline taxes in Denmark account for 64.9% of the total price paid by consumers. In row (2), we evaluate the deadweight loss at the sample values of the observable characteristics entering into the equation for each observation. We should expect a different result using this approach because we are now taking the average of a nonlinear function rather than evaluating the nonlinear function at the average. Using the individual values for each observation results in a deadweight loss of 0.706 DKK/l and

Table 6.1: Deadweight Loss

				Deadweight loss	
Model	$\mathbf{x}_{it}^a$	Elasticity	Mean	sd	
<i>Estimates from Section 5.1</i>					
(1)	Log-log	Sample avg.	−0.304	0.592	–
(2)	Log-log	Individual	−0.304	0.706	0.41
<i>Estimates from Section 5.2</i>					
(3)	Quantile	Sample avg.	Median: −0.233 ^b	0.556	–
(4)	Quantile	Individual	Median: −0.233 ^b	0.658	0.38

^a: The column $\mathbf{x}_{it}$ indicates whether observables were evaluated at the *sample average* values or at the *individual*-level observed values.

^b: For the quantiles, the elasticity is integrated out, thus using the entire set of quantile estimates (equation (4.6)).

a standard deviation of 0.41 DKK/l. This illustrates how the considerable heterogeneity in responsiveness maps into heterogeneity in the deadweight loss.

Rows (3) and (4) use the integrated deadweight loss based on the quantile regression estimates inserted in equation (4.6). Row (3) evaluates at the average characteristics, giving a deadweight loss of 0.556 DDK/l, and row (4) takes the sample average, giving 0.658 DDK/l. These are only about 7% lower than the corresponding numbers from (1) and (2).

Finally, we investigate the distributional consequences of the fuel tax by work distance. To do this, we use the individual-level predicted elasticities from the model with interactions from Section 5.3. The heterogeneity in the elasticity is going to create heterogeneity in the deadweight loss. In Figure 6.1, we have divided the work distance (for couples we take the maximum) into 20 quantiles and calculated the deadweight loss within each quantile. Note that 11 of these quantiles fall at a work distance of zero. As economic intuition would suggest, the graph shows that the most elastic households have the greatest deadweight loss; at the highest and lowest work distances, the loss is approximately 1 DKK/l, while at the middle part of the work distance distribution there are much smaller losses (around 0.2 DKK/l). These findings indicate that the deadweight loss depends strongly on work distance, a novel finding in the literature.

7 Robustness

In this section, we go through a number of different robustness checks we have performed. First, we will summarize robustness with respect to the way our sample is constructed, then regarding the empirical specification, and in Section 7.1 we will discuss endogeneity concerns. We will be focusing on the effects of the specific issues on our parameter of

Figure 6.1: Deadweight Loss by Work Distance




Table 7.1: Overview of Robustness Checks: Worst-case Elasticity Estimate

Name	Elasticity	Table
Years in the sample	$[-0.402; -0.279]$	C.1, C.2
Length of driving periods	$[-0.298; -0.275]$	C.4
Fuel type	$[-0.541; -0.257]$	C.8
Singles or couples	$[-0.318; -0.250]$	C.3
Time controls	$[-0.517; -0.304]$	C.5
Clustering	$p = 0.011$	C.5
Linear specification	n.a.	Fig. 7.1
Instrumenting with oil price	-0.368	D.1

interest, the fuel price elasticity, and how it is changed from the central estimate of -0.304 from table 5.1. Table 7.1 provides an overview of the different robustness checks, showing the highest and lowest elasticities that came out in each case; in many cases, the extreme elasticities are perfectly expectable, so we comment on them in the text below, going through each case in turn.

Regarding the selection of the time-window, we have chosen driving periods that start between 1998M07 and 2007M12. The latest periods in the sample therefore end in 2011. We have chosen this window to ensure the most homogeneous composition of driving periods (see figure A.4). Tables C.1 and C.2 show the implications of starting the period later or ending it earlier. The two most extreme specifications yield elasticities of -0.279 and -0.402 . Considering that our primary source of variation is over time, we do not consider these to be very large deviations from our preferred estimate. Moreover, our time controls are very flexible and identification of the fuel price elasticity therefore rests on the long time horizon. So when we remove years from the sample, we are reducing the

primary source of variation. In that sense, it is comforting that the elasticity is not overly sensitive to selection on time periods.

For our sampling scheme, we have focused on including as many driving periods as possible and controlling for characteristics rather than dropping atypical observations. For example, we have included all driving periods that are between 1 and 2.5 years or between 3.5 and 4.5 years long. A typical driving period should be 2 or 4 years, plus or minus 3 months, so in table C.4 we first control for periods of irregular length and then drop them; adding a control drops the elasticity to -0.298 and dropping the 1.3 million irregular periods drops it to -0.275 . Considering how many periods are dropped, we do not think this is a substantial change in the central elasticity.²⁹

A key decision in the empirical design has been to jointly model gasoline and diesel car users. Essentially, we are imposing the restriction that households in the two different types of cars respond similarly to the fuel price (gasoline or diesel respectively). Table C.8 shows the results when we estimate the preferred specification on the two segments separately as well as where we add an interaction between log fuel price and the diesel dummy. The interaction shows that the gasoline elasticity is -0.257 while that for the diesel segment is -0.392 . Estimating on separate samples yields corresponding elasticities of -0.268 and -0.541 . This shows two things; firstly, the central elasticity is not solely identified by the differential fuel prices of the two segments. Secondly, it highlights that the diesel segment is more price sensitive. This is consistent with the theoretical model in the sense that diesel drivers generally have longer work distances, making them closer to the switching threshold.

Similarly to having both fuel types, we have also included observatoins for both couples and single households. Table C.3 shows results from estimating on the couple and single subsamples separately. This yields elasticities of -0.318 and -0.250 respectively. There are many differences between couples and singles so the magnitude of these differences is within reasonable bounds. We explore these differences more when we look at interactions in section 5.3.

The year controls employed in the main specification are highly flexible, which might be a problem given that fuel prices primarily vary over time. At the same time, table 5.2 shows that controlling for time effects is important to capture the substantial decrease in mean driving that has occurred between 1998 and 2011, cf. Figure 3.1. Therefore, table C.5 shows estimation results where we vary the time controls. In one extreme, we can reduce the time controls to a single, common linear time trend. This produces an elasticity of -0.313 , which is very close to the baseline of -0.304 . Thus, our results are not sensitive to how we control for the shift in mean driving over time but it is important

²⁹Figure A.3 shows the distribution of the length of the driving periods in the sample. We note an excess mass of periods falling before their “planned” inspection date at 2 or 4 years. One explanation for this may be that households wanting to sell their vehicle can expedite the inspection if they want to sell the car in order to assure prospective buyers that the car is in good condition.

to control for this. We have also included controls for seasonality by adding the fraction of the period that falls in each month with April as the baseline. That is, if a period is precisely 2 or 4 years long, these 11 controls would all be equal to $\frac{1}{12}$. In table C.5 we also remove these month controls from the preferred specification and it does not affect the elasticity at the third decimal. In Table C.6 we use the *number* of times each month occurred in the driving period rather than the fraction of the driving period falling in each month (so for a period of precisely 2 years, these 11 controls all equal 2). This leaves the elasticity unchanged on the third decimal at -0.304 . The coefficient estimates from Table C.5 clearly show the seasonal pattern of driving in Denmark; driving is highest in the summer — in particular in July where most workers have most of their 5 annual weeks of holiday — and lowest in the winter months. Overall, our results are highly robust to the functional form for the general time trends in driving.


In all the results we have presented, the standard errors have been very tight due, in part, to the large sample size. However, we do not account for regional patterns in the driving variable. This is both because we have not been able to obtain fuel price data with spatial variation but also because spatial variation in prices may be endogenous to the spatial variation in fuel price elasticities, which we have shown to be substantial.³⁰ To explore whether this potential measurement error is contaminating our standard errors, we want to use clustered standard errors. As we use household fixed effects, we can only cluster by variables that nest the households completely over time.³¹ Therefore, since households move over time and are observed at different time periods, we can neither use spatial nor temporal clusters. Instead, we investigate the effects of clustering in a model without the household fixed effects. Table C.7 shows that even when we cluster at the start-year level, yielding only 10 clusters, we still have marginal significance ($p = 0.011$). So while the heteroscedasticity robust standard errors from the primary specifications may have a downwards bias, we do not believe it to be of great concern. We discuss this further in appendix section C.5.

Finally, we explore whether the linear functional form for fuel prices in the log-log specification is appropriate. To do this, we plot a semi-parametric demand curve of log VKT against log fuel price (figure 7.1). The graph is constructed following the Robinson (1988) double-residual approach; first, log VKT and log fuel price are both *residualized* by regressing out the effect of all the remaining regressors from the primary specification. Then, the residualized log VKT can be regressed nonparametrically against the residualized log fuel price, using a local polynomial regression. Standard errors naturally have

³⁰Many papers rely mainly on spatial variation in fuel prices and this may be fine if the local markets are sufficiently segmented. However, in Denmark the regions are so close that the variation might be endogenously set by fuel stations. Anecdotally, an app for iPhone has been made in recent years that lets commuters plot their commute path and then provides them with the cheapest spot to buy fuel.

³¹There are methods out there to allow non-nested clustering but that is beyond the scope of this robustness check.

Figure 7.1: Semi-parametric demand curve



to take into account of this two-step approach, which may for example be done following a bootstrap procedure vis-a-vis Blundell, Horowitz, and Parey (2012). We refrain from doing standard errors and instead focus on the point estimate. Figure 7.1 shows a clear linear form, which indicates that the log-log specification appears to be appropriate.

7.1 Exogeneity

In this section, we discuss two concerns regarding the exogeneity of our primary specification; endogeneity of fuel prices and of car characteristics.

In any econometric estimation of a demand curve, it is important to consider endogeneity of the price variable. First off, we argue that since Denmark is a small country in the global market for fuel, which is a highly traded good, and thus developments particular to Denmark are unlikely to affect fuel prices. Moreover, Danish fuel taxes have not been changed discretionarily in the period, except a few minor changes (see Figure A.7). However, local market shocks may affect local prices. Therefore, we instrument for the fuel price using the West Texas Intermediate crude oil price (see Section 3 and Appendix D). The 2SLS fuel price elasticity estimate is -0.368 . This difference is statistically significant and we see that the increase in elasticity is even larger without household fixed effects (from -0.298 to -0.511). The strong significance is because the average oil price is very strongly correlated with the fuel price. However, we do not view it as a fundamentally different estimate.³²

Next, we turn to the car type choice. The car characteristics of a household may be endogenous to the price elasticity of the household in the sense that households that

³²We have considered using an alternative set of IVs instead of the oil price. For example, we could use supply shocks in the US gulf region in terms of hurricanes and their projected effect on oil production. Since these are driven by meteorological forces, they are more plausibly exogenous.

are more financially vulnerable may choose to buy a more fuel efficient car in order to minimize their exposure to the fuel prices. In that sense, they are not responding to the fuel price per se when making their driving decisions but rather to the price per kilometer, defined as the fuel price divided by the fuel efficiency of the car. Such a selection story has been analyzed by many different authors with different modeling approaches and it underlies the discussion of the *rebound effect* (the effect of an exogenous change in fuel efficiency on driving).³³ However, there are three reasons why we do not believe this to be a large concern for the present analysis; firstly, it has been found in the literature that most of the adjustment to higher fuel prices comes from changes in driving and not in car characteristics. Bento et al. (2005b) find that 95% of the final response in gasoline consumption to a change in the gasoline tax in the US comes from a change in driving rather than a change in fleet fuel efficiency. For Denmark, Munk-Nielsen (2015) finds that when fuel prices increase by 1%, average fuel efficiency only increases by 0.1%. Thus, the most important component is understanding the response in driving.

Secondly, we believe that our controls for car characteristics capture the most important aspects of such selection, even though fuel efficiency and car price are not included. As mentioned in section 3, these variables are not available for car vintages older than 1997 but we can still use them to explore robustness to endogenous selection. We discuss this in detail in appendix section C.7. There, we argue that firstly, the included car characteristics account for a substantial portion of the variation in the omitted ones and secondly, that including the omitted characteristics does not change the elasticity on the subsample where they are available.³⁴

Thirdly, the adjustments in the car stock will most likely take place over a longer horizon. In that sense, the work we are doing is relevant for short to medium run policy analysis. For examining long-run implications of fuel taxes, one needs to take into account the interactions with the used car market and scrappage as well. If the policy has effects on the life time of cars, then that can potentially have huge environmental consequences (D’Haultf  uille, Givord, and Boutin, 2014).

8 Conclusion

In this paper, we estimate the fuel price elasticity of driving demand for Denmark. Our preferred specification yields an estimated elasticity of -0.304. Using quantile regression, we document an inverted U-shape in the fuel price elasticity, where the households with the shortest and longest work distances have the highest price sensitivities in the population.

³³See for example West (2004); Bento et al. (2005b, 2009); Gillingham (2013); Munk-Nielsen (2015) and the theoretical treatment by Chan and Gillingham (2015).

³⁴The elasticity is, however, quite different on that subsample compared to the full sample. This is mainly due to an over-sampling of newer cars in the early years. That is why we do not include the variables in the preferred specification.

From interactions with the fuel price variable, we are also able to find the same inverted U-shape in the fuel price elasticity based on the household’s work distance. Furthermore, our results indicate that ability to reduce driving in response to increasing fuel prices is mitigated by the availability of public transportation, which is nearly universally available in Denmark. We develop a theoretical framework that rationalizes the heterogeneity in the fuel price elasticity in a model with switching costs. “Tail households”, with high work distance and high driving, will tend to have high fuel expenditures if they commute to work by car so they stand to gain a lot from switching to public transportation. This group will therefore switch for lower increases in fuel prices for any fixed switching cost. In the lower end of the work distance distribution, the total driving is primarily made up of leisure trips. The finding that these households are more responsive indicates that leisure trips are very price sensitive.

By predicting elasticities at the individual level, we are able to visualize the spatial distribution of the fuel price elasticity on a map. This indicates that many tail households live in rural areas and commute to Copenhagen. Since public transportation is readily available in these areas, this is consistent with our intuition about the switching behavior. The second group of high-elasticity households with low driving and low work distances are highly represented in some of municipalities immediately to the north of Copenhagen, where there is good access to substitutes to transportation by car.

We conduct a simple counterfactual where we increase the fuel price by 1 DKK/liter. Our results indicate that the high work distance tail households, defined as the households in the top 5% of the conditional driving distribution, account for 14.4% of the total response in driving. Other studies have found that the primary adjustment to fuel prices is in driving rather than the fuel efficiency of new cars (Bento et al., 2009; Munk-Nielsen, 2015), so this result implies that the tail households will bear the greater part of the aggregate reduction in emissions in response to a fuel tax increase. We show that the fuel tax of 1 DKK/liter will imply a deadweight loss of 0.556 for the average person in our sample. This seemingly very large estimate can be explained by the high fuel tax level in Denmark. Finally, we show that the heterogeneity in the fuel price elasticity leads to heterogeneity in the deadweight loss of the tax, with the two groups of tail households bearing a four times larger deadweight loss than the households with medium work distances.

References

- Abrevaya, Jason and Christian M Dahl. 2008. “The effects of birth inputs on birthweight: evidence from quantile estimation on panel data.” *Journal of Business & Economic Statistics* 26 (4):379–397.

- Adamou, Adamos, Sofronis Clerides, and Theodoros Zachariadis. 2013. “Welfare Implications of Car Feebates: A Simulation Analysis.” *The Economic Journal* URL <http://dx.doi.org/10.1111/ecoj.12094>.
- Bento, Antonio, Maureen Cropper, Ahmed Mushfiq Mobarak, and Katja Vinha. 2005a. “The Effects of Urban Spatial Structure on Travel Demand in the United States.” *Review of Economics and Statistics* 87(3):466–478.
- Bento, Antonio, Lawrence Goulder, Mark Jacobsen, and Roger von Haefen. 2009. “Distributional and Efficiency Impacts of Increased US Gasoline Taxes.” *American Economic Review* 99(3):667–699.
- Bento, Antonio M., Lawrence H. Goulder, Emeric Henry, Mark R. Jacobsen, and Roger H. von Haefen. 2005b. “Distributional and Efficiency Impacts of Gasoline Taxes: An Econometrically Based Multi-market Study.” *American Economic Review* 95 (2):282–287. URL <http://www.aeaweb.org/articles.php?doi=10.1257/000282805774670536>.
- Blundell, Richard, Joel Horowitz, and Matthias Parey. 2012. “Measuring the Price Responsiveness of Gasoline Demand: Economic Shape Restrictions and Nonparametric Demand Estimation.” *Quantitative Economics* 3 (1):29–51.
- Brons, Martijn, Peter Nijkamp, Eric Pels, and Piet Rietveld. 2008. “A Meta-analysis of the Price Elasticity of Gasoline Demand. A SUR Approach.” *Energy Economics* 30 (5):2105–2122.
- Brownstone, David and Thomas Golob. 2010. “The Impact of Residential Density on Vehicle Usage and Energy Consumption.” *Journal of Urban Economics* 65 (1):91–98.
- Canay, Ivan A. 2011. “A simple approach to quantile regression for panel data.” *Econometrics Journal* 14 (3):368–386.
- Chan, Nathan and Kenneth Gillingham. 2015. “The Microeconomic Theory of the Rebound Effect and its Welfare Implications.” *Journal of the Association of Environmental and Resource Economists* 2 (1):133–159.
- Coglianesi, John, Lucas Davis, Lutz Kilian, and James Stock. 2015. “Anticipation, Tax Avoidance, and the Price Elasticity of Gasoline Demand.” *NBER Working Paper No. 20980* .
- Dahl, Carol and Thomas Sterner. 1991. “Analysing Gasoline Demand Elasticities: A Survey.” *Energy Economics* 13 (3):203–210.
- Davis, Lucas and Lutz Kilian. 2011. “Estimating the Effect of a Gasoline Tax on Carbon Emissions.” *Journal of Applied Econometrics* 26 (7):1187–1214.

- De Borger, B., I. Mulalic, and J. Rouwendal. 2013. "Substitution between cars within the household." *Tinbergen Institute Discussion Paper No. 13-158/VIII* .
- D'Haultfœuille, Xavier, Pauline Givord, and Xavier Boutin. 2014. "The Environmental Effect of Green Taxation: The Case of the French Bonus/Malus." *The Economic Journal* 124 (578):F444–F480. URL <http://dx.doi.org/10.1111/ecoj.12089>.
- DTU Transport. 2013. "Transportvaneundersogelsen." Tech. rep., DTU Transport. URL <http://www.modelcenter.transport.dtu.dk/Transportvaneundersoegelsen/TU-udgivelser>.
- Edelstein, Pual and Lutz Kilian. 2009. "How Sensitive Are Consumer Expenditures to Retail Energy Prices." *Journal of Monetary Economics* 56 (6):766–779.
- Espey, Molly. 1998. "Gasoline Demand Revisited: An International Meta-analysis of Elasticities." *Energy Economics* 20 (3):273–295.
- Frondel, Manuel and Colin Vance. 2013. "Re-Identifying the Rebound: What about Asymmetry?" *Energy Journal* 34 (4):43–54.
- Gillingham, Kenneth. 2013. "Selection on Anticipated Driving and the Consumer Response to Changing Gasoline Prices." *Yale University Working Paper* .
- . 2014. "Identifying the Elasticity of Driving: Evidence from a Gasoline Price Shock." *Regional Science & Urban Economics* 47 (4):13–24.
- Glaeser, Edward and Matthew Kahn. 2010. "The Greenness of Cities: Carbon Dioxide Emissions and Urban Development." *Journal of Urban Economics* 67 (3):404–418.
- Graham, Daniel and Stephen Glaister. 2004. "Road Traffic Demand Elasticity Estimates: A Review." *Transport Reviews* 24 (3):261–274.
- Grazi, Fabio, Jeroen van den Bergh, and Jos van Ommeren. 2008. "An Empirical Analysis of Urban Form, Transport, and Global Warming." *The Energy Journal* 29 (4):97–122.
- Grigolon, Laura, Mathias Reynaert, and Frank Verboven. 2015. "Consumer Valuation of Fuel Costs and the Effectiveness of Tax Policy: Evidence from the European Car Market." *Working Paper* .
- Hamilton, James. 2009. "Understanding Crude Oil Prices." *Energy Journal* 30 (2):179–206.
- Hughes, Jonathan, Christopher Knittel, and Daniel Sperling. 2008. "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand." *Energy Journal* 29 (1):93–114.

- Huse, Cristian and Claudio Lucinda. 2013. "The Market Impact and the Cost of Environmental Policy: Evidence from the Swedish Green Car Rebate." *The Economic Journal* URL <http://dx.doi.org/10.1111/eoj.12060>.
- Hymel, Kent M. and Kenneth A. Small. 2015. "The Rebound Effect for Automobile Travel: Asymmetric Response to Price Changes and Novel Features of the 2000s." *Energy Economics* forthcoming.
- Jacobsen, Mark. 2013. "Evaluating U.S. Fuel Economy Standards in a Model with Producer and Household Heterogeneity." *American Economic Journal: Economic Policy* 5(2):148–187.
- Kilian, Lutz and Daniel Murphy. 2014. "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil." *Journal of Applied Econometrics* 29 (3):454–578.
- Knittel, Christopher and Ryan Sandler. 2013. "The Welfare Impact of Indirect Pigouvian Taxation: Evidence from Transportation." *National Bureau of Economic Research Working Paper No. 18849* .
- Knittel, Christopher R. 2011. "Automobiles on Steroids: Product Attribute Trade-offs and Technological Progress in the Automobile Sector." *The American Economic Review* 101 (7):3368–3399.
- Koenker, Roger. 2004. "Quantile regression for longitudinal data." *Journal of Multivariate Analysis* 91 (1):74–89.
- . 2005. *Quantile regression*. Cambridge University Press.
- Levin, Lewis, Matthew Lewis, and Frank Wolak. 2014. "High-Frequency Evidence on the Demand for Gasoline." *mimeo, Stanford University* .
- Li, Shanjun, Joshua Linn, and Erich Muehlegger. 2014. "Gasoline Taxes and Consumer Behavior." *American Economic Journal: Economic Policy* 6 (4):302–342.
- Linn, Joshua. 2013. "The Rebound Effect for Passenger Vehicles." *RFF Working Paper* .
- Machado, José AF and José Mata. 2005. "Counterfactual decomposition of changes in wage distributions using quantile regression." *Journal of applied Econometrics* 20 (4):445–465.
- Mayeres, Inge and Stef Proost. 2013. "The taxation of diesel cars in Belgium - revisited." *Energy Policy* 54 (0):33–41. URL <http://www.sciencedirect.com/science/article/pii/S0301421511009670>.

- McFadden, Daniel. 1974. “The Measurement of Urban Travel Demand.” *Journal of Public Economics* 3 (4):97–122.
- Melly, Blaise. 2005. “Decomposition of differences in distribution using quantile regression.” *Labour Economics* 12 (4):577–590.
- Munk-Nielsen, Anders. 2015. “Diesel Cars and Environmental Policy.” *Working Paper* .
- Portnoy, Stephen. 1991. “Asymptotic behavior of the number of regression quantile breakpoints.” *SIAM journal on scientific and statistical computing* 12 (4):867–883.
- Robinson, Peter M. 1988. “Root-N-consistent semiparametric regression.” *Econometrica* :931–954.
- Small, Kenneth A. and Kurt van Dender. 2007. “Fuel Efficiency and Motor Vehicle Travel: The Declining Rebound Effect.” *Energy Journal* 28 (1):25–51.
- West, Sarah. 2004. “Distributional Effects of Alternative Vehicle Pollution Control Technologies.” *Journal of Public Economics* 88 (3-4):735–757.

A Data Details

In this part of the appendix, we describe the data in more detail and elaborate on the sample selection.

A.1 Cleaning of the Raw Data

We have generally focused on avoiding dropping observations and instead including controls to maximize our dataset coverage.

- After cleaning the data, we have 10,994,333 combinations of households and driving periods.³⁵
- Of these, we further make restrictions on the period, which leaves us with 7,254,893. The requirements are that the start of the driving period should be:

³⁵Note that a driving period may occur twice, if for example the car gets sold midway through; once with each household that contributed to the driving in that period. One of the most important sample selection criteria in the raw data cleaning was deleting cars with negative driving. We have experimented extensively with potential remedies for this problem but decided that it was too complicated to bring in these final observations. The issue is that we do not know if the negative driving for a period $[t_0; t_1]$ occurred because the odometer observation at period t_0 was incorrect and too large or the measure at t_1 was incorrect and too small.

- after July 1st, 1998: this is because tests were not mandatory in 1997 and the first half of 1998. Moreover, while we can impute driving going back earlier into the 90s by the first time the cars appear for inspection, these periods are typically very long and different from the later part of the sample. See figure A.4.
- before January 1st 2008: the last inspection occurs in September 2011, so allowing for tests periods starting later will systematically deselect new cars.
- After this, we delete observations where the length of the driving period, which we call *years to test*, is not either between 1 and 2.5 years or between 3.5 and 4.5 years. This is chosen to balance not getting too many observations with unexplained lengths of the driving periods while also accounting for the phase-in early, which lead to in particular a number of 1-year periods in 1999 and 2000. See section A.2.1. This leaves us with 6,877,185 driving periods.
- We have missing demographic variables for 277,294 rows, which brings us to 6,599,891 rows.
- Then, we delete 178 observations having VKT greater than 10,000 km per day, which we delete.³⁶
- Now, we impose the restriction that households be observed at least twice since our main specification will use household fixed effects. There are 744,267 households that we only observe once and thus delete (see figure A.2). This brings us to our final sample size of

$$N = 5,855,446.$$

Table A.1 shows a histogram of the start year of the driving period. The low number in the first year is due to the sample selection criterion keeping only periods starting after July 1st, 1998.


A.1.1 Car Ownership

The information about car ownership comes from the Central Motor Register. This register basically contains the license plate, vehicle identification number (VIN; uniquely identifying the vehicle across potentially different license plates and/or owners) and personal identification numbers (CPR numbers, allowing us to merge with other public registers). The ownership period is matched to the daily level.

In the raw data, we observe some problematic observations. When we observe a car ownership period without an ending but we observe another person owning the car from

³⁶The average of the 178 VKT observations is 308,364.7 km/day.

Figure A.1: Observations by start year



a later date, we assign that date as the ending of the ownership period for the first owner. Similarly for the reverse scenario where we have a prior complete ownership row but a later one where we observe the ending but not the start in which case we assign the former owner's end date as the start date.

We also see more problematic rows where there is an overlap of owners. In that case, we have no way of discerning which person truly owns the car and according to the data documentation such an observation should be impossible so we drop them from the dataset.

A.2 Household information


We are forced to drop households that we only observe once in the data due to the fixed effects specifications. We could keep such households for regressions without fixed effects and for certain graphs but we choose to give priority to having the same sample throughout the paper.

A.2.1 Driving Periods

The data on driving periods come from the Ministry of Transportation (MOT) tests that were introduced in 1997. These inspections are mandatory and must be performed at car ages 4, 6, 8, 10, 12, etc. This means that we have two types of driving periods; The *first* driving period is 4 years long (that is, it has *4 years to test*) and any *subsequent* driving period will be only 2 years long. The inspection date is set based on the date of the first registration of the car in Denmark.

MOT tests were originally performed by public authorities directly but in more recent years, this is done by private companies approved by the MOT. At the test, it is verified

Figure A.2: Number of Driving Periods per Household




whether the car is in safe condition for driving on the roads. As a part of the test, the odometer of the car is measured, giving us the km that the car has driven since last test or since purchase. A test may have four outcomes; 1) The car can be *approved*. 2) The car can be *conditionally approved*, meaning that certain repairs must be performed for the car to be in legal driving order but that no extra test will be required. 3) The car can be *approved after a re-inspection*, implying that repairs must be made and then the car must return for another test before 33 calendar days. Finally, 4) the car can be declared *not approved* in which case it will be illegal to drive the car and the police will withdraw the license plates.

In practice, the years to test may deviate with plus or minus three months, which is also where we see most of the observations. A person may choose to take the car in for inspection *earlier* than the set date if he wishes. Anecdotally, some people appear to choose to do this prior to selling the used car in order to give the buyer a signal that the car is in proper working order.

Figure A.3 shows the distribution of the driving period length. The vertical lines mark the sample selection. We have chosen to use only driving periods satisfying that either the driving period be between 1 year and 2.5 years or between 3.5 years and 4.5 years. The reason for going down to 1 year is that during the phase-in of the inspections in the earlier years, there were some irregularities in the test lengths so cutting at 1.5 would discard disproportionately more of the driving periods in the earlier years of our sample.

The first MOT tests were conducted in 1997. However, from January 1, 1997 to July 1, 1998, the tests were being phased in and moreover, it wasn't yet mandatory for all cars to come to the inspection. As shown in figure A.4, this implies that most driving periods starting in 1997 through mid-1998 were 4-year tests. From mid-1998, however, the tests were mandatory for all cars and we see in figure A.4 that the (smoothed) mean

Figure A.3: Years to Test Distribution



years to test drops down to its stable level just above 2. After 2008, we see that the mean years to test drops quickly down to 2. This is because the last observed MOT tests are on December 30, 2011, meaning that none of the new cars bought from January 1, 2008, are due for inspection before after our sample ends. Including all information from tests would mean that we would have zero new cars for the additionally gained observations in 2008 and 2009.

Note that the fluctuations seen in the graph reflect the patterns in new car purchases since a new car purchase will add another driving period of length 4 years to that point in time.

A.3 Variables Used

In this section, we explain all the variables used in the paper. Table A.1 provides details on the variables included in the regressions and table A.2 provides a list of all the variables used in this paper.

Table A.2: Variables used in the paper

Variable	Description
----------	-------------

VKT	Vehicle Kilometers Travelled in km per day. The variable is constructed by taking first-differences of the odometer readings from the dataset with vehicle inspections. For the first inspection we observe for a car, we assume that the odometer was zero at the time of the car's first registration in Denmark. This will be incorrect if the car was imported from abroad. However, then the car must have had a toll inspection, which we observe; we find that this does not impact our results.
Couple	Dummy for there being two members of the household (married or co-habiting, of opposite genders and having at less than 15 years of age difference).
Real gross income	The sum of gross incomes for the member(s) of the household. The variable comes from the income tax registers. The variable includes all government transfers such as pension payments, unemployment benefits, etc.
Real gross income (couples)	As above but equal to zero for singles.
Real gross income (singles)	As above but equal to zero for couples.

WD

Work distance. The variable is based on the Danish deduction for work distance. Any working household having further than 12 km each way to work can deduct a fixed amount per km. Thus, the measure will be equal to zero if the individual lives closer than 12 km from his or her work. Between 12 and 25 km, there is a rate and above 25 km, the rate drops to half. The rate changes over the period. The total deduction is the daily rate times the number of days worked. The variable is self-reported but the tax authorities have access to both the home and work addresses for the individual. The deduction is the rate times the distance times the number of days worked. We do not observe the number of days worked so we assume 225 work days, which corresponds to the number of days in a typical Danish year.³⁷ Figure A.5 shows the distribution of the work distance variable, both the full distribution including the mass point at a work distance of zero and the uncensored distribution, where the work distance is not censored. We clearly see the effects of the censoring. There is also positive mass on the interval (0; 12) km even though the deduction is only given if the actual work distance is above 12 km; this is due to the assumption about 220 work days per year. If an individual works part time, say 110 days, but has a distance of 20 km to work, then the variable will be equal to 10. The positive mass will therefore contain many part time employees. For validity, we can compare it to the continuous WD measure, available for a subset of the period (see Appendix A.4.4).

WD non-zero

Dummy for the WD measure being observed. Thus, this is essentially a dummy for the individual living further than 12 km from the work place.

WD (actual distance)

This is the actual distance from home to work. The variable comes from the Danish Technical University's Department of Transportation. It is calculated using a shortest-path algorithm and the National Transport model with GIS data on households and their work places. The variable is only observed for households where the work place is observed and not for 1998 or 1999. In total, it is observed for 76.17% of our estimation sample (79.61% of the observations between 2000 and 2008). We use this measure to validate the tax-based WD variable.

³⁷ For example, the official number of work days were; 224 in 2007, 226 in 2008, 225 in 2009 and 228 in 2010. Most unions follows these and most public sector employees.

# of kids	The number of kids aged less than 18 years living with the household.
Urban (dummy)	Dummy equal to one if the household lives in either Copenhagen, Frederiksberg, Aarhus, Aalborg og Odense municipalities, which constitute the major Danish urban areas.
Company car	Dummy equal to one if at least one member of the household has paid the tax penalty for having access to a company car. The use of company cars is restricted to avoid making it an alternative to buying your own car privately. The size of the tax depends on the value of the car. We collapse the variable to a dummy for having any car available to any of the members of the household. Individuals may have access to a company car and not pay this tax if the car is a “yellow license plate” car. These cars can have at most two seats and are typically vans used by craftsmen. The police enforces this very strictly and an individual caught using a company car privately and not paying the penalty is fined and some times forced to pay the registration tax.
Self employed	Dummy equal to one if the household has at least one self employed individual. This information comes from the tax registers.
# of periods observed	The number of driving periods observed for the household. Note that the other driving periods may be with different cars and that our sample selects only households with at least two driving periods.
Bus stops per km ²	The number of bus stops in the municipality in 2013 divided by the area of the municipality of residence at the start of the driving period in km ² . The data for this comes from the Travel Planner (http://rejseplanen.dk), which is a search engine for planning trips using public transportation. The data is unfortunately not available back in time so we use the 2013 data. Thus, the variable does not have time-series variation. The highest number of stops is 79.9 stops per km ² for Odense municipality and the lowest is Aaskov municipality with 0.3 stops per km ² .
Weight (ton)	The gross weight of the car in metric tonnes. This is the maximum allowed weight of the vehicle including cargo. The variable comes from the vehicle type approval documents.
Diesel	Dummy equal to one if the car uses diesel fuel. Note that the fuel price will then be based on the diesel price.

Van	Dummy equal to one if the vehicle is a van.
Percent owned of period	The fraction of the driving period where the car was owned by this household. That is, if the driving period starts on Jan 1st, 2001 and ends on Jan 1st 2003, but the car changed owner on Jan 1st 2002, this variable will be equal to 0.5 for both the observations of the two households driving the car.
Driving period length	The length of the driving period in years. For new cars, this will be 4 years and for older cars, it will be 2 years, both plus or minus 3 months and with some exceptions. Note that our sample selects on driving periods being either 1.0 to 2.5 years long or 3.5 to 4.5 years long.
Car age (ultimo)	Car age in years at the start of the driving period. Car age is defined as the time since the car's first registration in Denmark since we do not observe the actual production year of the. Thus, it is not actually the car age and it will be particularly far off for imported cars or for cars than have been imported by a dealer but only sold after a long time.
# cars / vans / motorcycles / mopeds / trailers owned	Continuous measure of the number of vehicles of the given type owned by the household. For example, if for a given household i and driving period t , the household owns another car for the entire duration of the period, then # of cars owned will be 2.0. If that other car is only purchased half-way through the driving period t , then it is equal to 1.5. That is, the variable is equal to the fraction of this driving period overlapping with the ownership of other vehicles.
First driving period	Dummy equal to one if it is the car's first driving period, i.e. the driving period's start date is equal to the first registration date of the car.
Fraction owned	For household i and driving period t , this is the percent of the driving period where household i is the owner. That is, if the car changes owner midway through, there will be an observation in the dataset for each of the two households owning the car and they will both have this variable set to 0.5.
Years to test	The length of the driving period in years (continuous variable). Due to our sample selection, this will be in $[1.0; 2.5]$ or in $[3.5; 4.5]$.

% of each month	This is a set of variables for each month equal to the % of the driving period taking place in each of the 12 months. Thus, if a driving period is precisely 2 or 4 years long, these will all be equal to $\frac{1}{12}$. We omit April as the reference group in regressions since the fractions will always sum to 1.
Year controls	These are variables for each year, 1998, ..., 2011, each equal to the % of the driving period falling in that year. In the preferred specification, we exclude year 2003 as the reference year and include an additional full set of year controls interacted with the diesel dummy to allow a separate time trend for diesels.
Period	In the regressions, this covers the “Percent owned of period” and the “Driving period length” variables.

A.4 Descriptives

A.4.1 Fuel Prices Over Time

Figure A.6 shows the real fuel price series for octane 95 (gasoline) and diesel over the sample period. The two most apparent features are the sharp price increase from 1999 to 2000 and the apparent convergence of the two price series over time.

Figure A.7 shows the composition of fuel price for gasoline and diesel in nominal prices. Gasoline has higher fixed taxes (the “Energy Tax”) throughout the period. To correct for this, the biannual ownership tax is higher for diesel cars. Over the period from July 1st 19998 to January 1st 2012, taxes make up 64.87% of the gasoline price and 55.24% of the diesel price.

A.4.2 Driving and Demographics

Figure A.8 shows the distribution of vehicles kilometers travelled (VKT) (cut at 200km/day). Note that there is still positive mass for very low VKT. This may for example be explained by vintage or specialty cars.

Figure A.9 shows the nonparametric relationship between VKT and four demographic variables: car age, income, work distance and household age. For car age, we show average driving by the age of the car at the start of the driving period. We use car age 0, 4, 6, 8, ... which are where the data density is highest due to the inspections being administered at these car ages. The relationship is monotonically decreasing from just over 55 km/day for new cars to about 35 km/day for 20 year old cars. For the income graph, we have divided the income distribution into 20 quantiles and plotted the average VKT within each quantile. The graph shows an inverted U-shape: driving increases from

Figure A.4: Years to Test by Start Date of the Driving Period




Table A.1: Regression Variables

Demographics	Log real income, log real income (singles), urban (dummy), WD (male), WD (female), WD (single), WD non-zero (male), WD non-zero (female), WD non-zero (single), age (male), age (female), age (single), age squared (male), age squared (female), age squared (single), number of kids, company car (dummy), self-employed (dummy), number of bus stops.
Car	Weight, Weight squared, diesel (dummy), van (dummy), car age (ultimo), number of cars owned, number of vans owned, number of motorcycles owned, number of mopeds owned, number of trailers owned.
Period	First driving period, years to test, percent owned of period.
Year controls	A full set of year controls (omitting 2003 as the baseline).
% of each month	A full set of month controls (omitting April as the baseline).

Figure A.5: The Distribution of Work Distance: Full Sample and Uncensored Subsample




Figure A.6: Real Fuel Prices




Figure A.7: Fuel Price Composition in the Sample Period






Figure A.8: The Distribution of Vehicle Kilometers Traveled



around 33 km/day at the lowest income quantiles to a top point of approximately 52 km/day for incomes around 800,000 2005 DKK. After this, VKT decreases in income. The work distance graph is constructed by dividing households into 40 quantiles based on work distance. More than half of these occur at a work distance of zero because our measure is censored at 12 km. The graph shows that average driving of the households with shorter than 12 km to their work place is 40 km/day. As work distance increases, driving increases monotonically up to about 70 km/day for work distances of just under 90 km. The final graph shows the relationship between driving and household age (the maximum within the household is used for couples). The average driving is somewhat stable at just over 50 km/day until around age 50, after which driving starts to decline rapidly down to an average of about 25 km/day for 80 year old households.

Figure A.9: Driving and Demographics: Nonparametric Relationships



The dots are placed according to equi-distant percentiles of the conditioning variables and the corresponding y-value represents the average VKT within that percentile-group. Finally, the dots are connected by linear line segments for illustration purposes.

Figure A.10: Bus and Train Stops in Denmark



A.4.3 Spatial Descriptives

Figure A.10 shows a map of Denmark where each dot represents a train or bus stop. The map comes from the Journey Planner (rejseplanen.dk), which is an online search engine for planning trips using public transportation.

Figure A.11 shows the number of observations by municipality. The four major urban areas clearly stand out: Copenhagen (east), Odense (center, on the island of Fyn), Aarhus (midway up on the eastern side of Jutland) and Aalborg (Northern part of Jutland).

Figure A.12 shows a map of Denmark where municipalities are colored by the average work distance of the households. We see that the households with high work distances tend to be in the major urban areas with a few exceptions. Note that we are doing this for the estimation sample; in particular, this means that we are conditioning on the households owning a car. As with the VKT map in Figure 3.3, this may help explain the perhaps surprising fact that the urban residents tend to have larger work distances than the rural ones. Instead, the interpretation should be that the households in the urban areas that choose to own a car tend to have longer work distances than their in the rural areas.

To focus on the tail households, we also show the 95th percentile of driving and work distance within municipalities. These are shown in Figures A.13 and A.14 respectively. The graphs show that the patterns from the corresponding graphs of the averages are mirrored in the 95th percentiles; the urban areas again stand out with high driving and high work distances.³⁸

³⁸Also, the municipality at the center of Sealand, Ringsted, stands out. This municipality has a major inter-regional train line crossing through and thus offers a good opportunity for households working in opposite ends of the country.

Figure A.11: Observations in the Estimation Sample by Municipality




Figure A.12: Average Work Distance by Municipality




Figure A.13: The 95th Percentiles of Driving Within Municipality





Figure A.14: The 95th Percentiles of Work Distance Within Municipality



A.4.4 Work Distance

In this subsection, we consider the validity of the work distance variable. The distribution of this variable is shown in Figure A.5, which shows the censoring of the variable. Table A.3 shows summary statistics for work distances of males, females and singles. It shows both the measure based on the tax deduction for work distance as well as the “actual work distance” variable, which measures the distance using GPS coordinates. The tax deduction is a deduction from taxable income and it is given as a fixed amount per kilometer per day but is equal to zero if the distance is shorter than 12 km. The number of days worked is not observed so we assume that all individuals work 220 days a year, which is very common in Denmark. Hence, if the individual actually worked fewer days, we will be undershooting the measure (which explains why the variable can take values below 12 km) and vice versa. The per kilometer rate varies over time and there is a kink in the schedule at 50km where it falls to half the rate.³⁹

Table A.3: Work Distance Variables


	count	mean	sd	p1	p10	p25	p50	p75	p90	p95	p99
WD, male	4550411	9.5932	18.51	0.0	0.0	0.0	0.0	15.5	32.2	44.7	80.7
WD, female	4550411	6.9385	13.77	0.0	0.0	0.0	0.0	10.7	24.8	33.7	58.3
WD, single	1305035	7.6966	16.87	0.0	0.0	0.0	0.0	9.0	27.5	39.3	75.1
WD non-zero, male	4550411	0.3493	0.48	0	0	0	0	1	1	1	1
WD non-zero, female	4550411	0.3137	0.46	0	0	0	0	1	1	1	1
WD non-zero, single	1305035	0.2917	0.45	0	0	0	0	1	1	1	1
Actual WD, male	3343884	20.3157	34.36	0.0	0.0	2.7	9.8	23.7	46.5	71.8	196.3
Actual WD, female	3094025	14.3657	22.45	0.0	0.6	2.8	8.1	18.1	32.1	45.0	99.3
Actual WD, single	813453	18.6009	32.87	0.0	0.1	2.6	8.6	21.1	42.0	66.1	183.4

To explore the validity of the work distance variable, we exploit the aforementioned *actual* work distance. We call this the actual work distance because our primary work distance variable is a daily rate so in that sense it includes a measure of the number of days worked in the year. In that sense, it is a better variable for measuring commuting and driving than simply the work distance. However, we can compare the distribution of driving according to the two variables to validate the measure. To make the comparison sensible, we do it for the subsample where both measures fall in the range [12km ; 100km] — this lower bound ensures that the tax-based measure is also observed, while the upper is to make the graph easier to read.

Table A.4 shows the frequency table for the two dummies for zero work distance for male and female respectively for the subsample of couples. We see a clear positive correlation between the spouses’ work distances. The coefficient of correlation between the

³⁹In some years, a small number of *fringe municipalities* (Danish: *udkantskommuner*) also had the full rate after the 50km threshold.

Figure A.15: Comparing the Two Work Distance Measures



two binary variables is 0.295 (if we only use one observation per couple, leaving 1,130,487 unique households, the correlation is 0.292). If we regress the work distance of the male on that of the female, the R^2 is 4.58% (4.53% with one observation per household) — so there is substantial variation within households.

Table A.4: Work Distance Within-Household

	WD _f = 0	WD _f > 0	Total
WD _m = 0	2,329,041	631,740	2,960,781
WD _m > 0	794,008	795,621	1,589,629
Total	3,123,049	1,427,361	4,550,410

B Other Results

This appendix contains two subsections; first, a number of econometric results supplementing the primary results from section 5. Second, welfare calculations to supplement the counterfactual analyses in section 6.

B.1 Supplementary Econometric Results

Table B.1 shows the coefficients pertaining to car characteristics and the driving period that were suppressed in table 5.1.

Table B.2 shows the fuel price elasticities used in figure 5.1.

Table B.3 shows the coefficients for the demographic variables for the quantiles 1, 50 and 99 in the panel quantile regression estimates. They show that many of the coefficients

Table B.1: Main results — Car and Period Controls

	OLS (1) No demo	(2) Base	Household FE (3) FE	(4) Main
$\log p^{\text{fuel}}$	-0.866*** (0.00509)	-0.298*** (0.0143)	-0.515*** (0.00722)	-0.304*** (0.0154)
New car	-0.00350* (0.00148)	0.0128*** (0.00148)	0.00838*** (0.00160)	0.0394*** (0.00164)
Percent owned of period	-0.189*** (0.000826)	-0.112*** (0.000862)	-0.0537*** (0.00106)	-0.0154*** (0.00110)
Driving period length	-0.0507*** (0.000634)	-0.0541*** (0.000645)	-0.0465*** (0.000681)	-0.0242*** (0.000725)
Weight (ton)	0.00214*** (0.00000523)	0.00169*** (0.00000506)	0.00166*** (0.00000799)	0.00167*** (0.00000798)
Weight squared	-0.000000471*** (1.35e-09)	-0.000000369*** (1.30e-09)	-0.000000354*** (2.00e-09)	-0.000000354*** (2.00e-09)
Diesel	0.316*** (0.000918)	0.311*** (0.00557)	0.228*** (0.00139)	0.259*** (0.00545)
Van	-0.236*** (0.00117)	-0.199*** (0.00115)	-0.204*** (0.00171)	-0.205*** (0.00170)
Car age (ultimo)	-0.0302*** (0.0000932)	-0.0275*** (0.0000911)	-0.0284*** (0.000140)	-0.0293*** (0.000141)
# cars owned	0.0482*** (0.000593)	-0.0202*** (0.000759)	-0.0581*** (0.00114)	-0.0501*** (0.00109)
# vans owned	0.0111*** (0.00124)	-0.0470*** (0.00122)	-0.0711*** (0.00183)	-0.0654*** (0.00179)
# motorcycles owned	0.0319*** (0.00101)	-0.00420*** (0.000905)	0.0102*** (0.00178)	0.0118*** (0.00178)
# mopeds owned	0.136*** (0.00138)	0.0415*** (0.00131)	0.0232*** (0.00218)	0.0204*** (0.00217)
# trailers owned	0.0123*** (0.000519)	0.0258*** (0.000983)	0.00334** (0.00106)	0.00595*** (0.00106)
Demographics	No	Yes	Yes	Yes
Year controls	No	Yes	No	Yes
% of each month	Yes	Yes	Yes	Yes
Household FE	No	No	Yes	Yes
N	5855446	5855446	5855446	5855446
R^2	0.198	0.340	0.175	0.180

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B.2: Fuel Price Elasticity by Conditional Quantile

Quantile	$\log p^{\text{fuel}}$	se
99	-0.609***	(0.0592)
95	-0.369***	(0.0223)
90	-0.289***	(0.0145)
85	-0.256***	(0.0114)
80	-0.238***	(0.00988)
75	-0.234***	(0.00890)
70	-0.228***	(0.00831)
65	-0.230***	(0.00796)
60	-0.231***	(0.00758)
55	-0.233***	(0.00742)
50	-0.233***	(0.00739)
45	-0.244***	(0.00755)
40	-0.250***	(0.00779)
35	-0.267***	(0.00808)
30	-0.284***	(0.00868)
25	-0.309***	(0.00942)
20	-0.330***	(0.0106)
15	-0.357***	(0.0124)
10	-0.402***	(0.0156)
5	-0.490***	(0.0233)
1	-0.559***	(0.0663)
For all quantile regressions:		
Demographics		Yes
Year controls		Yes
% of each month		Yes
Car		Yes
Period		Yes
Household FE (Canay, 2011)		Yes
N		5855446

do not vary over the conditional distribution of VKT. However, in particular the fuel price elasticity, work distance, company car and bus stops variables change.

For the results with interactions, not all the estimated coefficients were shown in Table 5.2 in the main results section. Table B.4 shows some of the omitted coefficients, namely the ones pertaining to car characteristics.

The interaction estimates may mask some interesting correlations; for example, the income term is negative and normally, we would expect high-income households to be less price-sensitive. To get a fuller picture, we instead want to look at the picture when all interaction effects are taken into account jointly. To do this, we compute the fuel price elasticity, $\gamma_{it} = \gamma \mathbf{x}_{it}^1$, for all observations. Then we split the sample in two based on whether γ_{it} is above or below the 5th percentile of the unconditional distribution of γ_{it} (the bottom 5% being the *most* responsive). Note that the percentiles of γ_{it} are different from the conditional quantiles of VKT that are on the x-axis of figure 5.1; they should instead be thought of as being on the y-axis of figure 5.1. Table B.5 shows summary statistics of household i with driving period t , based on where it falls in the distribution of γ_{it} .

Figure 5.3 shows the spatial pattern of the upper tail in the elasticity distribution. It shows for each municipality, the 10th percentile of the distribution of the elasticities, $\hat{\gamma}_{it}$, i.e. the highest decile in fuel price responsiveness in absolute terms. The pattern is remarkably similar to that shown in figure B.1, which shows the average elasticity within municipality rather than the 10th percentile.

In Figure B.2, we show the percent of the observations belonging to the “tail”, defined as the first decile in the distribution of $\hat{\gamma}_{it}$, i.e. the most responsive. We see that the largest frequency occurs in an equi-distant ring around Copenhagen and in higher frequency at the bottom of Sealand. The map also shows that some highly-elastic households are found to the immediate North of Copenhagen.

B.2 Supplementary Welfare Results

Figure B.6 shows the predicted changes in VKT by quantile from increasing the fuel price from 9 to 10 DKK/liter. The first column, labeled VKT_{it} , shows the quantiles of VKT in the data. The second column, $\text{VKT}_q(p_0)$ shows the predicted quantiles at $p_0 = 9$ using the panel quantile estimates. In the third column, we report the predicted change, $\Delta \text{VKT}_q \equiv \text{VKT}_q(p_0) - \text{VKT}_q(p_1)$, using each of the quantile estimates. In the final column, we report the share of the total response in driving attributable to each quantile. In other words, we have constructed the aggregated predicted response in driving as the weighted average, $\sum_q \Delta \text{VKT}_q$, where the weights correspond to each quantile’s share of the total population, that is $w_q = 0.01$ for $q = 0.01, 0.99$, $w_q = 0.04$ for $q = .05, .95$ and $w_q = .05$ otherwise.

Table B.3: Panel Quantile Regression for P01, P50 and P99: Demographics

	(1) Linear	(2) P01	(3) P50	(4) P99
$\log p^{\text{fuel}}$	-0.304*** (0.0154)	-0.559*** (0.0663)	-0.233*** (0.00739)	-0.609*** (0.0592)
<i>Work Distance (WD) controls</i>				
WD, male	0.00242*** (0.0000336)	0.00137*** (0.000106)	0.00260*** (0.0000118)	0.00347*** (0.0000947)
WD non-zero, male	0.0329*** (0.00107)	0.0797*** (0.00427)	0.0320*** (0.000476)	-0.00688 (0.00382)
WD, female	0.00303*** (0.0000443)	0.00245*** (0.000151)	0.00328*** (0.0000168)	0.00338*** (0.000135)
WD non-zero, female	0.0257*** (0.00111)	0.0950*** (0.00461)	0.0247*** (0.000513)	-0.0327*** (0.00412)
WD, single	0.00419*** (0.0000835)	0.00360*** (0.000216)	0.00448*** (0.0000241)	0.00562*** (0.000193)
WD non-zero, single	0.0724*** (0.00243)	0.138*** (0.00832)	0.0713*** (0.000927)	-0.0174* (0.00743)
<i>Age controls</i>				
Age, male	0.0212** (0.00813)	0.0224*** (0.00133)	0.0213*** (0.000148)	0.0199*** (0.00119)
Age, female	0.0468*** (0.00813)	0.0534*** (0.00132)	0.0469*** (0.000148)	0.0403*** (0.00118)
Age, single	0.0598*** (0.000939)	0.0631*** (0.000971)	0.0604*** (0.000108)	0.0549*** (0.000868)
Age squared, male	-0.0000930*** (0.0000112)	-0.000118*** (0.0000128)	-0.0000943*** (0.00000143)	-0.0000705*** (0.0000115)
Age squared, female	-0.000195*** (0.0000115)	-0.000275*** (0.0000134)	-0.000197*** (0.00000149)	-0.000117*** (0.0000120)
Age squared, single	-0.000206*** (0.00000767)	-0.000275*** (0.00000933)	-0.000213*** (0.00000104)	-0.000119*** (0.00000834)
<i>Other demographic controls</i>				
log gross inc (couple)	-0.0242*** (0.00162)	-0.0156*** (0.00304)	-0.0176*** (0.000339)	-0.0278*** (0.00272)
log gross inc (single)	0.0200*** (0.00288)	0.0268*** (0.00226)	0.0195*** (0.000252)	0.0144*** (0.00202)
Urban (dummy)	-0.0249*** (0.00284)	-0.0392*** (0.00519)	-0.0254*** (0.000578)	-0.0146** (0.00463)
# of kids	-0.0168*** (0.000650)	-0.0145*** (0.00146)	-0.0169*** (0.000163)	-0.0145*** (0.00130)
Company car	-0.0977*** (0.00216)	-0.312*** (0.00695)	-0.102*** (0.000775)	0.0601*** (0.00621)
Self employed	0.000712 (0.00136)	-0.0818*** (0.00426)	0.00334*** (0.000474)	0.0694*** (0.00380)
Bus stops per km ²	0.0000419 (0.0000548)	-0.0000421 (0.000103)	0.0000173 (0.0000114)	0.000300** (0.0000916)
Year controls	Yes	Yes	Yes	Yes
% of each month	Yes	Yes	Yes	Yes
Car	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Linear Fixed Effects (FE)	Yes	No	No	No
Canay (2011) FE	No	Yes	Yes	Yes
N	5855446	5855446	5855446	5855446

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B.4: Heterogenous Elasticity — Car Variables

	(1)	(2)	(3)
Mean elasticity	-0.253	-0.288	-0.238
$\log p^{\text{fuel}}$	-0.879*** (0.240)	-3.847*** (0.0862)	-4.698*** (0.283)
Weight (ton) $\times \log p^{\text{fuel}}$		0.00316*** (0.0000780)	0.00299*** (0.0000820)
Weight squared $\times \log p^{\text{fuel}}$		-0.000000633*** (1.90e-08)	-0.000000603*** (1.98e-08)
Diesel=1 $\times \log p^{\text{fuel}}$		-0.389*** (0.0321)	-0.439*** (0.0325)
Van=1 $\times \log p^{\text{fuel}}$		0.346*** (0.0185)	0.413*** (0.0200)
Car age (ultimo) $\times \log p^{\text{fuel}}$		0.0318*** (0.00104)	0.0295*** (0.00111)
# cars owned $\times \log p^{\text{fuel}}$		-0.232*** (0.0339)	-0.238*** (0.0378)
# vans owned $\times \log p^{\text{fuel}}$		-0.189*** (0.0487)	-0.177*** (0.0466)
# motorcycles owned $\times \log p^{\text{fuel}}$		-0.0404* (0.0164)	-0.0448** (0.0164)
# mopeds owned $\times \log p^{\text{fuel}}$		0.145*** (0.0227)	0.0768*** (0.0227)
# trailers owned $\times \log p^{\text{fuel}}$		0.0497*** (0.0135)	0.0399*** (0.0120)
Demo x p	Yes	No	Yes
Year controls	Yes	Yes	Yes
Month controls	Yes	Yes	Yes
Car	Yes	Yes	Yes
Period	Yes	Yes	Yes
Demographics	Yes	Yes	Yes
Household FE	Yes	Yes	Yes
N	5855446	5855446	5855446
R^2	0.182	0.182	0.185

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B.5: Summary Statistics by Tail Status

	Non-tail $\gamma_{it} \geq P05$	Tail $\gamma_{it} < P05$
Real gross income	566905.5 (554309.8)	709912.8 (1424203.3)
Real gross income (couples)	639735.8 (555530.1)	762801.5 (1350496.5)
Real gross income (singles)	320099.9 (473056.0)	350523.6 (1808309.5)
Couple	0.772 (0.419)	0.872 (0.334)
Age (oldest member)	49.74 (14.32)	51.62 (15.68)
WD non-zero	0.442 (0.497)	0.483 (0.500)
WD	11.41 (17.63)	27.03 (40.56)
WD (actual distance)	22.66 (34.65)	37.86 (47.08)
# of kids	0.783 (1.034)	0.346 (0.677)
Urban (dummy)	0.158 (0.362)	0.173 (0.375)
Company car	0.0290 (0.168)	0.125 (0.331)
Self employed	0.0934 (0.291)	0.188 (0.390)
# of periods observed	4.590 (2.359)	5.801 (6.419)
Bus stops per km ²	15.81 (18.37)	16.81 (19.30)
VKT	46.23 (39.18)	53.40 (55.33)
Weight (ton)	1682.2 (324.8)	1464.6 (368.0)
Diesel	0.132 (0.338)	0.360 (0.480)
Van	0.0816 (0.274)	0.0156 (0.124)
Percent owned of period	0.797 (0.296)	0.720 (0.334)
Driving period length	2.369 (0.875)	2.762 (1.068)
Car age (ultimo)	7.094 (5.145)	4.663 (5.018)
# cars owned	0.309 (0.489)	0.873 (1.548)
# vans owned	0.0506 (0.213)	0.113 (0.574)
# motorcycles owned	0.0515 (0.258)	0.0829 (0.394)
# mopeds owned	0.0277 (0.161)	0.0137 (0.112)
# trailers owned	0.303 (0.598)	0.284 (0.614)
Observations	5855446	

γ_{it} comes from the linear model with interactions.

Means of variables; standard deviations in parentheses.

WD (actual distance): Only available in some years.

WD: is zero (censored) when smaller than 12 km.

Figure B.1: Average Elasticity by Municipality





Figure B.2: Municipality Residents in Top 10% of the Elasticity Distribution



The results indicate that driving response in the top percentile accounts for 5.6% of the total change in the VKT distribution between the two price levels. Taken together, the top 5% account for 14.4% of the change in driving. This indicates the relative importance of the upper tail for determining the response in driving to fuel taxes.

Table B.6: Counterfactual VKT Responses by Quantile: $p_0 = 9$, $p_1 = 10$

Quantile	VKT _{it}	VKT _q (p_0)	Δ VKT _q	Share
.01	5.04	18.09	-1.0351	0.0077
.05	12.05	26.76	-1.3476	0.0400
.10	17.20	31.17	-1.2921	0.0480
.15	21.19	34.04	-1.2578	0.0467
.20	24.64	36.26	-1.2407	0.0460
.25	27.74	38.13	-1.2226	0.0454
.30	30.67	39.79	-1.1743	0.0436
.35	33.48	41.31	-1.1444	0.0425
.40	36.21	42.76	-1.1104	0.0412
.45	38.96	44.17	-1.1193	0.0415
.50	41.72	45.59	-1.1035	0.0410
.55	44.56	47.04	-1.1409	0.0423
.60	47.56	48.58	-1.1696	0.0434
.65	50.75	50.26	-1.2013	0.0446
.70	54.26	52.16	-1.2379	0.0459
.75	58.24	54.39	-1.3254	0.0492
.80	62.94	57.16	-1.4170	0.0526
.85	68.89	60.84	-1.6220	0.0602
.90	77.29	66.46	-1.9936	0.0740
.95	92.48	77.96	-2.9715	0.0882
.99	146.70	121.42	-7.5407	0.0560

Columns 4 and 5 consider the change from p_0 to p_1

Column 5 shows Δ VKT_q / $\sum_{q'} \Delta$ VKT_{q'}, i.e. the share of the total driving response attributable to each quantile.

The justification for the process above follows Melly (2005) and Machado and Mata (2005) and the exposition in section 4.3. Recall that we may think of the quantile regression model as a random parameter model, where each observation draws a quantile, $u_{it} \sim \text{Uniform}(0, 1)$, and then obtains the parameters as $\gamma_{it} = \gamma(u_{it})$ and $\theta_{it} = \theta(u_{it})$, where $\theta(\cdot)$ is the quantile function and θ contains all parameters except for the fuel price elasticity, γ . Now, we may estimate the counterfactual distribution of our outcome variable by taking independent draws of u to obtain the distribution of $\{\gamma(u) \log p + z_{it}\theta(u)\}$, where z contains all regressors except for the fuel price. Melly (2005) explains that this may be approximated using a grid over $[0; 1]$, such as the quantiles we have been using all along, and then weighting the values at each grid point according to the length of the interval, i.e. using w_q as described above.

C Robustness Checks

C.1 Stratifying on Time

Tables C.1 and C.2 shows the implications for the estimated fuel price elasticity of dropping certain years from the sample. We see that in particular if we drop observations where the driving period begins prior to 1999, the estimated elasticity drops a lot and loses significance. This makes a lot of sense in the light of figure A.6 given that the largest variation in fuel prices in the period is before and after the large increase in prices up through 1999. In other words, by restricting the sample to only include driving periods mostly covering the period after the jump in prices, there isn't a lot of variation in the fuel prices to identify the elasticity from.

Table C.1: Robustness: dropping initial years

	(1) Full	(2) 1999-	(3) 2000-	(4) 2001-
$\log p^{\text{fuel}}$	-0.304*** (0.0154)	-0.326*** (0.0165)	-0.384*** (0.0149)	-0.402*** (0.0153)
Year controls	Yes	Yes	Yes	Yes
% of each month	Yes	Yes	Yes	Yes
Car	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes
N	5855446	5681226	5235440	4675560
R^2	0.180	0.182	0.188	0.198

Note: In each column (2)–(4), data before year 97, 98, 99 are dropped respectively.

Rocust standard errors.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C.2 Stratifying on Couples or Singles

Table C.3 shows the results when estimating on the sample consisting exclusively of couples or singles.

C.3 Stratifying on the Length of the Period

In table C.4, we drop the driving periods that have years to test (length of the driving period) further than 3 months away from either 2 or 4 years. Recall that a normal test period will be 4 years for a new car and 2 years for a used car. However, during the phase-in of the inspections, cars were summoned for inspection for the first time and therefore did not necessarily drive the normal length early on. The results show that when

Table C.2: Robustness: dropping later years

	(1) Full	(2) -2006	(3) -2005	(4) -2004
$\log p^{\text{fuel}}$	-0.304*** (0.0154)	-0.258*** (0.0156)	-0.308*** (0.0171)	-0.279*** (0.0187)
Year controls	Yes	Yes	Yes	Yes
% of each month	Yes	Yes	Yes	Yes
Car	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes
N	5855446	5177147	4443035	3736630
R^2	0.180	0.173	0.166	0.161

Note: In each column (2)–(4), data after year 06, 05, 04 are dropped respectively.
Rocust standard errors.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table C.3: Robustness: dropping couples or singles

	(1) Base	(2) Only couples	(3) Only singles
$\log p^{\text{fuel}}$	-0.304*** (0.0154)	-0.318*** (0.0176)	-0.250*** (0.0323)
Year controls	Yes	Yes	Yes
% of each month	Yes	Yes	Yes
Car	Yes	Yes	Yes
Period	Yes	Yes	Yes
Demographics	Yes	Yes	Yes
Household FE	Yes	Yes	Yes
R^2	0.180	0.200	0.108
N	5855446	4550410	1305036

Note: columns (2) and (3) contain only couples or singles respectively.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

we remove these driving periods with non-standard length we find a numerically lower elasticity of -0.275 . In column (2), we include a dummy to control for the non-standard length but this doesn't change the fuel price elasticity much (-0.304).

Table C.4: Robustness: length of the driving period

	(1) Base	(2) Dummy	(3) Subsample
$\log p^{\text{fuel}}$	-0.304^{***} (0.0154)	-0.298^{***} (0.0154)	-0.275^{***} (0.0158)
Non-standard test length		-0.00278^{***} (0.000596)	
Year controls	Yes	Yes	Yes
% of each month	Yes	Yes	Yes
Car	Yes	Yes	Yes
Period	Yes	Yes	Yes
Demographics	Yes	Yes	Yes
Household FE			
R^2	0.180	0.180	0.192
N	5855446	5855446	4535353

Note: Standard test length: years to test is ± 3 months from either 2 or 4 years.

Elsewhere, sample selection requires VKT in $[1;2.5]$ or $[3.5;4.5]$ years.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C.4 Year and Seasonality Controls

Table C.5 shows the results when we change the way we control for time effects in decreasing complexity over the columns. The results show that even if we simplify down to a specification with only a linear time trend, our mean elasticity is unchanged. However, if we remove time controls entirely, the elasticity goes up.

In table C.6, we change the main specification to use the *number* of each month covered by the driving period rather than the *fraction* as we use in the main specification. Our mean elasticity is almost unchanged (from -0.373 to -0.372).

C.5 Clustered Standard Errors

In table C.7, we consider the effects of clustering standard errors at various levels in the specification without household fixed effects. The reason for this is that we cannot cluster on time or municipality because these are not stable within households over time. The results show that even when we cluster at the start-year level, giving 10 clusters, we still get marginal significance ($p = 0.011$). However, this does indicate that the standard errors reported in the fixed effects regressions may be biased towards zero, implying that we are

Table C.5: Robustness: year controls

	(1)	(2)	(3)	(4)	(5)
$\log p^{\text{fuel}}$	-0.304*** (0.0154)	-0.309*** (0.0124)	-0.303*** (0.0123)	-0.313*** (0.00691)	-0.517*** (0.00722)
Linear time trend				-0.0414*** (0.000360)	
Year controls (gas)	Yes	Yes	Yes	No	No
Year controls (diesel)	Yes	No	No	No	No
% of each month	Yes	Yes	No	No	No
Car	Yes	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes	Yes
N	5855446	5855446	5855446	5855446	5855446
R^2	0.180	0.180	0.180	0.177	0.174

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

in reality less precise than the heteroscedasticity robust standard errors would seem to indicate.

C.6 Fuel Type

Table C.8 explores heterogeneity in the fuel price elasticity by the fuel type of the car. Note that when we have household fixed effects, removing one or more rows will drop households entirely if they end up with one or zero remaining periods. Thus, we are removing some of the “switchers” who have responded on the extensive margin of choosing a different vehicle, which we do not model separately in this paper. Therefore, we might expect the elasticity estimate to be impacted by such sample selection.

We see that allowing the elasticity to vary by fuel type results in a larger (in absolute value) mean estimate (-0.398), while the positive coefficient on the interaction of the diesel dummy and the log fuel price implies a smaller elasticity for the diesel drivers (-0.362). Estimating only on the subsamples of each fuel type yields higher elasticities in both subsamples; -0.408 for gasoline and -0.452 for diesel. The reason for this lies in the substitution going on over time; as fuel prices increase, more and more households select into the diesel segment because diesel cars generally are cheaper to use but cost more up front (see e.g. Munk-Nielsen, 2015). Note also that the diesel sample is very small — also, the observation count doesn’t reflect the true number of observations as 1-observation households are still in the 790,390 observations [**todo**: get proper obs count]

Table C.6: Robustness: month controls

	(1) Fraction	(2) Sum
$\log p^{\text{fuel}}$	-0.304*** (0.0154)	-0.304*** (0.0154)
Feb	-0.152*** (0.0394)	-0.00366*** (0.000866)
Mar	-0.0973 (0.0513)	-0.000226 (0.000826)
May	-0.0312 (0.0517)	0.00143 (0.000852)
Jun	0.0515 (0.0404)	0.00344*** (0.000862)
Jul	0.231*** (0.0429)	0.00791*** (0.000890)
Aug	-0.0445 (0.0421)	0.00114 (0.000862)
Sep	0.00781 (0.0410)	0.00255** (0.000820)
Oct	-0.0541 (0.0412)	0.00105 (0.000829)
Nov	-0.141*** (0.0423)	-0.00183* (0.000841)
Dec	-0.174*** (0.0440)	-0.00257** (0.000937)
Apr		0.00199* (0.000851)
Year controls	Yes	Yes
Car	Yes	Yes
Period	Yes	Yes
Demographics	Yes	Yes
Household FE	Yes	Yes
N	5855446	5855446
R^2	0.180	0.180

(1): The share of the driving period falling in each month.

(2): The number of months covered by the driving period.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table C.7: Robustness: clustered standard errors — pooled regression

	(1)	(2)	(3)	(4)
	No cluster	Municipality	Mun., start year	Start year
$\log p^{\text{fuel}}$	-0.298*** (0.0143)	-0.298*** (0.0161)	-0.298*** (0.0345)	-0.298* (0.117)
Year controls	Yes	Yes	Yes	Yes
% of each month	Yes	Yes	Yes	Yes
Car	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes
Household FE	No	No	No	No
R^2	0.340	0.340	0.340	0.340
N	5855446	5855446	5855446	5855446

Note: Clustering levels = none, municipality, (municipality, start year), start year.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table C.8: Robustness: elasticity by fuel type

	(1)	(2)	(3)	(4)
	Base	Interaction	Gas only	Diesel only
$\log p^{\text{fuel}}$	-0.304*** (0.0154)	-0.257*** (0.0191)	-0.268*** (0.0194)	-0.541*** (0.0260)
Diesel=1 $\times$ $\log p^{\text{fuel}}$		-0.135*** (0.0279)		
Year controls	Yes	Yes	Yes	Yes
% of each month	Yes	Yes	Yes	Yes
Car	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes
R^2	0.180	0.180	0.140	0.135
N	5855446	5855446	5018019	837427

In columns 3 and 4, only a single set of time controls is included.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

C.7 Fuel Efficiency and Car Price

In this section, we will argue why the exclusion of the fuel efficiency and car price characteristics does not bias our estimate of the fuel price elasticity. First, we will show that adding the control (in the subsample where the variable is observed) does not change the fuel price elasticity, and then we will show that the included variables account for a substantial portion of the variation in the omitted characteristics.

In table C.9, we show the results of accounting for fuel efficiency and car price (merchant suggested retail price, MSRP). The reason why these variables are not included in the main specifications is that they are only available for a subset of the period. The data source for these variables is the Danish Automobile Dealer Association (DAF). This dataset has been merged to the vehicle identification numbers used by the Motor Register and the authors gratefully acknowledge Ismir Mulalic at DTU Transport for help.

The merge is incomplete and there are some vehicle identifiers in the Motor Register that get no matches. In particular, MSRPs for cars vintages before 1997 is unavailable. This means that any analysis involving fuel efficiency and MSRP will produce a biased sample, including only new cars in the earlier years and then gradually including older cars.

The results in table C.9 indicate precisely that the sample where the characteristics are observed is different from the estimation sample used throughout this paper; the fuel price elasticity increases from -0.304 to -0.591 . This can be explained by there being more households with newer cars; from the interaction results, we saw that households who tend to have newer cars tend to also be more price sensitive (see table 5.2).

Including the fuel efficiency variable lowers the elasticity marginally from -0.591 to -0.582 . Further including the MSRP leaves this almost unchanged (-0.584). We take this as an indication that the included car characteristics capture the most important facets of the selection process. Of course, this test is not perfect given that we are unable to use fuel efficiency for the full sample.

Finally, as mentioned in section 7.1, we will briefly argue that the included car characteristics capture most of the variation in the omitted ones. That is, our prior is that gross weight, fuel type, a dummy for van and the car vintage will explain most of the variation in the two variables; this follows the line of thought behind the efficient production frontiers for car producers as estimated by (Knittel, 2011). Basically, producers trade off weight and engine power against fuel efficiency along an efficient frontier (with some statistical noise around). In that sense, it is not strictly necessary to include all variables. For the sample where the characteristics are available, we find that a regression of fuel efficiency on weight, weight squared and dummies for diesel, van and vintage yields an R^2 of 0.72. The similar R^2 for MSRP is 0.55. Thus, the included variables capture much of the variation in the omitted ones.

Table C.9: Robustness: controlling for fuel efficiency and car MSRP

	(1) Base	(2) Subsample	(3) Control	(4) Controls
$\log p^{\text{fuel}}$	-0.304*** (0.0154)	-0.591*** (0.0166)	-0.582*** (0.0166)	-0.584*** (0.0166)
Fuel efficiency in km/l (improved)			-0.00249*** (0.000338)	0.00166*** (0.000343)
price_new				0.000000478*** (9.93e-09)
Weight (ton)	0.00167*** (0.00000798)	0.00196*** (0.0000125)	0.00193*** (0.0000131)	0.00176*** (0.0000133)
Weight squared	-0.000000354*** (2.00e-09)	-0.000000383*** (3.30e-09)	-0.000000380*** (3.33e-09)	-0.000000365*** (3.34e-09)
Diesel	0.259*** (0.00545)	0.214*** (0.00766)	0.228*** (0.00793)	0.195*** (0.00787)
Van	-0.205*** (0.00170)	-0.229*** (0.00225)	-0.232*** (0.00228)	-0.136*** (0.00278)
Car age (ultimo)	-0.0293*** (0.000141)	-0.0195*** (0.000233)	-0.0201*** (0.000247)	-0.0178*** (0.000248)
# cars owned	-0.0501*** (0.00109)	-0.0258*** (0.00135)	-0.0259*** (0.00135)	-0.0271*** (0.00137)
# vans owned	-0.0654*** (0.00179)	-0.0798*** (0.00220)	-0.0796*** (0.00220)	-0.0836*** (0.00223)
# motorcycles owned	0.0118*** (0.00178)	0.00889*** (0.00216)	0.00883*** (0.00216)	0.00859*** (0.00217)
# mopeds owned	0.0204*** (0.00217)	0.0161*** (0.00272)	0.0161*** (0.00272)	0.0162*** (0.00271)
# trailers owned	0.00595*** (0.00106)	0.00905*** (0.00128)	0.00900*** (0.00128)	0.00909*** (0.00128)
Year controls	Yes	Yes	Yes	Yes
% of each month	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes
R^2	0.180	0.202	0.202	0.205
N	5855446	3035301	3035301	3035301

(2), (3) and (4) restricts the sample to fuel efficiency and car MSRP being observed.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure D.1: Danish Fuel Prices and the WTI Oil Price





Figure D.2: Comparing Two Oil Price Series



D Instrumental Variables

In this section, we present results from instrumenting for the fuel price. Our primary instrument is the WTI crude oil price in USD per barrel. The price is converted to DKK using the spot USD price and then deflated using the Danish CPI. Figure D.1 shows the oil price together with the Danish real fuel prices. We have also included the Brent oil prices and the two price series are shown together in figure D.2 and clearly follow each other very closely.

Table D.1 shows the main two-stage least squares results, instrumenting log real fuel price with log real WTI oil price.

Table D.2 shows the first stage results run using the micro data. Note that the very high R^2 of 98% is partially due to the fact that overlapping periods are repeated. [**todo:** do this using the daily observations alone]

Table D.1: Instrumental Variables Results

	(1) OLS	(2) FE	(3) 2SLS	(4) 2SLS FE
$\log p^{\text{fuel}}$	-0.298*** (0.0143)	-0.304*** (0.0154)	-0.511*** (0.0148)	-0.368*** (0.0161)
Year controls	Yes	Yes	Yes	Yes
% of each month	Yes	Yes	Yes	Yes
Car	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Demographics	Yes	Yes	Yes	Yes
Observations	5855446	5855446	5855331	5855296

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$


Table D.2: Instrumental Variables Results: First Stage

	(1) Simple	(2) Full
diesel	-0.973*** (0.000299)	-0.855*** (0.000700)
log_oil	0.176*** (0.0000187)	0.177*** (0.0000552)
diesel_log_oil	0.147*** (0.0000519)	0.124*** (0.000129)
All controls	No	Yes
Household FE	No	No
N	5855331	5855331
R^2	0.972	0.982

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure E.1: Nonparametric Demand Curve



E Nonparametrics and Semiparametrics

In this section, we apply nonparametric and semiparametric techniques to explore the demand curve in greater detail.

E.1 Nonparametric Demand Curves

Here, we consider the simple relationship,

$$\log \text{VKT}_{it} = m(\log p_{it}^{\text{fuel}}) + u_{it},$$

where $m(\cdot)$ is an unknown function. For the relationship to make sense, we must assume that $\mathbb{E}(u_{it} | \log p_{it}^{\text{fuel}}) = 0$. This will obviously not be true but from the graph we can still learn whether the story of a strong price response is well supported even by the raw data. We apply a local linear regression estimator and plot the estimated function for diesel and gasoline cars separately due to the large difference in mean driving between these two groups. The resulting curves are shown in figure E.1. The demand curves are clearly downward sloping but each have a flat (or almost flat) region; For gasoline cars in the (log) price range below 2.1 and for diesel cars in the range $[1.95; 2.05]$.

E.2 Semiparametric Demand Curve

In this section, we instead consider the equation,

$$\log \text{VKT}_{it} = m(\log p_{it}^{\text{fuel}}) + X_{it}\beta + u_{it},$$

where X_{it} is a $1 \times K - 1$ vector of household demographics and time dummies. This specification is like the main OLS regression considered in this paper with two exceptions; It does not control for household-specific fixed effects but it allows $\log p^{\text{fuel}}$ to enter nonlinearly without placing any functional form restrictions on $m(\cdot)$. The function $m(\cdot)$ is estimated using the double residual method by Robinson (1988) — the method is described in section E.2.1. The resulting graph is shown in 7.1. The graph clearly demonstrates a predominantly linear relationship.

E.2.1 The Robinson (1988) Double Residual Method

Consider the semiparametric regression,

$$y_i = m(z_i) + x_i' \beta + u_i,$$

where y_i, z_i, u_i are scalars, x_i is $K \times 1$ and it is assumed that $\mathbb{E}(u_i | x_i, z_i) = 0$. The estimator proposed by Robinson (1988) proceeds in three steps;

Step 1: Compute

$$\begin{aligned} \tilde{y}_i &= y_i - \hat{m}_y(z_i), \\ \tilde{x}_{ik} &= x_{ik} - \hat{m}_{x_k}(z_i), \end{aligned}$$

where the functions $\hat{m}_y(\cdot), \hat{m}_{x_k}(\cdot)$ are the orthogonalized y_i and x_{ik} respectively, defined by

$$\begin{aligned} \hat{m}_y(z) &= \sum_{i=1}^N w_i(z) y_i, \\ \hat{m}_{x_k}(z) &= \sum_{i=1}^N w_i(z) x_{ik}, \\ w_i(z) &= \frac{K_h(z - z_i)}{\sum_{i=1}^N K_h(z - z_i)}. \end{aligned}$$

Step 2: Estimate the linear coefficients, β , as

$$\hat{\beta} = (\tilde{X}' \tilde{X})^{-1} \tilde{X}' \tilde{Y},$$

where $\tilde{X}$ and $\tilde{Y}$ are the stacked versions of $\tilde{x}_i$ and $\tilde{y}_i$.

Step 3: The unknown function, $m(\cdot)$, can now be estimated as the usual nonparametric estimator of

$$\tilde{v}_i = m(x_i) + u_i,$$

where $\tilde{v}_i = y_i - x_i' \hat{\beta}$.

Standard errors: The 95% CI bounds shown in the figures around $\hat{m}(\cdot)$ are found by ignoring the first-stage estimation of $\hat{\beta}$ — in other words treating $\tilde{v}$ as data.

Chapter 2

Diesel Cars and Environmental Policy

Diesel Cars and Environmental Policy

Anders Munk-Nielsen*

December 3, 2015

Abstract

In this paper, I measure the costs of environmental taxation of car ownership and usage in Denmark. Using full population Danish register data covering 1997–2006, I estimate a discrete-continuous model of car choice and usage that explicitly allows households to select cars based on expected usage conditional on observed and unobserved heterogeneity. I validate the model using a major Danish reform in 2007 which prompted a substantial shift in the characteristics of purchased cars unique to the Danish setting compared to the rest of Europe. Through counterfactual simulations, I find that both Danish reforms in 1997 and 2007 were cost-ineffective at reducing CO₂ emissions compared to a fuel tax. Moreover, I find that the diesel market share responds strongly to taxation but that environmental goals can be reached both with and without a large diesel share in the fleet.

Keywords: Car taxation, fuel taxation, environmental policy, discrete/continuous choice estimation.

JEL codes: D12, H23, Q53, Q58, L98.

*Department of Economics, University of Copenhagen, anders.munk-nielsen@econ.ku.dk. I would like to thank David Brownstone, Hamish Lowe, Bo Honoré, Aureo de Paula, John Rust, Kenneth Gillingham, Mogens Fosgerau, Jesse Burkhardt, Bertel Schjerning, Søren Leth-Petersen, Jeppe Druedahl and Thomas Jørgensen for comments and feedback. This paper is part of the IRUC research project financed by the Danish Council for Strategic Research (DSF). Financial support is gratefully acknowledged. All remaining errors are my own.

1 Introduction

Since 1980, greenhouse gas emissions from the Danish transport sector have increased from 10 to 15 mio tons CO₂ annually while all remaining sectors together have reduced emissions from 55 to 30 mio tons. In Denmark as well as the rest of the developed world, a consensus is emerging that emissions from the transport sector must be decreased if environmental goals are to be reached. The goal of this paper is to measure the cost-effectiveness of environmentally motivated tax policies that have targeted car choice and use.

Towards this end, I estimate a structural 2-period discrete-continuous model of new car purchase and subsequent usage by Danish households. My dataset covers all new car purchases for the period 1997–2006 as well as subsequent driving over a 4-year period and detailed demographics from the Danish registers. In 2007, a major Danish reform was implemented, followed by substantial changes in the characteristics of newly purchased cars. In particular, the diesel share of new cars in Denmark increased remarkably compared with other European countries at that time. My sample period stops before the reform but I know the response from external sources and can use this to validate my model.

My results contribute to the understanding of the costs of environmental car taxation. The model gives predictions on car choices and subsequent driving, allowing me to analyze the impact of counterfactual policy scenarios on tax revenue, substitutions in the new car market, total driving, fuel demand and CO₂ emissions. I find that a simple fuel tax would have been more efficient per ton of CO₂ than both the 1997 and 2007 reforms were. Other studies found fuel taxes to be more effective compared with taxes that target car characteristics (Grigolon, Reynaert, and Verboven, 2015) and with emissions standards (Jacobsen, 2013).

I also contribute with new insights regarding the increasing diesel share. This has received attention by policy makers as awareness has increased about the negative health effects of local air pollution from diesel cars.¹ A key descriptive fact is that diesel car drivers tend to drive on average 60.0% more than gasoline car drivers. I therefore estimate a high-dimensional discrete-continuous model that explicitly accounts for selection based on observed and unobserved heterogeneity in driving. To my knowledge, I am the first to empirically explore the rise in the diesel share accounting for endogeneous selection. My findings indicate that the diesel share is highly sensitive to the way that car taxes discriminate between gasoline and diesel cars. Environmentally motivated car taxes tend to target the fuel efficiency but must correct for the inherently higher efficiency of diesel cars. I show that the diesel cars are neither necessary nor sufficient for environmental goals. To shed light on what the diesel would be in absence of discriminatory taxation

¹In 2012, the World Health Organization moved diesel exhaust to their list of carcinogens — substances that are definitely known to cause lung cancer.

based on fuel type, I counterfactually equalize car taxes and fuel taxes for gasoline and diesel cars and find a level slightly higher than that in 2006, but lower than for most other European countries.

My findings complement existing knowledge on car choice and usage due to the unique nature of my setting; by studying a small open economy without domestic car production and using a reform to explore the validity of the model, I can address some of the issues that are inherent in many of the classic studies of car taxation. Firms respond to car tax policies for example by changing their portfolios (Reynaert, 2014). Policies that affect a small market such as the Danish will tend to provide a smaller incentive for automakers to change their portfolios, reducing this supply side concern. Similarly, the market is too small for shocks that are unique to Denmark to affect global fuel prices.

There may, however, still issues with common demand shocks across countries, such as increasing urbanization. Therefore, another strength is the access to full population detailed register data, including demographic information on work distance, income. In addition to accounting for changing urbanization patterns, this allows me to model household driving very precisely. Thereby, I can also give an accurate estimate of the response in driving to an exogenous increase in fuel efficiency (the so-called *rebound effect*), which has been widely debated in the literature (e.g. Small and Van Dender, 2007; Bento et al., 2009; Gillingham, 2012; Hymel and Small, 2015). I estimate the rebound effect for Denmark to be -0.30 .

The rest of the paper is organized as follows; Section 1.1 discusses the contributions from this paper in the context of related literature. Section 2 presents the institutional setting and the data and presents some preliminary descriptive evidence. Section 3 lays out the theoretical model while Section 4 gives the empirical strategy for estimation and discusses identification. Section 5 presents the estimates and structural elasticities. Section 6 contains the counterfactual policy simulations and section 7 concludes. Appendix A contains a list of the notation used throughout the paper as well as the core equations of the structural model for easy reference.

1.1 Related Literature

I mainly contribute to the literature on the cost of environmental policies in the car market. Recently, a number of papers have emphasized European settings. D’Haultf uille, Givord, and Boutin (2014) study the French *Bonus/Malus* reform of 2008 which is a feebate similar to the Danish one. They find that the reform had a negative environmental impact, mainly because it led to more cars being sold at the extensive margin. My model conditions on entry into the new car market so I make no claims on the extensive margin results. Adamou, Clerides, and Zachariadis (2013) counterfactually study the impact of a feebate, finding that the reform needs to look more like a fee than a rebate in order to be

optimal. Grigolon, Reynaert, and Verboven (2015) find that fuel taxes are more efficient than vehicle taxes in reducing fuel usage than taxes working through the fuel efficiency of cars. Using cross-country market-level data, they find that discriminatory fuel taxes and differences in fuel efficiency alone explain 40% of the differences across countries. My results indicate that discriminatory ownership and purchase taxes may well account for a substantial part of the remaining 60%. Mabit (2014) also uses Danish data and analyzes the 2007 reform that is also under study in this paper and finds the changes in car characteristics occurring in the period to be as important as the reform.

A number of other studies consider more small-scale reforms, typically affecting smaller segments. These are generally found to be cost-ineffective. Huse and Lucinda (2013) consider a Swedish reform affecting only highly efficient green cars using a BLP model. They find that the implicit price of CO₂-emissions from that reform was far above the social cost of carbon in Sweden. Beresteanu and Li (2011) and Chandra, Gulati, and Kandlikar (2010) study incentive schemes aimed at hybrid cars in the US and Canada and both find them to be cost-ineffective.

The papers cited above all target the demand side of the market but a large American literature focuses on supply side instruments, primarily the Corporate Average Fuel Economy (CAFE) standards. These require car makers to reach a certain weighted average fuel economy across their sold cars, subject to a number of technical details. Goldberg (1998) is one early study of CAFE standards utilizing joint modeling of car choice and usage, finding that policies targeting the car choice are favorable to fuel taxes. Building on the framework by Bento et al. (2009), recent work by Jacobsen (2013) compares the cost-effectiveness of CAFE standards and fuel taxes, finding the latter to be the more effective. Reynaert (2014) and Clerides and Zachariadis (2008) are among the few papers studying the effects of the European fuel economy standards, announced in 2007 and to be fully binding by 2015. Reynaert (2014) focuses on the responses of the European automakers, finding that they primarily respond by technology adoption.

A different strand of literature looks at the fuel type of the purchased cars, focusing on the choice of diesel vs. gasoline. This is a much more prevalent option in the European than the American context and the diesel market share increased substantially up through the early 1990's, following the introduction of the direct injection or common rail technology. Miravete, Moral, and Thürk (2014) study this in the Spanish setting, finding that the policy treatment of diesel vs. gasoline in Europe functioned in effect as a subsidy to European car makers. On the methodological side, Verboven (2002) uses within-model variation between car models that only differ in using gasoline or diesel fuel for identification in a BLP framework. Grigolon, Reynaert, and Verboven (2015) also consider heterogeneity in driving but assume a zero fuel price elasticity of driving. My paper is the first to my knowledge to study the dieselization while estimating the driving decision simultaneously.

Endogenous selection of consumers into car types based on individual driving demand has been emphasized in recent work. This paper builds on Gillingham (2012) who introduces endogenous selection both based on observables, unobservables and explicitly on expectations about future fuel prices. The model builds on Dubin and McFadden (1984). Some work has used 2-step approaches to integrating type choice and usage (e.g. Goldberg 1998; West 2004; D’Haultf  uille, Givord, and Boutin 2014), while more recent work has promoted simultaneous estimation (e.g. Bento et al. 2009; Feng, Fullerton, and Gan 2013; Jacobsen 2013 and in particular Gillingham, 2012). The model explicitly accounts for the selection effect required to identify the so-called *rebound effect*, namely the effect on driving of increasing fuel efficiency (see e.g. Small and Van Dender, 2007).

In terms of the data used, this paper is novel in applying micro data on car choice and usage matched with household-level demographics for the full Danish population over a long period of 9 years. Many papers in the car demand literature have only used market-level data (e.g. Berry, Levinsohn, and Pakes 1995; Miravete, Moral, and Thurk 2014; Reynaert 2014; Verboven 2002). The papers using micro-level data either use survey data (West 2004; Bento et al. 2009; Jacobsen 2013), often with only a limited number of years, or do not observe household demographics at the micro level (e.g. Gillingham, 2012).

Two major aspects of car demand that I do not tackle in this paper are multi-car households, dynamics and myopia. Even though the data would allow it, I choose not to include 2-car households in this study.² This is to make sure the choiceset in the model remains tractable. Since only 12.1% of Danish households between 18 and 65 years own 2 or more cars, I capture the largest segment this way (see Figure B.8).

A recent literature has looked at the question of whether consumers correctly take into account future savings in fuel cost when making a car purchase.³ I make no claims to answering this question but will follow the empirical work indicating that that consumers are rational and time-consistent when they make their vehicle and driving decisions. However, I will allow some flexibility in consumer expectations about future fuel prices.

Finally, some authors have emphasized the dynamics of vehicle ownership decisions, opting for a fully dynamic structural model.⁴ While this facilitates the study of important

²Some of the only studies focused on modeling multi-car households are Spiller (2012); Borger, Mulalic, and Rouwendal (2013); Wakamori (2011). Bento et al. (2009) take a different approach, considering each car as a *choice occasion*. An alternative approach in my setting would be to ignore knowledge about other cars and consider the two instances as independent or to add a control.

³The findings have been mixed with some support for myopia (Allcott and Wozny, 2012) and some against (Busse, Knittel, and Zettelmeyer (2013); Sallee, West, and Fan (2010); Grigolon, Reynaert, and Verboven (2015)). The interested reader is referred to the literature review by Greene (2010) which documents that there has been extremely mixed evidence in the empirical literature. Another strand of literature emphasizes certain behavioral aspects that I will not consider in this paper; Gallagher and Muehlegger (2011) find that tax incentives working through the purchase price are more effective than ones working through income tax deductions, and Li, Linn, and Muehlegger (2014) find that driving responds more strongly to fuel taxes than to changes in the fuel product price.

⁴Many recent dynamic models build on the optimal replacement model by Rust (1987). These models are much better suited to looking at issues like vehicle scrappage (Adda and Cooper, 2000; Schiraldi,

aspects such as the used-car market, scrappage and ownership durations one must trade off complexity elsewhere in the model and it is central to maintain a high-dimensional choicset to accurately fit in the effects of the policies considered. As most other non-dynamic papers, the model presented in this paper conditions on entry into the new car market. If the reforms change substitutions between the used and new car market, such effects will be ignored. In that sense, the focus of this paper is purely on the substitution patterns in the car market.

2 Background and Data

In this section, I will first describe the institutional setting in Denmark, focusing on the taxation of cars in the period. I then discuss the data, explaining the different data sources and the construction of the final dataset. Finally, I present descriptive evidence on car choice and driving.

2.1 Institutional Setting

Car taxation in Denmark consists of three elements; a registration tax, a bi-annual ownership tax and fuel taxes. The registration tax is paid at the time of purchase and is a linear function of the purchase price with a kink,

$$\tau_t^{\text{reg}}(p^{\text{gross}}) = 1.05 \cdot \min(K_t, p^{\text{gross}}) + 1.80 \cdot \max(0, p^{\text{gross}} - K_t),$$

where K_t is a politically set kink, $\tau_t^{\text{reg}}(\cdot)$ denotes the registration tax and p^{gross} is the raw car price including VAT (25%) but net of deductions.⁵ Consequently, taxes make up just over 160% of the purchase price of the average Danish car. The second tax, ownership tax, is paid twice a year and depends on the fuel efficiency (in kilometers per liter, km/l) of the car according to a schedule that is updated irregularly over the period and accounted for in the estimation. There is a separate schedule for diesel cars where the tax rate is higher for any given level of fuel efficiency. This balances the fact that diesel cars on average have higher fuel efficiency than gasoline cars. The third tax element, fuel taxes, are comprised of a fixed and a proportional component and the total fuel tax amounts to 68.0% of the gasoline price, averaged over my sample period (58.5% for diesel). The composition of taxes and product price for the gasoline and diesel prices are shown in

2011), and the used car market (Adda and Cooper (2000); Schiraldi (2011); Chen, Esteban, and Shum (2010); Gavazza, Lizzeri, and Rokestkiy (2014); Gillingham et al. (2013); Stolyarov (2002) Chen, Esteban, and Shum, 2010; Gavazza, Lizzeri, and Rokestkiy, 2014; Stolyarov, 2002; Gillingham et al., 2013). Such issues are beyond the scope of this paper.

⁵Deductions are given for example for installed safety equipment which are not observed in the data and therefore ignored in this paper. Anecdotally, some deductions are larger than the cost of installing the equipment, meaning that the equipment is universally adopted.

Figure B.4.

There were two major reforms of interest in the sample period; A change in the bi-annual tax in 1997 and a change in the registration tax in 2007. My data does not cover both before and after either of these reforms. All cars first registered before July 1st 1997 have their bi-annual tax rate set according to the weight (and still follow that scheme) while those first registered after that date follow the fuel efficiency. The 2007 reform was a so-called *feebate*, working through the registration tax and giving a rebate to green cars and added a fee to inefficient cars. The rebate was DKK 4,000 per unit of km/l over the pivot (16 km/l for gasoline cars and 18 km/l for diesel cars). The corresponding fee was slightly lower, at 1,000 DKK per km/l. Figure 2.1 shows the prompt change in the fuel efficiency of new cars after the reform is introduced and 2.2 shows an even greater change in the diesel share of newly purchased cars. From the European Automobil Dealer Association, I have access to the diesel share in other European countries, which is also shown in 2.2, highlighting that the response was unique to Denmark.⁶

2.2 Data

The dataset contains all new cars purchased between July 1st 1997 and December 31st 2006 and is based on matched Danish administrative data. The car ownership information comes from The Central Motor Register, which holds license plate ownership information. Driving information comes from the mandatory safety inspection which all cars must attend four years after purchase. At this test, it is evaluated whether the car is in safe condition and the odometer is measured and recorded. Therefore, the driving data comes from a 4-year period following purchase. Demographic informations on the car owners and the remainder of their household is obtained by matching the personal identifier (CPR number) with the Danish registers. The most important variable is the computed work distance measure (described in appendix B.3). This measure captures the product of the work distance and the number of days that the individual goes to work, regardless of the mode choice. Households are only eligible for the deduction if they are working and their private address is further than 12 km from the address of their primary work place, which is the case for a little under half of the individuals. Appendix B.3.1 provides details on this unique variable.

A car type in the data is defined as a unique Vehicle Type Approval number. These are identifiers assigned by the Ministry of Transportation when a car is approved for import and sale in Denmark. They vary at a finer level than the traditional (make-model-year) in some respects, since any change in the vehicle that might alter safety aspects of the car in

⁶I have been unable to get the similar fuel efficiency numbers for other European countries. I expect that the Danish response is unique in relation to the timing but that the general trend is certainly shared across countries. The source for the Danish diesel share and average fuel efficiency post-2007 is Statistics Denmark's aggregate statistics (statistikbanken.dk), but I do not have this information in my micro data.

operation require a new approval for import. The identifier does not contain information on the make year, however. Car characteristics are merged using this identifier. An important limitation of the data is that I do not observe the age of the car; instead, I observe the year the car was first registered in Denmark and use this to construct the age, assuming that the car is not an imported used car. Imports of used cars are not a big problem for my setting because the high Danish car taxes imply that the used-car prices are generally very high. I have access to new car prices and depreciation rates are available from a dataset maintained by the Danish Automobile Dealer Association (DAF). The depreciation rates are used by used car dealers in Denmark when they make an offer on a used car of a given age in normal condition and the new car prices are merchant suggested retail prices (MSRPs). Fuel prices are available at the daily level from the Danish Oil Industry Association (EOF; www.eof.dk). These prices are recommended retail prices for the entire country so local variations and price wars do not show up in the data.⁷ In Appendix B.3.2, I show that the product prices of both types of fuel track international oil prices very closely (Figure B.5). All tax rates are taken directly from the law texts using www.retsinformation.dk with the exception of fuel taxes, that come from EOF.

As many of the classic car choice papers, the emphasis of this paper is on the new car market. While car ownership is observed for used cars, prices and characteristics are only available for cars purchased from 1997 and forward.


In order to evaluate the welfare consequences of the counterfactual policies, one needs a measure of the marginal external costs of driving. These are taken from DTU Transport (2010) and shown in Appendix B.2.⁸ The key thing to note about externalities is that the per-kilometer externality of congestion and accidents are far larger than environmental externalities (this has been emphasized by e.g. De Borger and Mayeres, 2007).

The final estimation sample contains $N = 128,910$ new car purchases by Danish couples in 1997–2006. The sample selection is described in details in Appendix B.1. To ensure demographic heterogeneity, I have selected only households consisting of couples. Adding singles could easily be done but would require many additional parameters and they account for less than 20% of all new purchases. I also deselect cars with missing observations as well as car types that are purchased fewer than 30 times. The final dataset has a total of $J = 1,177$ different cars to choose from. Even so, the choiceset facing a single household is much smaller than this because no car was available in all sample years. Working with a choiceset of this high dimensionality in a discrete choice setting is challenging but it allows me to implement and explore the tax system very precisely.

⁷In the literature estimating the demand for driving, many papers rely on spatial variation in fuel prices for identification. This would not be appropriate for Denmark, however, since the country is so small that it would be hard to establish regions that would avoid trading across markets.

⁸I have recalculated from a per kilometer to per liter externality in terms of air pollution from CO₂ and other particle emissions.

Figure 2.1: Fuel Efficiency of Newly Purchased Cars in Denmark



Source: Statistics Denmark, avg. fuel efficiency of newly purchased cars.
The horizontal line marks the 2007 reform in April.
The dots show yearly averages over the middle date of the year.


2.3 Descriptive Evidence

In the period 1997–2006, the fuel efficiency of newly purchased cars increased from just over 13 kilometers per liter (km/l) to 16 km/l, as shown in Figure 2.1. The figure furthermore shows a sharp change occurring right when the Danish feebate of 2007 was implemented in 2007. However, in the same period there was a drastic increase in diesel car sales, which made up 3.0% of all new cars sold in 1997 but had increased to 26.3% by 2006. Furthermore, this number increased to 38.4% in 2007. While the increasing trend over the period was shared by many other European countries, where the average diesel share increased from 22.3% to 50.8%, the jump in 2007 is absent for those countries. Figure 2.2 shows the diesel share of new purchases for Denmark together with 4 other countries and the Western European average. The common trend naturally opens the question of how much of the changes in characteristics was driven by changes in demand, supply and policy.

Table 2.1 shows summary statistics for the estimation sample both in terms of cars and households. Regarding average work distance variable, this is zero if the household has less than 12 km to work. The reported averages of 11.8 km for males and 8.12 for females are therefore the averages of this censored variable.

To get a first grasp of the conditional correlations in the data, Figure B.13 shows the distribution of driving for gasoline car drivers and for diesel car drivers. The average gasoline car drives 49.2 km per day while the average diesel car drives 78.8 km per day. This is confirmed in Table B.4; the table shows regressions where car characteristics of the purchased vehicles are regressed on the demographic variables of the households purchasing them. The estimates indicate that an increase in the male’s work distance

Figure 2.2: Diesel Cars — Fraction of Total New Car Sales in European Countries



of one standard deviation is associated with the probability that the household buys a diesel car by 5.1 percentage points. The coefficient on real household income is positive for weight, engine power (kW) and size (cc) and the real price. This means that richer households tend to buy larger and more powerful cars. Figure B.12 visualizes the spatial dimension of this and shows that the urban regions of Denmark have low work distances and low diesel shares while some of the regions with the longest work distances also have a higher prevalence of diesel cars. Moreover, the figure shows that there are rural regions in the eastern part of Denmark where diesel cars are very popular in spite of work distances being lower. Appendix B.3 contains more descriptive statistics.

In Appendix B.4, I present detailed descriptives for the fuel price development over time. Fuel prices have increased by 23.0% and 33.7% for gasoline and diesel fuel respectively. This has mainly been driven by changes in the product price as Danish fuel taxes rates have been largely unchanged over the period (cf. Figure B.4). The fact that the diesel share has increased in spite of diesel fuel prices growing faster than gasoline prices indicates that either the characteristics or the differential tax rates of diesel cars have changed even faster in a favorable direction. Finally, even though the relative price of diesel to gasoline has increased from 80.6% to 89.2% over the period, there has substantial gyrations in the relative price year to year (Figure B.6).

More detailed descriptives are presented in appendix B.3 but to paraphrase, the only household demographic that appears to predict diesel purchase is the home-work distances of each of the spouses. This variable is also an important predictor of the household's vehicle kilometers travelled (VKT) and elasticity of driving with respect to the price per kilometer (PPK). The variable is rarely available in empirical studies and often considered to be the main component of household fixed effects in driving equations.

Table 2.1: Summary Statistics — Shortened Names in Parentheses

Car Variables			
	N	Mean	Std.
Fuel efficiency (km/l, e)	128,910	14.68	2.56
Weight (tons, q^{weight})	128,910	1,660.80	201.63
Horsepower (kW, q^{kw})	128,910	70.71	16.94
Displacement (cc, q^{displace})	128,910	1,580.08	265.40
Diesel (%)	128,910	0.1108	0.31
Price (2005 DKK, p^{car})	128,910	219,284.20	66,522.11
Depreciation factor (δ)	128,910	0.8741	0.0118
Units Sold	128,910	228.20	213.48
Demographic Variables			
	N	Mean	Std.
Work distance, male (WDm)	128,910	11.80	19.63
Work distance, female (WDf)	128,910	8.12	14.84
Gross income (2005 DKK, inc)	128,910	701,058.5	456,223.5
Number of kids (nkids)	128,910	0.9866	1.07
Unemployment, male (unempm)	128,910	0.0859	0.28
Unemployment, female (unempf)	128,910	0.1616	0.37
Age, male (agem)	128,910	43.99	10.12
Age, female (agef)	128,910	42.00	10.27
Male income %	128,910	0.5894	0.13
Urban area (bigcity)	128,910	0.2084	0.41

3 Model

In this section, I outline the decision model of the households. I first present the functional form of the two-period utility function. I then solve for optimal planned driving in the second period, before inserting this back into the first-period utility to derive the expected utility of choosing a given car.

The model builds on the discrete-continuous selection model literature going back to Dubin and McFadden (1984). The idea is that the usage in the second period comes out of Roy's identity. This type of framework was applied to car choice and usage by Bento et al. (2009) and Gillingham (2012). The model presented below is based closely on the latter but with the extension of allowing household demographics to affect driving not only through the price sensitivity parameter but also through the mean driving.

3.1 Household Utility

The model is a two-period model; in the first period, t_1 , the household purchases a car of type j at the price p_j^{car} under uncertainty about fuel prices in the future. In the second period, t_2 , fuel prices are realized and the household makes its driving decision. Households enter the new market at different points in time and thus face different sets of available cars, $\mathcal{J}_{t_1}$, and different fuel prices. In the implementation, t_1 is the calendar year in which the household enters the new car market, i.e. $t_1 \in \{1997, \dots, 2006\}$. The driving period length is four years, because the first mandatory safety inspection at which the odometer is measured in the data occurs after four years. At the end of the second period, four years later, the car is sold at the used-car price given by $\delta_j^4 p_j^{\text{car}}$, where δ_j is a car-specific annual depreciation factor obtained from the Danish Automobile Dealer Association (the δ_j is 0.874). There is no outside option of not owning a car and there are no used cars in the choiceset.⁹ In that sense, the model conditions on entry into the new car market but remains agnostic about why and when this entry occurs.¹⁰

The utility function takes the form

$$u_{ij} = u_{ij1} + \beta^4 \mathbb{E}(u_{ij2}),$$

where β is the annual discount factor (fixed at 0.95) which is raised to the power four because there are four years between purchase and driving period. Both of the period-

⁹The main reason for not having an outside option because this simple quasi-linear two-period model is not well-suited to deal with the inherently dynamic problem of purchasing a car, which represents a major investment in Denmark on account of the large taxes. I ignore the used-car market partly due to missing data on car characteristics, which would heavily skew my sample over time, and partly due to the dimensionality; including that many more car types would force be to reduce the dimensionality of the choiceset.

¹⁰One could imagine a fully dynamic optimal stopping problem where the consumer in each period considers replacing his current car, e.g. Schiraldi (2011). However, then it would be computationally very challenging to have a choiceset of $J = 1,177$ cars.

utilities are quasi-linear in the consumption of the composite outside good. First-period utility takes the form

$$u_{ij1} = \gamma_i \left(y_{it_{1i}} - p_j^{\text{car}} - 4\tau_j \right) + u^{\text{own}}(j),$$

where $u^{\text{own}}(j)$ is utility from owning a car but not related to the driving, τ_j is the annual tax and y_{it} denotes household income in period t . The parameter γ_i scales the utility of money relative to that of driving and it varies across households according to $\gamma_i \equiv \gamma'_z z_i$, where z_i is a vector of household demographics. For the primary results, I let $u^{\text{own}}(j) = \alpha'_0 q_j$, where q_j is a vector of observable characteristics for the car such as weight, engine power but not including fuel efficiency, e_j , which is restricted to enter the model through the cost structure.¹¹ This term shifts mean utilities of buying a given car in a way that is unrelated to the driving utility so as to better fit market shares.

In the second period, the household must choose how many kilometers to drive, x . The second-period utility is given by

$$u_{ij2} = \gamma_i \left(y_{it_2} + \delta_j^4 p_j^{\text{car}} - \frac{p_{jt_2}^{\text{fuel}}}{e_j} x \right) + \alpha_{1ij} x + \alpha_2 x^2,$$

where e_j is the fuel efficiency of car j in kilometers per liter, $p_{jt_2}^{\text{fuel}}$ is the price of fuel (gasoline or diesel depending on the fuel type of car j), and α_{1ij} is a parameter that affects the utility of driving an extra kilometer. This parameter is heterogeneous and correlated with demographics and car characteristics as follows:

$$\alpha_{1ij} \equiv \alpha_{10} + \alpha'_{1z} z_i + \alpha'_{1q} q_j + c_i.$$

The variable c_i is a time-constant random effect that is independent of z_i and q_j and captures heterogeneity in the utility of driving that is unobserved by the econometrician but observed by the household. The assumption that utility from driving is quadratic yields a computationally attractive form for optimal driving as we shall see. It implies theoretically a bliss point in driving but in the application, all households were far below this point. The coefficient α_2 has also been allowed to vary over i and j but the additional parameterization did not improve model fit so I chose the more parsimonious specification.

3.2 Solving the Consumer's Problem

In period t_2 when the household makes its VKT choice, x , it conditions on the purchased car. Thus, optimal driving maximizes u_{ij2} conditional on j . Interior solutions must

¹¹In future work, it would also be interesting to include information on parents' automobile choice where available in the registers as persistence in brand preference within a family has been documented in the literature (Anderson et al., 2013).

therefore satisfy the first-order condition;¹²

$$x = -\frac{1}{2\alpha_2} \left(\alpha_{1ij} - \gamma_i \frac{p_{jt_2}^{\text{fuel}}}{e_j} \right) \equiv x_{ij}^*(p_{jt_2}^{\text{fuel}}). \quad (3.1)$$

Thus, optimal driving is characterized by a linear equation, where car characteristics shift the level of driving and household demographics shift both the level and the price sensitivity of driving. In particular, note that the unobserved driving type, c_i , shifts the level of driving. The linear form conveniently allows me to relate the structural parameters to reduced-form regressions of VKT on the price per kilometer, defined as the fuel price divided by the fuel efficiency, $p_{jt_2}^{\text{fuel}}/e_j$, since the scaled parameters, $-\frac{\alpha_{1ij}}{2\alpha_2}$ and $\frac{\gamma_i}{2\alpha_2}$, are identified by the driving equation. This is also useful for finding good starting values.

When I insert the optimal driving rule from (3.1) back into the full utility function I obtain an expression that can be computed based on data. Due to the quasi-linearity, the income term, $\gamma_i(y_{it_{1i}} + \beta^4 y_{it_{2i}})$, does not vary over j and so can be dropped from the specification. Instead, income is allowed to enter through both the heterogeneous parameters, α_{1ij} and γ_i , to capture correlations with taste patterns and leisure activities. The expected utility of choosing car j is

$$\begin{aligned} u_{ij} = & -\gamma_i 4\tau_j + \gamma_i \left[1 - (\beta\delta_j)^4 \right] p_j^{\text{car}} + u^{\text{own}}(j) \\ & + \beta^4 \mathbb{E} \left\{ -\gamma_i \frac{p_{jt_2}^{\text{fuel}}}{e_j} x_{ij}^*(p_{jt_2}^{\text{fuel}}) + \alpha_{1ij} x_{ij}^*(p_{jt_2}^{\text{fuel}}) + \alpha_2 \left[x_{ij}^*(p_{jt_2}^{\text{fuel}}) \right]^2 \middle| p_{jt_1}^{\text{fuel}} \right\}. \end{aligned} \quad (3.2)$$

All that remains is to specify the household's expectations at time t_1 about fuel prices at time t_2 conditional on fuel prices at time t_1 s. In the literature, many implementations have used static expectations, whereby the expectation in (3.2) collapses to a single number. Gillingham (2012) uses a unit root and also allows consumers to use prices of futures on fuel in their forecast. He finds that it makes little difference to his results. I have implemented both static expectations, perfect foresight and a unit root with a drift. For the unit root, the expectation in equation (3.2) must be solved by numerical integration. I do this using Gauss-Hermite quadrature, which performs extremely well for univariate integrals. As it turns out, the specification of the fuel price expectations do not greatly impact my main results. There are two intuitive reasons for this; firstly, the variation in fuel efficiency in the choiceset is larger than the variation in fuel prices over time. Secondly, the quasi-linear utility function implies that consumers are risk neutral. In a model with diminishing marginal utility of money or credit constraints, concerns about fuel prices rising too much might push the household down to low levels of consumption and high curvature. The non-linearity of such a model could yield much greater differences depending on the expected fuel prices.

¹²At the estimated parameter values, the model only predicts strictly positive VKT for all households.

One simplification that the fuel price expectation structure has imposed in all cases is that the gasoline and diesel price processes are not modeled jointly by the households. The price of diesel has moved from 80.6% to 89.2% of the gasoline price over the period, but there are substantial fluctuation year to year (cf. Figure B.6). The forecasts from a bivariate time-series process of the two fuel prices, possibly including oil prices, might yield an improvement but I leave this for future work.

I will conclude the model section with a brief discussion of the assumptions imposed by the model. The quasi-linearity of the model affords a lot in terms of simplifying the model solution. An alternative interpretation, due to Bento et al. (2009), is that the model considers the problem of a household *renting* a car for four years; since the household pre-commits to selling the car again and there is no uncertainty about future car prices, the analogy is very clear. This simplification admits more complexity elsewhere. Moreover, curvature is more likely to make a difference for the decision about *when* to go on the new car market; households might choose to postpone car purchases simply due to the fear of becoming unemployed and receiving a large negative income shock. Since this is beyond the scope of this paper, I choose to focus on having a highly detailed model of the car choice conditional on entry. Instead, I rely on capturing some of these effects by allowing income to change the marginal utility of money and driving by including it in γ_i and α_{1ij} to capture some of these effects. This is similar to how many papers in the literature following Berry, Levinsohn, and Pakes (1995) have done it.

Computationally, the main challenge with the implementation of the model is the dimensionality of the choicest, $\mathcal{J}$. Avoiding aggregating cars has the advantage of clarity as well as precision in terms of calculating tax revenue and other counterfactual outcomes that rely on the precise characteristics of individual cars; such details might get lost in aggregation. The model has been implemented in `c`, which has yielded a considerable speedup over Matlab in particular due to parallelization and explicit utilization of the sparsity structure of $\mathcal{J}$ due to some car models not being available in all years.

4 Empirical Strategy

In this section, I first outline the econometric methodology and derive the likelihood function. I then discuss where the identifying variation is coming from in the data and comment on the implementation of the estimator. Finally, I outline how I simulate from the model and calculate counterfactual outcomes based on the estimated parameters.

4.1 Econometric Methodology

The econometric methodology follows Gillingham (2012). The dataset contains for each household the discrete car choice, d_i , and the continuous driving choice, x_i . Furthermore,

it contains the realized average fuel price over the household's driving period, $p_{jt_{2i}}^{\text{fuel}}$, and finally the vector of demographic variables, z_i . The subscript i in period t_{2i} is to remind the reader that there is cross sectional variation in the fuel price insofar as two households' periods do not perfectly overlap. Fuel prices also vary with j depending on the fuel type of the car. Other than that, the year of purchase gives the annual fuel prices that year and the choiceset and characteristics of the cars available in that year.

To obtain non-degeneracy of the model, an error term is added to both choice margins; an IID Gaussian measurement error to the optimal driving equation and an IID Extreme Value term to the conditional utility, u_{ij} . The observed driving for household i , x_i , is therefore written as

$$x_i = x_{id_i}^*(p_{d_it_{2i}}^{\text{fuel}}) + \eta_i, \quad \eta_i \sim \mathcal{N}(0, \sigma_x^2),$$

This means that the partial likelihood contribution for the observed driving is given by

$$f_x(x_i|\theta) = \frac{1}{\sigma_x \sqrt{2\pi}} \exp \left\{ -\frac{[x_i - x_{id_i}^*(p_{d_it_{2i}}^{\text{fuel}})]^2}{2\sigma_x^2} \right\}, \quad (4.1)$$

where the dependence of predicted driving, $x_{id}^*(\cdot)$, on the unobserved type, c_i , is subsumed.

For the type choice, the full utility for household i from choosing car $j \in \mathcal{J}_{t_1}$ becomes

$$\tilde{u}_{ij} = u_{ij} + \varepsilon_{ij}, \quad \varepsilon_{ij} \equiv \frac{1}{\lambda} \tilde{\varepsilon}_{ij} \quad \tilde{\varepsilon}_{ij} \sim \text{IID Extreme Value}.$$

I will discuss the scale parameter, λ , in greater detail below. The probability that car j maximizes household i 's utility is therefore given by

$$\Pr_i(j|\theta) = \frac{\exp(u_{ij}/\lambda)}{\sum_{j' \in \mathcal{J}_{t_1}} \exp(u_{ij'}/\lambda)}.$$

I will estimate one version of the model where $c_i = 0$ for all i . In that model, the full log-likelihood contribution for household i becomes

$$\ell_i^{\text{full}}(\theta) = f_x(x_i|\theta) \Pr_i(d_i|\theta).$$

In the general case, I will assume that $c_i \sim \mathcal{N}(0, \sigma_c^2)$ and the likelihood gets the typical integrated likelihood form similar to the mixed logit:

$$\ell_i^{\text{sim}}(\theta) = \int f_x(x_i|\sigma_c c; \theta) \Pr_i(d_i|\sigma_c c; \theta) d\Phi(c),$$

where Φ is the Gaussian cdf. The conditioning on the individual effect $c = \sigma_c c_i$ is made explicit in both $f_x(\cdot)$ and $\Pr_i(\cdot)$ in the equation as a reminder that it enters into α_{1ij} and thus in both optimal driving and choice-specific utilities. In this sense, the c_i variable

has the interpretation of a random effect. The univariate integral will be computed using Gauss-Hermite quadrature.¹³

The model has been implemented in the programming language `c` using Matlab’s interface, `Mex`. For optimizing the likelihood function, I have alternated between using a gradient based (quasi-Newton) and a gradient-free (Nelder-Mead) solver with semi-analytic numerical gradients (exploiting the linear structure of the random coefficients) and BHHH approximation of the Hessian due to Berndt et al. (1974).¹⁴

The logit scale parameter, λ , is not identified in the outset because the scale of utility can be moved up and down by α_2 . However, I found that the likelihood was more easy to manage numerically with a re-normalization setting $\alpha_2 := -1$ and estimating λ instead. Unfortunately, the likelihood function turned out to be extremely flat in the direction of λ . Instead, I estimated the model over a grid of λ -values and picked the λ that produced the best fit for the data while also giving sensible elasticities. If I allow the optimizer to choose λ freely, the optimizer terminates without convergence at a λ value of just over 100,000, at which point the model produces zero elasticities (to the fifth decimal) on all margins. I discuss this issue in greater detail in Appendix C and outline a potential model extension that would allow me to estimate the scale parameter jointly with the remaining parameters. This approach involves estimating car type fixed effects vis-a-vis Berry, Levinsohn, and Pakes (1995).

4.2 Identification

The model relies on both cross-sectional and time-series variation as well as within-household variation. The variation in fuel prices and the choicest set over time identifies how households substitute between available cars under different circumstances. The parameters in the utility function are moreover tied down by there being two observed outcomes for each household; the discrete car choice and the continuous driving choice. In that sense, the model intuition is not far from a Heckman selection model; the exclusion restrictions are the fuel prices at the time of purchase, the choicest set available at the time of purchase as well as the structure of the model. In essence, the model imposes the strict cross-sectional restriction that consumers value money in a similar fashion when making car purchase decisions and driving decisions. The driving decisions should be thought of as covering several years and not the daily driving decisions, where households can switch purchases over the week days in response to daily variation in prices.

¹³For the results presented here, only 8 nodes were used. Future work is under way using more nodes. Comparing quadrature with simulation using simple, smooth functions and univariate integrals, it was found that quadrature attains the same level of precision as simulation using five to ten times more evaluations of the integrand. This point was also highlighted by Dubé, Fox, and Su (2012) and Judd and Skrainka (2011).

¹⁴Whenever the gradient-based solver would get stuck, unable to improve the likelihood along the gradient direction, the Nelder-Mead solver proved useful in breaking free of the local optimum.

There has been a considerable increase in fuel prices in my sample period which, as discussed in section B.3, has arguably been driven by world-market factors. To leverage variation from changes in the tax rates over the period, I have explicitly coded the annual tax rates, τ_j , and included those in the model. The characteristics of available cars have also changed substantially due to technological progress over the ten-year period, which has made cars more fuel efficient for any given level of car weight. These sources of variation are fine to the extent that the changes in car makers' portfolios is driven by tax policies or demand side effects in other, bigger markets. However, there may of course be common trends in demand across countries leading to this. For example, urbanization patterns across many developed countries have followed similar patterns with more households moving to the urban areas. My work distance variable will capture such trends, so in terms of the driving equation, I am more worried about correlated trends in leisure driving. In related work, Gillingham and Munk-Nielsen (2015) explore many different sources of variation to estimate the medium-run, 1-year elasticity of driving with respect to fuel prices and find a central elasticity of -0.30 with household fixed effects. This is very close to what I find when I take into account selection, even though I don't include fixed effects.

In terms of the discrete car choice, the model can be thought of as a mixed logit with a particular functional form imposed on the choice-specific utilities. In much of the literature on car choice the driving equation is not considered but there will often be either sophisticated nesting structures on the logit errors or car specific fixed effects in the Berry, Levinsohn, and Pakes (1995) or both (Grigolon, Reynaert, and Verboven, 2015). In future research, it would be interesting to see these features integrated in a discrete-continuous choice model. I propose such an extension in Appendix C but leave the estimation to future research.

4.3 Simulating From the Model

As with most structural models, it is essential to be able to simulate counterfactual behavior from the model. Essentially, we want to compute simple statistics characterizing the final market outcome of making changes to taxes, prices or the characteristics of cars. These outcomes might be the CO₂ emitted, tax revenue, the average fuel efficiency, etc. Formally, suppose we are interested in some outcome ω_{ij} . Then define the average expected outcome as

$$\tilde{\mathbb{E}}(\omega|\theta) \equiv \frac{1}{N} \sum_{i=1}^N \sum_{j \in \mathcal{J}_i} \text{Pr}_i(j|\theta) \omega_{ij}. \quad (4.2)$$

This is the average (over households) weighted average (over available choices weighted with conditional choice probabilities) outcome.

Note that in the computation of (4.2), I need to take a stand on the stochastic variables

in the model; η_i , ε_{ij} and c_i . The measurement error is set to zero, $\eta_i := 0$. Since I am weighting by conditional choice probabilities, the expression is implicitly an expectation over ε_{ij} . Lastly, c_i is set to zero for all households; instead, one could integrate out the random effect unconditionally, but given the quasi-linearity and the linear driving equation, it is unlikely that such efforts would yield very different results.¹⁵ Standard errors have not been computed for the expected outcomes.

Two examples of outcomes of particular interest require an extra comment. Firstly, the CO₂ emissions; These are calculated using the kg of CO₂ that is emitted by the combustion of a liter of each fuel,¹⁶ yielding the following CO₂ emissions (in kg) conditional on choosing car j and realized fuel price p_{jt2i}^{fuel} ,

$$\text{CO}_{2,ij} \equiv \left(\mathbf{1}_{\{j \text{ is gas}\}} 2.392^{\text{kg/l}} + \mathbf{1}_{\{j \text{ is diesel}\}} 2.64^{\text{kg/l}} \right) \frac{x_{ij}^*(p_{jt2i}^{\text{fuel}})}{e_j}.$$

Setting $\omega_{ij} := \text{CO}_{2,ij}$ in (4.2) gives the average expected CO₂ emissions. The analysis emphasizes CO₂ emissions to focus on the environmental aspects but might as well have emphasized fuel use; the two are proportional.

Secondly, the tax revenue can be calculated conditional on car purchase and subsequent usage. The conditional total tax revenue, τ_{ij}^{total} , is given by

$$\tau_{ij}^{\text{total}} \equiv \tau_j^{\text{fuel}} \frac{p_{jt2i}^{\text{fuel}}}{e_j} x_{ij}^*(p_{jt2i}^{\text{fuel}}) + \tau_j^{\text{reg}}(p_{tj}^{\text{car}}) + 4\tau^{\text{annual}},$$

where $\tau^{\text{reg}}(\cdot)$ gives the registration tax and τ_j^{fuel} is the fuel taxes in pct. of the total fuel price. Setting $\omega_{ij} := \tau^{\text{total}}$ in (4.2) gives the average expected tax revenue for the government.

Lastly, following Small and Rosen (1981) and Gillingham (2012), the model yields the usual “logsum” welfare measure defined as

$$\text{CS} \equiv \frac{1}{N} \sum_{i=1}^N \log \left[\sum_{j \in \mathcal{J}_i} \exp(u_{ij}) \right], \quad (4.3)$$

which can be used to evaluate the welfare impacts on consumers from changing parameters of the choiceset such as car characteristics or tax rates. It should be noted though that since there is no outside option, this welfare measure does not take into account that

¹⁵The reason why the random effect makes a difference in estimation is that here, information from both periods are employed simultaneously and thus the simulated likelihood will apply the highest weight to the region of the support of c_i that best rationalize household i ’s two decisions. An alternative approach that might yield different results would be to integrate out c_i *conditional on choices*; this is in line with the approach outlined in Train (2009, ch. 11). That strategy has some similarities with a latent class model where one can compute the probability that $c_i = c_q$ for some quadrature node q , and use these weights in counterfactual simulations.

¹⁶These numbers come from www.ecoscore.be (and are confirmed by www.environment.gov and www.epa.gov).

households may choose not to own a car.

5 Results

In this section, I present the estimation results. I start by presenting the structural parameter estimates and discussing these. To assess the validity, I discuss the driving equation and relate it to a partial estimation of the driving parameters alone as well as to what has been found in the literature. To get a better intuitive grasp of the model behavior at the estimated parameters, I compute a number of relevant outcomes and calculate the elasticities of these with respect to exogenous variables. Finally, I discuss robustness and consider alternative specifications of the fuel price expectations process.

Table D.1 shows the structural estimates from the preferred specification allowing random effects ($c_i \neq 0$) and where consumers have perfect foresight with respect to fuel prices. I will discuss the fuel price expectations later. The coefficients have the expected signs; households with higher work distances tend to drive more (α_{1z} -coefficients are positive) and be more price-responsive in their driving (γ_z -coefficients also positive, increasing the magnitude of the utility of money).¹⁷ Urban households tend to drive their cars less and older households also drive less. Heavy cars tend to be driven more intensively as indicated by $\alpha_{1,\text{weight}}$ and $\alpha_{1,\text{weight}^2}$ both being positive; this is consistent with car weight proxying for unobserved luxury characteristics. The term, α_0 , captures utility from the car ownership that are unrelated to driving. The parameters entering into α_0 tend to be very large, but recall that they should be divided by the λ -value of 10,000. The diesel coefficient ($\alpha_{0,\text{diesel}}$) is negative; this indicates that there is some characteristic about diesel cars that keeps households from buying them even though their other characteristics make them more attractive than a given gasoline car. Finally, note that the dispersion in the unobserved driving type, c_i , is estimated to be 16.09, while the standard error on the driving equation measurement error is 21.95. This indicates that the endogenous selection of car type based on other factors still play a considerable role even though work distance is accounted for.

Recall from section 4 that the VKT equation can be estimated separately, using the partial likelihood function from equation (4.1). Table 5.2 shows the elasticities of VKT with respect to the fuel efficiency, the weight of the car, and the fuel price.¹⁸ Elasticities are computed numerically for each observation using finite differences and reports both

¹⁷Gillingham and Munk-Nielsen (2015) explore precisely this feature of the data, finding that it high-driving households switch from driving to car to using other modes of transport when fuel prices increase. The behavior is consistent with a model of switching costs in changing transport mode to work from private car to public transportation.

¹⁸The estimated linear equation regresses VKT on demographics and car characteristics as well as demographics interacted with the price per kilometer, defined as the fuel price (gasoline or diesel depending on the car) divided by the fuel efficiency.

Table 5.1: Estimated parameters

Fixed Parameters				
Parameter	Value			
β	0.95			
λ	10000			
α_2	-1			
Model: Perfect foresight, random effects.				
General Parameters				
Parameter	Estimate		t	
σ_x	16.093		(69.12)	
σ_α	21.951		(31.77)	
Demographics				
	γ_z		α_{1z}	
Parameter	Estimate	t	Estimate	t
Constant	47.596	(35.22)	74.927	(14.88)
Age	-8.447	(-18.97)	8.901	(8.71)
Age ²	7.363	(15.88)	-15.168	(-14.39)
Work distance, male	8.170	(18.95)	17.889	(69.45)
Work distance, female	1.079	(19.46)	9.684	(108.20)
Income	-9.457	(-31.44)	-8.768	(-39.94)
Number of kids	1.453	(11.65)	-0.458	(-2.93)
Urban dummy	-0.210	(-1.48)	-1.412	(-10.09)
Car Parameters				
Parameter	Estimate		t	
$\alpha_{0,\text{weight}}$	124074.734		(41.91)	
$\alpha_{0,\text{weight}^2}$	-5009.689		(-5.67)	
$\alpha_{0,\text{kw}}$	-413.653		(-25.53)	
α_{0,kw^2}	5.114		(46.83)	
$\alpha_{0,\text{displace}}$	-194.172		(-0.15)	
$\alpha_{0,\text{displace}^2}$	4976.559		(13.12)	
$\alpha_{0,\text{diesel}}$	-4235.595		(-24.99)	
$\alpha_{1,\text{weight}}$	18.876		(3.12)	
$\alpha_{1,\text{weight}^2}$	10.189		(5.64)	

Table 5.2: Elasticities of VKT From the Structural Model — Partial Estimates and Preferred Specification

Partial Likelihood (using $f_x(\cdot)$)			
	Fuel efficiency	Weight	Fuel Prices
Mean	0.718	1.323	−0.725
Std.	0.426	0.339	0.431
Preferred Specification (using $\ell^{\text{sim}}(\cdot)$)			
	Fuel efficiency	Weight	Fuel Prices
Mean	0.279	0.858	−0.282
Std.	0.085	0.171	0.086

the average and standard deviation of the elasticity across observations.¹⁹ The elasticity of driving with respect to the fuel efficiency from the partial model is −72.5%. This central elasticity, when properly identified, is what Small and Van Dender (2007) refer to as the *rebound effect*. This estimate is fairly close to the approximately −80% that Frondel, Peters, and Vance (2008); Frondel, Ritter, and Vance (2012) find using German data. The estimate from the full model accounting for selection, however, is −28.2%. Gillingham (2012) finds a bias in the same direction but smaller in magnitude with a rebound effect of −21% dropping to −15% when selection is accounted for. Bento et al. (2009) find a mean elasticity of −35% which also controls for selection. Note that the elasticity with respect to fuel price and fuel efficiency are the same (except for the sign and direction) since they only enter the model together in the price per kilometer.²⁰ Finally, the estimates in Table 5.2 indicate that weight has a large effect on driving with an increase in weight of 1% being associated with an increase in driving of 0.858%. This implies that to understand the impacts of a car reform on driving and thereby emissions, it is not enough to just focus on the fuel efficiency; the weight of the chosen vehicles can also have strong effects on the final driving.

To get a better grasp of the implied behavior by the structural elasticities, Table 5.3 shows a range of economic outcomes simulated from the model in column (1) by the method described in Section 4.3. The table also shows elasticities of these outcomes with respect to four different variables in columns (2)–(5), computed using finite differences.

Column (2) shows the relative change in each expected outcome when the fuel efficiency of each car in the choicest is increased by 1%. For the fuel efficiency of the chosen vehicles, this has an elasticity of 0.90 so that the average expected fuel efficiency is 0.9% higher. This implies that when technological progress makes cars more fuel efficient, households substitute away some of this for other characteristics; the weight increases by 0.09%, the

¹⁹The dispersion in the elasticity is driven by the dispersion in the computed coefficient $\hat{\gamma}_i \equiv \hat{\gamma}'_z z_i$.

²⁰Gillingham (2012) allows e_j to shift the mean u_{ij} by putting it in the term $\alpha'_0 q_j$ in (3.2).

Table 5.3: Structural Elasticities — Quasi, Perfect Foresight, Random effect

	Levels	Elasticities			
	(1) Baseline	(2) Fuel efficiency	(3) Weight	(4) Fuel prices	(5) O95 prices
Consumer welfare					
CS	114970.09	0.25	1.29	-0.25	-0.20
Total taxes					
E(total taxes)	146623.83	0.08	0.40	-0.08	-0.06
Ownership tax					
E(Regtax revenue)	106556.44	0.23	0.27	-0.23	-0.14
E(Owntax revenue)	11093.62	0.29	0.31	-0.28	-0.16
Fuel tax					
E(O95 revenue)	25115.85	-0.49	0.63	0.50	-0.02
E(Diesel revenue)	3857.92	-0.89	2.98	0.88	2.33
Driving/fuel use					
E(VKT)	79663.89	0.30	1.01	-0.30	-0.19
E(litre O95)	4340.92	-0.49	0.63	-0.50	-1.01
E(litre D)	891.32	-0.89	2.98	-0.12	2.33
E(litre D—urban)	188.02	-0.86	3.04	-0.14	2.24
E(kg CO2)	12736.56	-0.57	1.06	-0.43	-0.39
Characteristics					
E(fe)	15.92	0.90	-0.00	0.10	0.20
E(we)	1.70	0.09	1.15	-0.09	-0.04
E(kw)	77.08	0.24	0.13	-0.23	-0.25
E(displace)	1.65	0.18	0.12	-0.18	-0.16
E(% diesel)	18.49	-0.16	1.86	0.15	2.33
E(% diesel—urban)	3.89	-0.14	1.88	0.13	2.24

The model is quasi-linear with perfect foresight and random effects (σ_α is estimated).

The baseline column is expected outcomes, all other are elasticities.

(1): baseline 2006 scenario, (2) fuel efficiency up by 1%, (3): weight up by 1%,

(4): all fuel prices up by 1%, (5): only O95 up by 1%.

Counterfactuals are run on 2006 data.

engine power (kW) by 0.24% and the diesel share falls by 0.18%. More interesting, the CO₂ elasticity is -57%, so that a 1% improvement in fuel efficiency does not give a full 1% improvement in CO₂ emissions. This is partly due to consumers switching away from efficient cars and partly due to consumers driving longer since the cost of driving an extra km is now lower. This result has huge implications for climate policy since it means that in order to reduce CO₂ emissions by 1%, the required improvement in fuel efficiency is approximately 1.75%.

Column (3) shows the effects of increasing the weight of all cars by 1%. This increases VKT by 1.01% and CO₂ by 1.06%. Note that the elasticity of driving with respect to car weight was even stronger in the partial equation, indicating that selection is at play.

Column (4) shows the effects of increasing the real fuel price at the pump by 1%.²¹ The most notable result here is that tax revenue *falls*, indicating that the Danish taxes are at the wrong side of the Laffer curve's top; While revenue from fuel taxes increase, revenue from the registration and the ownership tax fall by much more because households buy different types of cars. CO₂ emissions fall by 0.41%, which should be compared to the intensive-margin response of 0.28% implied by Table 5.2.²²

Finally, column (5) increases gasoline prices by 1% but keeps diesel prices constant. The result is a 2.33% change in the probability of purchasing a diesel car (and thus of the diesel market share). This gives a first indication that the diesel market share is highly sensitive to cost differences.

Based on the elasticities of CO₂, tax revenues and welfare with respect to fuel prices, it is possible to compute the marginal cost of CO₂ reductions from a fuel tax. Back of the envelope calculations indicate, that a reduction of one ton of CO₂ would cost society 7389.90 DKK.²³ This number is far above the Social Cost of Carbon of 260 DKK per ton as suggested by the US Environmental Protection Agency. The high cost is perhaps not surprising given how high the tax level is in Denmark.

To examine robustness, the model has also been estimated assuming static expectations and a unit root as described in section 3.2. These different specifications gave quite similar results in terms of elasticities and implications for the counterfactual simulations

²¹Note that to obtain this using taxes, one would have to take into account supplier responses. For the US, Marion and Muehlegger (2011) find a pass-through to consumers of almost 100% but given the substantially higher taxes in Denmark, that conclusion might not be valid here. Nonetheless, I abstract from the question of passthrough.

²²Table 5.2 conditions on car choice so any given relative change in driving will produce the same relative change in fuel consumption and thus in CO₂ emissions.

²³The required change in fuel prices to reduce CO₂ by 1 ton is approximately $\Delta p = (\mathcal{E}_{\text{CO}_2,p} \text{CO}_2/p)^{-1} = (0.43 \frac{12.7 \text{ ton}}{8.5 \text{ DKK/l}})^{-1} \cong 1.55 \text{ DKK/l}$. This implies an approximate change in consumer surplus and taxes of

$$\begin{aligned} \Delta \text{CS} &= \text{CS} \times \mathcal{E}_{\text{CS},p} \times \frac{\Delta p}{p} = 114,970.09 \times -0.25 \times \frac{1.55}{8.5} \cong -5248,13 \text{ DKK} \\ \Delta \text{Taxes} &= 146,623.83 \times -0.08 \times \frac{1.55}{8.5} = -2141,78 \text{ DKK}. \end{aligned}$$

so the perfect foresight model was chosen. The results with static expectations are shown in Appendix D.1; the elasticities of the relevant quantities relatively unchanged compared the corresponding ones for the model with perfect foresight, although the driving response is -0.39 instead of -0.30 . This implies a higher reduction in driving, and the cost per ton of CO_2 for the fuel tax implied by these estimates is correspondingly lower: 5843.02 DKK. The key to understanding the difference between the parameters estimated under the two sets of assumed price expectation formation is the realized movements in fuel prices (see Figure B.3); prices have been increasing throughout the period and the likelihood conditions on the same car and driving choices. Therefore, if consumers knew that prices would increase yet did not choose an even more fuel efficient car to curb fuel costs, it must be because they valued the fuel savings less relative to the other car characteristics. I have chosen the perfect foresight model as the preferred specification because the in-sample fit of the diesel share is better (Figure D.3). However, I think that for future research it might be more fruitful to focus on modeling the time-series development in the relative price of gasoline and diesel; Figure B.6 shows that the relative price of diesel to gasoline has gyrated around an increasing trend and gyrations appear to show up in the predicted diesel share. Figure D.2 illustrates that the model's over- and under-predictions seem to be correlated with the gyrations.

6 Counterfactual Policy Simulations

In this section, I present a sequence of counterfactual policy simulations. I start with a discussion of the model structure and assumptions and what they imply for the applicability of the counterfactuals. I then present a counterfactual simulation, implementing the out-of-sample 2007 reform in-sample. Next, I assess the role of the 1997 reform in driving the increase in diesel cars in Denmark. Finally, I present a counterfactual exploring the diesel share in absence of discriminatory ownership and fuel taxes.

6.1 External Validity

The strength of the model is in understanding how households trade off between available cars in the choicest set in characteristic space and how this interacts with driving decisions. In that sense, the model is well-suited for understanding how car tax policies feed into driving behavior and the related externalities. The high-dimensional choicest set makes the model precise in terms of modeling the tax system and leveraging policy variation. However, the computational cost of this dimensionality is that the model conditions on entry into the new car market. This means that all the simulated effects are for the average household; the model is uninformative as to changes in the number of households (or cars per household). Moreover, restricting the model to new cars effectively eliminates

real-world substitution alternatives in the form of the used-car market and the outside option of not owning any car. Ignoring substitution options for the consumers will inflate my estimate of the consumer loss related to increasing taxation.²⁴

The second main restriction of the model is that supply side responses to the proposed reforms are ignored, i.e. assuming a 100% passthrough of taxes. In reality, profit maximizing car sellers in oligopolistic competition will likely change the relative prices of cars in their portfolios. In defense of this assumption, Adamou, Clerides, and Zachariadis (2013) find little difference between their simulation results when they use their estimated supply side pricing function or simply assuming 100% passthrough in a European context. For fuel taxes, Gallagher and Muehlegger (2011) find that passthrough in the US to consumers is approximately 100%. Moreover, given the small size of the Danish market relative to other European countries, auto makers are unlikely to change their production to cater to Denmark.

6.2 The 2007 Reform: Model Validation and Policy Evaluation

As described in 2.1, the 2007 reform was a feebate, meaning that it gives a rebate to green cars and puts a fee on inefficient cars. The *pivot point* of the reform, differentiating green cars from dirty ones, was set to 16 km/l for gasoline cars and 18 km/l for diesel cars. Recall that 2007 is not in the estimation sample because driving information is only available for a small number of cars purchased in this year.

Table 6.1 shows the implications of implementing the 2007 feebate in 2006. Most importantly, the diesel market share goes up from 18.5% to 24.5%, an increase of 32.3%.²⁵ The true response to the 2007 reform was an increase in the diesel share of 46.0%. In other words, the model can explain two-thirds of the relative shift in the diesel share. Similarly, the model predicts the average fuel efficiency to increase by 7.04% whereas the actual response to the reform was 5.73%. In this case, the model overshoots but as Figure (2.1) illustrates, the fuel efficiency continues to increase in the following years, increasing by an additional 7.63% in 2008. I view these as good out-of-sample fits.²⁶

Regarding the predicted environmental impact of this reform, the average expected

²⁴This is because in the model, the consumer has no choice but to shift around in the choicest. In a more realistic model, the consumer can also choose the outside option or used cars. Instead of being forced to absorb higher taxes, the consumer has the option of not owning a car. Since this alternative is unaffected by fuel taxes, the consumer surplus measure in (4.3) will drop less when that alternative is available.

²⁵One important note to make in this regard is that the diesel share in the sample in 2006 is 18.5% whereas in the full population it is 21.8%. As discussed in appendix B.1, this is due to diesel cars being over represented in the car types that are only purchased by very few households and therefore dropped from the sample. I expect that these niche cars would be hard to fit in this model framework.

²⁶I have been unable to find data to produce a graph comparing fuel efficiencies across European countries similarly to how Figure 2.2 shows diesel penetration rates. My impression is that change in fuel efficiency in 2007 for Denmark is still uniquely large but not as different from the rest of Europe as is the case for the diesel share.

Table 6.1: Counterfactual Simulations — The 1997 and 2007 Reforms

	(1)	(2)	(3)	(4)
	Baseline	1997	2007	Internalization
Consumer welfare				
CS	114,970.09	99,607.67	115,989.89	115,569.79
Total taxes				
E(total taxes)	146,623.83	176,422.24	134,398.55	146,854.69
Ownership tax				
E(Regtax revenue)	106,556.44	117,238.77	98,175.80	107,363.31
E(Owntax revenue)	11,093.62	26,613.14	9,813.59	9,131.02
Fuel tax				
E(O95 revenue)	25,115.85	31,519.34	21,695.96	23,698.46
E(Diesel revenue)	3,857.92	1,050.99	47,13.20	6,661.89
Driving/fuel use				
E(VKT)	79,663.89	78,391.96	78,740.62	79,518.56
E(litre O95)	4,340.92	5,447.67	3,749.84	4,095.94
E(litre D)	891.32	242.82	1,088.92	1069.70
E(litre D—urban)	188.02	50.40	230.59	230.87
E(kg CO2)	12,736.56	13,671.87	11,844.36	12,621.51
Characteristics				
E(fe)	15.92	14.69	17.04	16.12
E(we)	1.70	1.73	1.64	1.71
E(kw)	77.08	89.93	70.07	76.64
E(displace)	1.65	1.86	1.54	1.65
E(% diesel)	18.49	4.97	24.48	23.28
E(% diesel—urban)	3.89	1.03	5.18	5.03

The counterfactuals are run on data for 2006.

1997: The green ownership tax is replaced with the weight based annual tax.

2007: The 2007 feebate reform is implemented on 2006 data.

Internalization: Annual and registration taxes for diesels are set in the same way as gasoline cars but the diesel price is increased by 1.923 DKK/l.

CO₂ emissions fall by 892.2kg or 7.0%. Some of this comes through the intended channel of improved fuel efficiency which increases by 7.0%, but recall from table 5.3 that this only translates into approximately $0.57 \cdot 7.0\% = 4.0\%$ reductions in CO₂. In particular, the reform as a by-product reduces weight by 3.5% which translates into less driving, yielding an additional $1.01 \cdot 3.5\% = 3.5\%$ in CO₂ reductions. In other words, the reform's impact on the weight of the chosen vehicles is almost as important as the intended impact via fuel efficiency.

In terms of welfare, the 2007 reform increased consumer surplus but decreased taxes by much more. Even accounting for the lowered driving and thus lower non-CO₂ externalities, the societal cost of the predicted reduction in CO₂ was 11,886.99 DKK/ton. This is a 60.9% higher cost per ton of CO₂ than that of the fuel tax, cf. section 5, and even further from Social Cost of Carbon of 260 DKK/ton. It is not uncommon to find high implied costs of CO₂ savings in the literature, e.g Beresteanu and Li (2011) and Huse and Lucinda (2013), although my estimates are exceptionally high. However, the feebate is asymmetric with a higher rebate than fee; in light of Adamou, Clerides, and Zachariadis (2013) it is not surprising that it is in-effective.

Given that the model fits the shift to diesels, the next question is which part of the policy design led to this shift. The pivot point of 16 km/l for gasoline and 18 km/l is an obvious candidate given that the median difference between gasoline and diesel cars is higher than 4 km/l. I therefore implement a counterfactual where the pivots instead are set to 16 km/l and 20 km/l. The results of this counterfactual are shown in Table D.2; here, diesel share only increases marginally by 6.2%. Moreover, this alternative version of the reform yields 91% of the CO₂ reductions of the actual reform with almost identical consumer surplus and tax revenue. This provides evidence that the CO₂ reductions achieved by the feebate were not simply due to a shift to diesel cars.


6.3 The 1997 Reform: The Role of Taxation in the Dieselization

The 1997 reform changed the annual tax from being based on the weight of the car to being based on the fuel efficiency (see section 2.1). However, cars first registered before July 1st 1997 still follow a weight-based scheme. In this counterfactual, I compute the annual tax for all cars based on that scheme instead of the actual, fuel efficiency based scheme.²⁷ The average expected outcomes in 2006 under this counterfactual are shown in column (2) of Table 6.1. Figure 6.1 shows the predicted diesel share year by year in the sample.²⁸ The results show that while the diesel share would still have increased,

²⁷There might be many other counterfactuals equally interesting as the alternative to the 1997-reform; if for example the rates were changed over time to encourage scrapping of vintages from before 1997. That does not, however, appear to be the case.

²⁸The in-sample fit of the diesel share (the "Predicted" curve in Figure 6.1) fluctuates around the observed diesel share. The deviations are timed along with the movements in the relative price of diesel to gasoline (Figure B.6). One way to improve the fit might be to add a bivariate forecast, since the

Figure 6.1: Predicted Dieselization From the Baseline Model vs. the Weight Tax Counterfactual.



the increase would have been substantially lower. In 2006, the predicted share is 4.97%, which is substantially below the baseline of 18.49%. The reason for this difference is that the post-1997 tax regime rewards high fuel efficiency while the pre-1997 regime punished heavy cars. Since diesel cars are inherently more fuel efficient and tend to be heavier, they are likely to benefit from this. Already in 1997, it is clear that the new tax scheme favors diesel cars; for the average car in the choicset in 1997, the actual annual tax of a diesel car was 8.0% higher than the average annual tax for a gasoline car. However, under the counterfactual, pre-1997 regime the diesel car would be paying an 18.8% higher annual tax. Under the counterfactual, the predicted diesel share is substantially below the predicted under the actual tax regime (Figure 6.1) and this is driven mainly by this difference moving even further in favor of gasoline cars over the period.²⁹ The fact that the difference in the average annual tax increases over time also helps to explain why the response in the diesel share following the 1997 reform in Figure 2.2 is not a drastic shift as is the case for the 2007 reform.

6.4 The “Optimal” Diesel Share

For both the 1997 and 2007 reforms, I have found that the reforms were both misaligned in their differential treatment of diesel and gasoline cars, causing a change in the status-quo diesel share. Given that there are discriminatory elements both in fuel taxes and

relative deviations appear to be strongly mean-reverting around a trend.

²⁹By the end, however, the average diesel car in the choicset would have paid 61.8% more, had it followed the old scheme, while the actual, post-1997 annual tax only imposed a 23.4% higher annual tax on diesel owners.

ownership taxes, I explore the question: what would be the free-market outcome if the only discrimination in taxes was due to differences in externalities? To answer this question, I implement a counterfactual on the 2006 data. The only source of differences in externalities between diesel and gasoline cars is related to the fuel usage; the burning of diesel fuel emits slightly more CO₂ and emits harmful local air pollutants that gasoline does not.³⁰ Therefore, I first equalize ownership taxes, setting those for diesel cars equal to those for gasoline cars. Fuel taxes are not equal in the outset since gasoline has a higher fixed component of the taxes (see Appendix B.3.2). Therefore, I first equate fuel taxes by increasing the fixed component for diesel fuel up to the level gasoline and then add an additional per-liter tax equal to the per-liter external cost. The estimates of marginal external costs are taken from DTU Transport (2010) (see Appendix B.2). Assuming a 100% passthrough to consumers, I can simulate whether the diesel market share would be above or below the baseline level for Denmark in absence of discrimination — this exercise is similar to internalizing an externality using a Pigovian tax, except that the baseline gasoline tax is not necessarily optimal. In this sense, I do not claim to find the optimal diesel share but rather an improvement over the status quo.

The results are shown in column (4) of table 6.1. The central conclusion is that the predicted diesel share increases by 25.9% based on the 2006 diesel share (from 18.49% to 23.28%). This puts the predicted counterfactual diesel share between the 2006 and 2007 levels. Note that any incomplete passthrough would directly dampen this effect. An interesting additional conclusion that can be drawn from this counterfactual is that the proposed policy appears to represent an unambiguous improvement; Consumer surplus and tax revenue go up, CO₂ emissions go down and VKT also goes down (so externalities from congestion and accidents also decrease). However, these improvements are very small economically. This counterfactual indicates that when the added externalities of diesel cars are priced (subject to the externality prices that I have used), the added value of those cars (in terms of efficiency, for example) relative to their price makes them a valuable part of the car fleet.

7 Conclusion

In this paper, I estimate a structural discrete-continuous model of car choice and usage, allowing endogenous selection into car types based on expected future driving. The model is estimated using high quality full population register data for Denmark covering 1997–2006. To validate the estimates, I exploit the Danish car taxation reform of 2007 which prompted clear changes in new car type decisions immediately, unique to Denmark, in

³⁰In 2012, the World Health Organization moved diesel fumes to the list of substances that are known to cause lung cancer. There is regulation in place, effectively requiring diesel cars to be fitted with particle filters to reduce this type of pollution. These are taken into account by the external cost estimates.

particular in the diesel market share. Implementing the 2007 reform counterfactually in 2006, I find that the model is able replicate the strong responses to the reform in terms of the diesel share and the fuel efficiency.

A consistent finding is that Danish households have responded very strongly to the tax incentives given by the 1997 and the 2007 reform. The implication is that both reforms were highly cost-ineffective ways of obtaining CO₂ reductions compared to a fuel tax, mainly due to foregone tax revenue. A central mechanism behind this is that according to simulations from the model, a 1% technological increase in the fuel efficiency of all cars only translates into a 0.57% reduction in CO₂ emissions; this is partly due to households substituting these fuel savings away for larger, more luxurious cars and partly due to the *rebound effect*, whereby households being pushed towards more efficient cars in turn drive them more intensively (at an elasticity of -30%). This greatly limits the effectiveness of environmental policies. Additionally, my results indicate that the effects of car taxes on driving that work through the weight of the chosen car may be at least as important as those working through the fuel efficiency.

To evaluate the two tax reforms of the period, I compare their cost-effectiveness to a fuel tax. I find that fuel taxes are much more effective. However, the cost per ton of CO₂ is still many times larger than the social cost of carbon, possibly due to the high level of taxes in Denmark in the outset. In particular, I find that increasing fuel taxes may *lower* tax revenue if they are increased; while they do increase fuel tax revenue, this is offset by an even larger drop in car taxes as consumers shift away from the luxury segment.

Another finding is that the reforms were responsible for most of the increase in the diesel share that occurred in my sample period. In particular, the Danish feebate reform in 2007 could have been designed differently to yield 91% of the CO₂ reductions but with only a minor increase in the diesel share. Nevertheless, I also show that the societal gains from diesel cars outweigh their negative aspects and that the diesel share in 2006 is close to the optimal level for the Danish setting.

Appendix

A Notation and Core Equations

This section is meant as a quick reference to give an overview of the model and the notation used in this paper.

The notation is as follows,

- j — car type (e.g. 2003 Volvo V70 Turbo Diesel),
- d_i — the chosen car type by household i ,
- x — vehicle kilometers travelled (VKT, a generic decision variable),
- x_i — the observed driving for household i (conditioning on d_i),
- $x_{ij}^*(p^{\text{fuel}})$ — the optimal driving rule,
- e_j — fuel efficiency of a car of type j in km/l,
- p_{tj}^{car} — price of a new car of type j in year t ,
- p_{tj}^{fuel} — fuel price (the subscript j is there to distinguish diesel or octane),
- γ_i — utility of driving relative to outside consumption (household-specific),
- z_i — household attributes correlated with driving utility,
- y_{it} — household income in period t ,
- β — discount factor (fixed at 0.95),
- δ_j — vehicle-specific depreciation rate (e.g. 0.8),
- α_{1ij}, α_2 — utility from driving is quadratic in VKT with these coefficients,
- α_0 — coefficients on q_j ; Utility from car j that is not related to driving,
- ε_{ij} — IID extreme value type II shock (to the car type choice utility),
- η_i — measurement error in the VKT equation,
- ζ — coefficients used in the linear interpretation of optimal driving.

The full utility can be written as

$$u_{ij} = \gamma_i \left[1 - (\beta \delta_j)^4 \right] p_{jt_{1i}}^{\text{car}} - 4\gamma_i \tau_j + u^{\text{own}}(j) \\ + \beta^4 \mathbb{E} \left\{ -\gamma_i \frac{p_{jt_{2i}}^{\text{fuel}}}{e_j} x_{ij}^*(p_{jt_{2i}}^{\text{fuel}}) + \alpha_{1ij} x_{ij}^*(p_{jt_{2i}}^{\text{fuel}}) + \alpha_2 \left[x_{ij}^*(p_{jt_{2i}}^{\text{fuel}}) \right]^2 \right\}.$$

where

$$\begin{aligned}\gamma_i &= \gamma'_z z_i, \\ u_{ij}^{\text{drive}}(x) &= \alpha_{1ij}x + \alpha_2 x^2, \\ \alpha_{1ij} &= \alpha_{10} + \alpha'_{1z} z_i + \alpha'_{1q} q_j + c_i, \quad c_i \sim \mathcal{N}(0, \sigma_c^2).\end{aligned}$$

The driving rule, $x_{ij}^*(p_{jt}^{\text{fuel}})$, is given by

$$x_{ij}^*(p_{jt}^{\text{fuel}}) = -\frac{1}{2\alpha_2} \left(\alpha_{1ij} - \gamma_i \frac{p_{jt}^{\text{fuel}}}{e_j} \right).$$

In the estimation, z_i contains mean spouse age, age squared, work distance for both spouses, real gross income, the number of kids and a dummy for living in a major urban area (Copenhagen, Odense, Aarhus or Aalborg). The characteristics, q_j , are vehicle total weight, engine displacement in cc, engine horsepower in kW and squares of all these variables and a dummy for diesel. To keep the number of parameters down, only the total weight and its square was used in α_{1ij} — the remaining were close to insignificant and greatly increased estimation running time.

B Data

B.1 Sample Selection

Table B.1 shows how the sample size (new car purchases) gradually drops from the initial 311,057 cars to 128,910 as different sample selection criteria are imposed. The first criterion states that the household purchasing the car must own it for at least 90% of the 4-year driving period. This causes the most dramatic reduction in sample size because many households sell the car within this period. Figure B.1 shows a histogram of the fraction of the 4-year period that the purchasing household owns the car for the full sample of 311,057 purchases (disregarding the mass point at 100%). This shows that the share declines steadily down from 90% to 0%. The choice of 90% is to emphasize the need for accurate data on the driving to ensure that the selection on anticipated driving is pinned down by the data. Future work should look checking the sensitivity of the results to reducing the 90%.

The second criterion deselects 2-car households but allows a second car to be present for up to 50% of the period.

The third criterion deselects certain car types from the choice set by deleting purchases of cars that were purchased fewer than 30 times in the period 1997–2006. This has a very unfortunate implication in that diesel cars are heavily over represented in this group. Table B.2 shows the implications on the sample size (N), the number of cars ($|\mathcal{J}|$) and

Table B.1: Sample Selection

	(1) New cars	(2) Owns $\geq 90\%$	(3) Ncars ≤ 1.5	(4) #sold ≥ 30	(5) Final sample
1997	14,500	8,866	8,252	6,453	6,019
1998	45,075	27,986	24,895	22,248	21,374
1999	42,260	25,846	22,540	20,165	19,525
2000	30,070	17,699	15,350	12,764	12,461
2001	23,774	12,182	10,389	8,057	7,893
2002	28,648	16,305	14,035	11,611	11,016
2003	22,733	12,516	10,774	8,961	8,600
2004	29,535	16,552	14,095	11,901	11,548
2005	36,722	22,794	18,999	15,863	15,490
2006	37,740	24,670	19,793	15,458	14,984
<i>N</i>	311,057	185,416	159,122	133,481	128,910

(2): The family owns the car at least 90% of driving period,

(3): The family may own another car but no more than 50% of the driving period of this car,

(4): At least 30 of this car sold in full sample, (5): final sample.

Figure B.1: Fraction of the Driving Period Where the Original Owner Still Owns the Car




Table B.2: Deselecting Cars That are Rarely Sold and the Resulting Diesel Share

Threshold	Diesel % in 06	$ \mathcal{J} $	N
30	18.5%	1,177	128,007
20	19.6%	1,518	136,977
10	20.6%	2,105	144,820
5	21.0%	2,783	149,112
0	21.8%	7,572	154,089

Table B.3: Marginal External Costs per Km Travelled by Fuel Type^a

	Gas	Diesel	Unit
Noise	0.0478	0.0478	DKK/km
Accident	0.2095	0.2095	DKK/km
Congestion	0.3368	0.3368	DKK/km
Infrastructure	0.0097	0.0097	DKK/km
Air pollution	0.1352126	0.668475	DKK/liter
Climate	0.19764	0.21000	DKK/liter

^a Source: DTU Transport (2010). Note that only air pollution and climate depend on the fuel type.

the diesel market share in 2006 of setting this limit to 20, 10, 5 and 0 respectively. The true market share in 2006 was 21.8% but the restriction on the choice set results in a share of just 18.5%. However, bringing this up towards the truth increases the size of the choice set immensely, making estimation computationally very burdensome.

The final criterion makes routine checks such as dropping extreme observations (outside of the 0.1th or 99.9th percentiles) or rows with missing or senseless values.

B.2 Marginal External Costs of Driving

In this subsection, the marginal external cost estimates used for welfare calculations and for the construction of the diesel internalization counterfactual in section 6.4 are described. The cost estimates are taken from DTU Transport (2010) and they are provided by a major Danish research institution and used by Danish policy makers. The external costs of driving a km in a gasoline and diesel car respectively are reproduced in table B.3.

Two things are worth noting; Firstly, pollution and climate change costs are dwarfed by the congestion and accident externalities. While this particular externality is not well addressed with the model applied in this paper because it depends critically on when and where the driving takes place, it does mean that an increased traffic flow should be highly discouraged.

Secondly, the only place where diesel car externalities are different from those of gaso-

line cars is in terms of air pollution and climate change. The difference in climate change externalities stem from the fact that diesel cars typically drive farther per litre of fuel (a sales-weighted average of 18.1 versus 13.5 km/l for in 2006) while diesel only contains 10.4% more CO₂ per litre than gasoline does (2.640 kg/l 2.392 kg/l). The difference in air pollution comes primarily from particulate matter. For the Belgian context, Mayeres and Proost (2013) report that particulate matter makes up 85.0% of all emissions-related externalities per ton of diesel, far more than the externalities from SO₂ and NO_x. In fact, the marginal externality of diesel air pollution depends crucially on the population density. Since a dummy for living in one of the four largest Danish cities is already in the model, the expected diesel use and diesel market share has been calculated conditional on urban residence. It turned out that urban diesel use and purchases followed the overall numbers quite closely for the reforms considered here.

B.3 Descriptives


B.3.1 Work Distance

The work distance variable is the only one that is not taken directly from the register data. This one is calculated based on the travel tax deduction which comes from the personal tax registers. In Denmark, anyone living further than 12 km from their work place is eligible for a deduction depending on the distance times the number of days worked. The deduction is regardless of the number of hours worked and regardless of the type of transportation actually used by the worker. The deduction is a linear function of the km travelled above 24 (to and from work) but the rate drops to half after 100 km. In 2005, for example, it was DKK 1.68 for each km above 24 but below 100 and 0.84 for each km above 100. The rate was changed each year and twice in 2000. Moreover, as a part of a larger Danish reform in 1998 dubbed the Whitsun package, there was an adjustment to give a lift for the low-paid.

Note that in order to construct a work distance measure, one needs to know the number of days worked which is not observed. Therefore, it is assumed that everyone work 225 days a year.³¹ Note, however, that this only means that the work distance variable may be imprecise for the actual distance to the work but still precise about the variable of interest that is the annual km commuted to work. Figure B.2 shows the distribution of the constructed work distance measure for the *larger dataset* from which I take my estimation sample; in the left panel, the full distribution is showed. This clearly shows the censoring with a large mass point at zero. The right panel removes these zeros and shows the remaining distribution. There is a clear discontinuity at a work distance of 12 km, consistent with the fact that this is the threshold for eligibility. For all the observations below 12 km, we know that their actual work distance is larger than 12 km

³¹The official numbers for public sector employees in 2007–2010 were 224, 226, 225 and 228.

Figure B.2: Work Distance Distribution



but that they must have worked fewer than 225 days. For example, individuals with part time employment can be expected to fall there.

B.3.2 Fuel Prices

Figure B.3 shows the development in gasoline and diesel prices in Denmark in 2005 DKK. Prices have generally been increasing and moreover, it appears that diesel prices were converging on gasoline prices up towards 2008. Figure B.4 shows the price composition for both types of fuel; the fixed tax rate (dubbed the “Energy Tax”, which is split up into a CO2 tax in 2005) is fairly constant over the period with the exception of 1999 for gasoline and 2000 for diesel, where it is increased by 12% for both fuel types. In other words, most of the variation in fuel prices in Denmark comes from the product price. Figure B.5 shows the product price for Octane 95 and Diesel fuel together with the Western Texas Intermediate crude oil price. This figure shows that the prices have tracked the oil price very closely over the period, in particular for diesel fuel.


B.3.3 Car Characteristics

Figure B.7 shows the fraction of diesel cars in the register data (i.e. also data not included in my estimation sample). It shows the increase in the diesel market share that appears to really start increasing after 1997. The larger share of diesel cars with vintages in the 1980s can either be due to higher market share there or due to a different scrappage pattern for diesel cars then.

Figure B.8 shows the number of cars owned per household by year. The graph is based on a dataset containing all households and cars. The figure indicates that even though there has been an increasing share of households owning more than 1 car, the share is still extremely small compared to for example the US.

Figures B.9–B.11 show the development in median characteristics of sold cars. The most notable development is the increasing trend in weight for both types of fuel that has occurred all the way back to the 80’s. In this paper, weight proxies for the quality of the

Figure B.3: Real Price of Fuel, 1992–2011



Source: The Danish Oil Industry Association (eof.dk).

Figure B.4: The Composition of the Price of Gasoline and Diesel Fuel




Figure B.5: Gasoline and Diesel Product Price And Crude Oil Prices




Figure B.6: Relative Fuel Prices




Figure B.7: Diesel Share in Denmark by Vintage





Figure B.8: Number of Cars Owned per Household



Note: each bar shows the share of households in a given year that own the particular number of cars.

Figure B.9: Median Characteristics Over Time — Weight and Fuel Efficiency





Figure B.10: Median Characteristics Over Time— Engine Power and Displacement



car by measuring comfort and the carrying capacity of the car. Similarly, fuel efficiency has gone up dramatically but here we see that while it has been somewhat monotone for gasoline cars, almost all the growth for diesel cars occurred in 1997–99. Two things are worth noting there; Firstly, only 17 diesel cars are in the sample in 1997 so we are talking about very small numbers. Secondly, the advent of the Common Rail injection technology which quickly became standard in all diesel engines was the main reason for this. Apart from improving performance in terms of fuel efficiency, it also greatly improved the torque of the cars (which is not in my data) and changed the sound signature, making it more appealing to many consumers (according to an car salesman I have talked to).

The development in engine displacement, horse power and purchase price are much more erratic. This underlines the advantage of the chosen empirical model where all these characteristics are used in the household's comparison across cars, rather than focusing on each characteristic separately.

To better grasp the overall patterns in what car characteristics certain households end up with, table B.4 shows the estimates from regressing each car characteristic on household demographics. The results are much as one would expect with for example richer households purchasing more powerful and luxurious cars. It also shows some ambiguity

Figure B.11: Median Characteristics Over Time— Real Price (2005 DKK)





Figure B.12: Spatial Illustration: Municipality-averages of Work Distance and Diesel Frequency



in the effect of work distance — if males have a long work distance, they tend to prefer having a more comfortable ride whereas females tend to go for a more fuel efficient, smaller car.

The patterns shown in Table B.4 also show up clearly in the spatial patterns; Figure B.12 shows two maps of Denmark where the municipalities have been colored according to the average of the work distance (the maximum within the household) of Danish car-owning households in the left panel and the frequency of diesel cars in the right panel. These maps are drawn for a larger sample containing all car-owning households. The patterns show two interesting aspects. Firstly, there positive relationship between work distance and diesel ownership is in line with both urban areas and rural areas of Sealand (the big western Island); urban areas have low work distances and low diesel shares and vice versa for the rural areas of Sealand. However, in the Eastern part of the country, there appears to be low work distances and high diesel frequencies. These areas have very different employment patterns from the greater Copenhagen region, which is most likely a part of the explanation.

Table B.4: Car characteristics of new cars

	(1) Km/l	(2) Weight	(3) Diesel	(4) kW	(5) Displace	(6) Real price
$p^{\text{fuel}}_{\text{(O95)}}$	0.415*** (19.33)	-0.0193*** (-10.84)	0.0410*** (14.93)	-1.923*** (-12.53)	-4.645 (-1.92)	-8612.1*** (-14.37)
GDP (2005=1)	-11.39*** (-17.05)	-0.165** (-2.98)	-1.771*** (-20.71)	-3.053 (-0.64)	-978.8*** (-13.01)	-244123.3*** (-13.11)
Age (m)	-0.0136 (-1.48)	0.00373*** (4.88)	0.000654 (0.55)	0.473*** (7.18)	4.907*** (4.72)	1520.1*** (5.90)
Age squared (m)	0.0000800 (0.78)	-0.0000430*** (-5.03)	-0.0000207 (-1.57)	-0.00549*** (-7.45)	-0.0593*** (-5.10)	-17.83*** (-6.20)
Age (f)	-0.0400*** (-4.89)	0.00491*** (7.23)	-0.00118 (-1.13)	0.335*** (5.73)	4.149*** (4.50)	1516.1*** (6.64)
Age squared (f)	0.000306*** (3.30)	-0.0000445*** (-5.80)	-2.68e-08 (-0.00)	-0.00325*** (-4.92)	-0.0430*** (-4.13)	-14.48*** (-5.62)
Work dist. (m)	0.0150*** (44.34)	0.000136*** (4.86)	0.00262*** (60.59)	-0.00892*** (-3.70)	0.555*** (14.59)	84.66*** (9.00)
Work dist. (f)	0.0178*** (39.57)	-0.000444*** (-11.91)	0.00238*** (41.42)	-0.0512*** (-15.95)	-0.160** (-3.16)	-59.57*** (-4.76)
Income	-0.000000245*** (-16.83)	2.45e-08*** (20.33)	-6.92e-09*** (-3.71)	0.00000296*** (28.43)	0.0000434*** (26.49)	0.0166*** (40.84)
Male inc %	0.00797 (1.42)	-0.000305 (-0.66)	0.000721 (1.00)	0.00768 (0.19)	0.287 (0.45)	121.5 (0.78)
# kids	-0.303*** (-37.23)	0.0417*** (61.88)	0.0000115 (0.01)	1.244*** (21.43)	24.10*** (26.33)	7462.6*** (32.94)
Urban dummy	0.0198 (1.22)	-0.00707*** (-5.26)	-0.00564** (-2.72)	-0.754*** (-6.51)	-8.729*** (-4.78)	-2115.9*** (-4.68)
Unemployed (m)	0.260*** (11.16)	-0.0407*** (-21.06)	-0.00916** (-3.07)	-3.412*** (-20.47)	-53.62*** (-20.41)	-15345.8*** (-23.59)
Unemployed (f)	0.170*** (9.50)	-0.0146*** (-9.86)	0.00574* (2.51)	-1.426*** (-11.14)	-21.18*** (-10.49)	-5694.4*** (-11.40)
Linear time trend	0.434*** (43.37)	0.0181*** (21.76)	0.0449*** (35.06)	1.082*** (15.11)	14.93*** (13.23)	7397.0*** (26.49)
Constant	21.45*** (42.70)	1.653*** (39.69)	1.231*** (19.16)	65.59*** (18.28)	2230.9*** (39.43)	404939.1*** (28.91)
N	128910	128910	128910	128910	128910	128910

For variable labels, m denotes male and f denotes female.

Same sample as the one used for the two-period model.

(1) Fuel efficiency in km/l, (2) weight in tons, (3) LPM for diesel,

(4) engine power in kW and (5) displacement in cc.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure B.13: Driving by Fuel Type





Figure B.14: Median VKT vs Fuel Price Over Time for Gas and Diesel



B.3.4 Descriptive Evidence on Driving


Figure B.13 shows the driving distribution for diesel car drivers and gasoline car drivers. The distribution for diesel car drivers is shifted strongly towards higher driving.

Figure B.14 shows median vehicle kilometers travelled (VKT) against median fuel price over time for gasoline cars (left panel) and diesel cars (right panel). Both figures show that the typical car purchased in later years ends up driving less than in earlier years and that fuel prices have been increasing. This is consistent with a negative fuel price elasticity.

The corresponding figures where the price per kilometer (PPK, p_{jt}^{fuel}/e_j) is used are shown in figure B.15 and here the picture is much less clear picture because fuel efficiency also increases over time. This is precisely the selection effect at play where consumers are moving to more fuel efficient cars to counteract the increasing fuel prices.

Table B.5 shows the results from regressing VKT on PPK, car characteristics and household demographics. The most central result is that the mean estimated PPK-elasticity depends very strongly on whether a different mean driving is allowed for diesel

Figure B.15: Median VKT vs Price Per Kilometer (PPK) Over Time for Gas and Diesel



car households (which decreases the mean elasticity from $-.74$ to $-.30$). This is further emphasized by the fact that the elasticity is -0.16 when estimated on the gasoline sample only and $-.39$ on the diesel subsample. Gillingham and Munk-Nielsen (2015) explore the heterogeneity in the fuel price elasticity on household demographics and the interested reader is referred to that paper.

C Joint Estimation of the λ -parameter

In this section, I discuss the issue with the estimation of the logit error term scaling parameter, λ , and present an idea for estimating a more flexible extension of the model that might facilitate joint estimation. I first discuss the problem, providing intuition about the λ parameter and why the maximum likelihood estimate is so high. I then argue that car fixed effects can be the cause of the problem and that controlling for these may solve the issue of the high λ . In light of this, I conclude with an outline a strategy for incorporating car type fixed effects into the model in the spirit of Berry, Levinsohn, and Pakes (1995) for future research.

As mentioned in Section (4), I have chosen to normalize $\alpha_2 = -1$ and estimate λ . When I estimate the model, I first estimate the reduced-form driving parameters subject to the normalization α_2 . I then use these parameters as the starting values for the full, joint optimization. However, the likelihood function is increasing in λ up to the point where λ is so high that the model just predicts uniform choice probabilities for all choices. Recall that in logit models where the choice-specific utilities are non-linear (for example the present model or dynamic discrete choice models), the λ is sometimes identified and then it acts as a *smoothing* parameter. In some sense, it is analogous to the bandwidth in a Nadaraya-Watson kernel density estimator; in one extreme, when $\lambda \rightarrow 0$, the choice probabilities converge to an indicator function for the highest utility choice, $\Pr(j) = \mathbf{1}\{j = \arg \max_{j'} u_{j'}\}$. In the other extreme, when $\lambda \rightarrow \infty$, we choice probabilities become uniform, $\Pr(j) = 1/|\mathcal{J}| \forall j$, where $|\mathcal{J}|$ is the number of choices available. In intuitive

Table B.5: VKT Regressions — Price per Kilometer (PPK) Elasticity

	(1) Simple	(2) Diesel dummy	(3) Year FE	(4) Only gas	(5) Only diesel
Price per km	-50.50*** (-21.52)	-20.32*** (-8.29)	-16.82*** (-6.17)	-10.01** (-2.83)	-65.73 (-1.94)
GDP	-41.98*** (-29.57)	-35.10*** (-24.69)	-57.10*** (-12.38)	-50.23*** (-10.68)	-82.38*** (-3.88)
Age (m)	0.518*** (10.62)	0.538*** (11.12)	0.536*** (11.09)	0.415*** (8.51)	1.256*** (6.15)
Age squared (m)	-0.00777*** (-13.57)	-0.00792*** (-13.91)	-0.00789*** (-13.87)	-0.00677*** (-11.85)	-0.0138*** (-5.59)
Work dist. (m)	0.353*** (119.98)	0.348*** (118.89)	0.348*** (118.99)	0.333*** (104.54)	0.381*** (47.49)
Work dist. (f)	0.340*** (87.05)	0.334*** (85.83)	0.334*** (85.92)	0.356*** (84.04)	0.260*** (24.27)
Income	-0.00000250*** (-19.67)	-0.00000239*** (-18.89)	-0.00000234*** (-18.54)	-0.00000231*** (-18.68)	-0.00000367*** (-4.55)
# kids	0.236*** (3.61)	0.223*** (3.43)	0.219*** (3.37)	0.128 (1.93)	0.515* (2.05)
Urban dummy	-1.131*** (-8.84)	-1.106*** (-8.70)	-1.092*** (-8.60)	-1.262*** (-9.88)	0.973 (1.80)
Unemployed (m)	0.492** (2.64)	0.438* (2.36)	0.459* (2.48)	0.565** (3.03)	-0.499 (-0.64)
Unemmployed (f)	-0.0700 (-0.49)	-0.0930 (-0.65)	-0.0784 (-0.55)	-0.0474 (-0.33)	-0.410 (-0.74)
Km/l	0.610*** (6.07)	-0.257* (-2.51)	-0.118 (-1.02)	0.117 (0.77)	-1.475 (-1.90)
Weight	0.0329*** (67.23)	0.0209*** (36.60)	0.0210*** (36.45)	0.0229*** (38.60)	0.0116*** (5.36)
Engine power	-0.0368*** (-5.52)	0.0426*** (6.18)	0.0428*** (6.19)	0.0483*** (6.67)	0.101*** (3.82)
Engine size	0.0140*** (33.71)	0.00367*** (7.55)	0.00352*** (7.23)	0.00180*** (3.36)	0.00178 (1.14)
Diesel dummy		17.99*** (40.22)	18.09*** (40.37)		
Constant	28.02*** (10.49)	43.76*** (16.30)	66.88*** (11.82)	54.31*** (8.78)	157.5*** (5.85)
Year FE	No	No	Yes	Yes	Yes
N	128007	128007	128007	114623	13384
R^2	0.348	0.356	0.357	0.235	0.216
Avg. elasticity	-0.744	-0.300	-0.248	-0.158	-0.392


In variable names, m denotes male and f denotes female.

Column 4 contains only gasoline cars and 5 only diesels.

Year FE: for each year, a dummy for whether the driving period covers the year.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure C.1: Market Shares in the Characteristics Space



terms, λ indicates how *precise* the model is, since it does not alter the ordering of the conditional utilities of the alternatives.

With this intuition at hand, it is easier to understand why the likelihood function is maximized for such a high value of λ . For given values of the remaining structural parameters, the model will tend to assign similar choice probabilities to cars that are in the same region of the choicest. However, Figure C.1 illustrates that this is not the case in the data. The figure shows a scatter plot of the cars that are available in the 2006 choicest. The x-axis denotes fuel efficiency in km/l and the y-axis denotes weight in metric tons while the coloring of the dots indicates the market share of each car in 2006. The figure shows that there are cars that are very close in characteristics with very dissimilar market shares. This will all else equal point towards characteristics not being important for determining the market shares of cars. However, the cross-equational restrictions implied by the model structure are such that to reduce the importance of the characteristics, the driving predictions will be altered. Therefore, I conjecture that the high value of λ is a way for the optimizer to reduce the importance of characteristics in predicting market shares without resulting in a bad fit of the driving equation.

If the explanation I have outlined above is correct, then controlling for car fixed effects should solve the problem. However, the big question for future work on this is whether they will suck up all the variation and result in a model where car choice becomes equally unresponsive to changes in policy; something which a priori must be wrong in light of the stark changes in the fuel efficiency and the diesel share following the 2007 reform.

For future research, I will now outline a potential strategy for estimating an extension of the model presented in this paper that allows for fully flexible car type fixed effects, $u^{\text{own}}(j) = \xi_j$. This is in line with the agenda of the Berry, Levinsohn, and Pakes (1995) literature, emphasizing the importance of unobserved car characteristics correlated with

price (and possibly other characteristics).

This model with product-level fixed effects may be estimated in two ways; A direct approach would be to simply estimate all the $J - 1 = 1,176$ dummies with maximum likelihood. Estimating such a large number of parameters would not be feasible using numerical derivatives, but with analytic derivatives and the BHHH approximation of the Hessian, complexity only increases linearly in the number of parameters.

An alternative approach is to apply a fixed point like that proposed by Berry (1994). Let $\Gamma : \mathbb{R}^{J-1} \rightarrow \mathbb{R}^{J-1}$ be the operator defined by $\Gamma(\xi^{[i]}) = (\Gamma_1(\xi^{[i]}), \dots, \Gamma_{J-1}(\xi^{[i]}))$, where

$$\Gamma_j(\xi^{[i]}) = \xi_j^{[i]} + \sum_{t \in \mathcal{T}_j} \varpi_{jt} \left[\log s_{jt}^{\text{data}} - \log s_{jt}^{\text{pred}}(\xi^{[i-1]}) \right],$$

where s_{jt} is the market share for car j in year t , $\mathcal{T}_j$ is the set of years where car j was available and

$$\varpi_{jt} = \frac{N_t}{\sum_{t \in \mathcal{T}_j} N_t},$$

where N_t is the number of households going on the market in year t . Letting $\tilde{u}_{ij} \equiv u_{ij} - u^{\text{own}}(j)$, the predicted market share is given by

$$s_{jt}^{\text{pred}}(\xi^{[i-1]}) = \frac{1}{N_t} \sum_{i=1}^{N_t} \frac{\exp \left[(\tilde{u}_{ij} + \xi_j^{[i-1]}) / \lambda \right]}{\sum_{k \in \mathcal{J}_t} \exp \left[(\tilde{u}_{ik} + \xi_k^{[i-1]}) / \lambda \right]}.$$

This gives rise to the following algorithm;

Algorithm: A Berry (1994) fixed point.

Initialization: Set $\xi_j^{[0]} := \sum_{t \in \mathcal{T}_j} \varpi_{jt} \log s_{jt}^{\text{data}}$ and pick a reference car, j_0 , for which $\xi_{j_0} := 0$.

Iteration: Given $\xi^{[i-1]}$, let $\xi^{[i]} = \Gamma(\xi^{[i-1]})$. Continue until $\|\xi^{[i]} - \xi^{[i-1]}\| < \epsilon$.

Recently, there has been some debate about numerical concerns with the implementations of algorithms using nested fixed points such as Berry (1994); Berry, Levinsohn, and Pakes (1995); Rust (1987). Dubé, Fox, and Su (2012) have emphasized the importance of using a tight inner-loop tolerance (ϵ) to avoid numerical noise spilling out into the outer loop. They suggest using the MPEC approach (Judd and Su, 2012). Instead, I follow the approach by Iskhakov et al. (2015) and use analytic derivatives for the inner loop, replacing the fixed-point iteration shown above with a root-finding solver for the quadratic system of non-linear equations,

$$\xi - \Gamma(\xi) = 0.$$

By using the analytic Jacobian of the operator Γ , which has a computationally simple form, I find that the solver converges in 13 iterations to machine precision.

D Additional Results

Table D.1 shows the structural elasticities from the preferred specification. The results are estimated based on a model with perfect foresight that allows random effects ($c_i \neq 0$). For the presented set of estimates, α_2 was fixed to -1 , but very recently, I have successfully estimated that coefficient as well without it significantly changing the results.

Table D.2 shows the results from the baseline on the 2006 data as well as the 2007 counterfactual implemented in 2006 (same as column (3) of table 6.1) and an additional simulation of the 2007 reform where the pivot point of diesel cars is moved from $18^{\text{km/l}}$ to $20^{\text{km/l}}$. The motivation is that the pivot point for gasoline cars is $16^{\text{km/l}}$ but a typical diesel car drives about 4 km further per liter of fuel than a gasoline car. In that sense, the pivot of $20^{\text{km/l}}$ should provide a better balance in the incentives.

In figure D.1 is shown the observed diesel share, the simulated diesel share from the model and a counterfactual simulation where both fuel price time series are kept at the 1997 level. The figure shows that the diesel share would have been higher in the later years if fuel prices had not changed. Two important points should be noted; Firstly, since the model conditions on entry into the new car market, raising or lowering fuel prices, for all cars will not change results as drastically as if more households were allowed to switch into car ownership. Nonetheless, raising fuel prices will lower expected driving and utility so given the convex utility in driving, some consumers will move towards more fuel efficient vehicles and therefore also diesel cars. This is also why, in the structural elasticities in table 5.3 we saw that when all fuel prices go up by 1%, the diesel share grows by 0.15%.

Secondly, the more important implication of holding fuel prices at the 1997 level is that the *relative* price of gasoline to diesel is kept constant. Figure D.2 plots two time series. On the left axis is the expected price of gasoline divided by the expected price of diesel (under perfect foresight — i.e. the fuel prices that are driving expectations) for a household going on the market in the given year and on the right axis is the predicted diesel market share for the year divided by the observed share. The figure shows that the tendency of the model to over or under-predict the diesel share is systematically related to the relative fuel prices. For example, the predicted share has two particularly striking periods; In 99–00, the prediction moves from over to under the observed share, coinciding with a sharp jump down in the relative price (diesel caught up with gasoline). In 05, the model has a kink down, under-predicting the diesel share. This coincides with a sharp jump down in the relative price from 117.9% to 110.9%, making diesels less favorable. Note that the predicted to observed share is not shown for 1997 because it is 432%. This extreme number is due to the observed share being quite close to zero in that year.

Table D.1: Estimated parameters

Fixed Parameters				
Parameter	Value			
β	0.95			
ψ	1			
λ	10000			
Model: Perfect foresight, quasi-linear, random effects.				
General Parameters				
Parameter	Estimate		t	
σ_x	16.093		(69.12)	
σ_α	21.951		(31.77)	
Demographics				
	γ_z		α_{1z}	
Parameter	Estimate	t	Estimate	t
Constant	47.596	(35.22)	—	(—)
age	-8.447	(-18.97)	8.901	(8.71)
agesq	7.363	(15.88)	-15.168	(-14.39)
WDm	8.170	(18.95)	17.889	(69.45)
WDf	1.079	(19.46)	9.684	(108.20)
inc	-9.457	(-31.44)	-8.768	(-39.94)
nkids	1.453	(11.65)	-0.458	(-2.93)
city	-0.210	(-1.48)	-1.412	(-10.09)
Car Parameters				
Parameter	Estimate		t	
α_{10}	74.927		(14.88)	
α_{20}	-1.000		†	
$\alpha_{0,\text{weight}}$	124074.734		(41.91)	
$\alpha_{0,\text{weight}^2}$	-5009.689		(-5.67)	
$\alpha_{0,\text{kw}}$	-413.653		(-25.53)	
α_{0,kw^2}	5.114		(46.83)	
$\alpha_{0,\text{displace}}$	-194.172		(-0.15)	
$\alpha_{0,\text{displace}^2}$	4976.559		(13.12)	
$\alpha_{0,\text{diesel}}$	-4235.595		(-24.99)	
$\alpha_{1,\text{weight}}$	18.876		(3.12)	
$\alpha_{1,\text{weight}^2}$	10.189		(5.64)	

†: Fixed parameter, see section ??.

Table D.2: Simulation of the 2007 Feebate Reform — The Role of the Diesel Pivot

	(1)	(2)	(3)
	Baseline	2007	2007 alt.
Consumer welfare			
E(CS)	114970.09	115989.89	115363.51
Total taxes			
E(total taxes)	146623.83	134398.55	134482.53
Ownership tax			
E(Regtax revenue)	106556.44	98175.80	97702.93
E(Owntax revenue)	11093.62	9813.59	9779.74
Fuel tax			
E(O95 revenue)	25115.85	21695.96	23122.16
E(Diesel revenue)	3857.92	4713.20	3877.70
Driving/fuel use			
E(VKT)	79663.89	78740.62	78323.44
E(litre O95)	4340.92	3749.84	3996.34
E(litre D)	891.32	1088.92	895.89
E(litre D—urban)	188.02	230.59	189.60
E(kg CO2)	12736.56	11844.36	11924.39
Characteristics			
E(fe)	15.92	17.04	16.75
E(we)	1.70	1.64	1.64
E(kw)	77.08	70.07	70.72
E(displace)	1.65	1.54	1.54
E(% diesel)	18.49	24.48	19.63
E(% diesel—urban)	3.89	5.18	4.15

2007: The feebate reform of 2007 is implemented in 2006.

2007 alt.: As 2007 but the diesel pivot is 20 km/l instead of 18 km/l.

Figure D.1: Counterfactual Simulation: The Diesel Share Under Constant Fuel Prices





Figure D.2: Relative Fuel Prices and Relative Market Share Error



D.1 Static Expectations

Table D.3 shows the structural elasticities from the model estimated imposing the assumption of static expectations. The elasticity of driving with respect to PPK is now -39% as opposed to -30% for the perfect foresight estimates, indicating that to fit the data, the estimates must emphasize monetary costs more in this version of the model. Similarly, when the fuel efficiency of all cars in the choice set go up by 1%, the expected fuel efficiency goes up by 0.93% as opposed to 0.90% with perfect foresight. In other words, consumers are still substituting away some technological gains in fuel efficiency for other engine characteristics but not as much as earlier. And in particular, as PPK rises, the expected diesel share now falls. Finally, as the weight of all cars goes up by 1%, the expected weight now goes up by 1.58%, as opposed to just 1.15% earlier and the expected driving response (allowing for changes on the extensive margin) goes up by 1.71% as compared to 1.01% under static expectations.

In short, the estimates from the model imposing static expectations imply that money matters more to consumers and that the weight of the car also matters more for how much it is driven.

Figure D.3 compares the diesel share predictions from the models that impose perfect foresight and static expectations respectively with the observed diesel share. The movements in the two are highly similar but there is a slight tendency in the later years for the static expectations prediction to be slightly below the other.

Figure D.4 shows the 1997 counterfactual simulation using the estimates imposing static expectations. It shows that the conclusion from the perfect foresight model still holds; The counterfactual simulation where the 1997 reform was never imposed show a dramatically smaller diesel share in all years (but still an increase over time).

Table D.3: Structural Elasticities — Static Expectations

	(1) Baseline	(2) Fuel efficiency	(3) Weight	(4) Fuel prices	(5) O95 prices
Consumer welfare					
CS	64412.21	0.43	2.32	-0.43	-0.33
Total taxes					
E(total taxes)	139431.76	0.05	1.24	-0.05	0.05
Ownership and registration tax					
E(Regtax revenue)	101066.75	0.19	1.11	-0.19	-0.02
E(Owntax revenue)	9999.10	0.23	1.37	-0.23	0.03
Fuel/RUC tax					
E(O95 revenue)	23801.00	-0.55	0.85	0.55	-0.04
E(Diesel revenue)	4564.90	-0.43	6.32	0.41	2.08
Driving/fuel use					
E(VKT)	81183.20	0.39	1.74	-0.39	-0.21
E(litre O95)	4113.67	-0.55	0.85	-0.44	-1.03
E(litre D)	1054.66	-0.43	6.32	-0.58	2.08
E(litre D—urban)	225.43	-0.40	6.40	-0.61	1.99
E(kg CO ₂)	12624.18	-0.52	2.04	-0.47	-0.34
Characteristics					
E(fe)	16.18	0.93	-0.18	0.06	0.15
E(we)	1.70	0.10	1.58	-0.10	-0.01
E(kw)	72.07	0.14	0.71	-0.14	-0.06
E(displace)	1.53	0.10	0.58	-0.10	-0.00
E(% diesel)	19.77	0.25	4.36	-0.26	2.08
E(% diesel—urban)	4.20	0.27	4.39	-0.29	1.98

Elasticities based on estimates imposing static expectations

(2): Relative changes when e_j increases by 1% for all j .

(3): Relative changes when weight_j increases by 1% for all j .

(4): Relative changes when fuel prices increase by 1%.

(4): Relative changes when gasoline prices increase by 1%.

All numbers are averages weighted with CCPs.

Figure D.3: Diesel Share Predictions — Comparing the Perfect Foresight and Static Expectations Predictions





Figure D.4: 1997 Counterfactual — Static Expectations



References

- Adamou, Adamos, Sofronis Clerides, and Theodoros Zachariadis. 2013. “Welfare Implications of Car Feebates: A Simulation Analysis.” *The Economic Journal* URL <http://dx.doi.org/10.1111/econj.12094>.
- Adda, Jérôme and Russell Cooper. 2000. “Balladurette and Juppette: A Discrete Analysis of Scrapping Subsidies.” *Journal of Political Economy* 108 (4):778–806. URL <http://www.jstor.org/stable/10.1086/316096>.
- Allcott, Hunt and Nathan Wozny. 2012. “Gasoline prices, fuel economy, and the energy paradox.” *NBER Working Paper* .
- Anderson, Soren T., Ryan Kellogg, Ashley Langer, and James M. Sallee. 2013. “The Inter-generational Transmission of Automobile Brand Preferences: Empirical Evidence and Implications for Firm Strategy.” Working Paper 19535, National Bureau of Economic Research. URL <http://www.nber.org/papers/w19535>.
- Bento, Antonio M, Lawrence H Goulder, Mark R Jacobsen, and Roger H Von Haefen. 2009. “Distributional and efficiency impacts of increased US gasoline taxes.” *The American Economic Review* :667–699.
- Beresteanu, Arie and Shanjun Li. 2011. “Gasoline Prices, Government Support, And The Demand For Hybrid Vehicles In The United States.” *International Economic Review* 52 (1):161–182. URL <http://dx.doi.org/10.1111/j.1468-2354.2010.00623.x>.
- Berndt, Ernst R, Bronwyn H Hall, Robert E Hall, and Jerry A Hausman. 1974. “Estimation and inference in nonlinear structural models.” In *Annals of Economic and Social Measurement, Volume 3, number 4*. NBER, 653–665.
- Berry, S., J. Levinsohn, and A. Pakes. 1995. “Automobile prices in market equilibrium.” *Econometrica: Journal of the Econometric Society* 63 (4):841–890.
- Berry, Steven T. 1994. “Estimating discrete-choice models of product differentiation.” *The RAND Journal of Economics* :242–262.
- Borger, Bruno De, Ismir Mulalic, and Jan Rouwendal. 2013. “Substitution between Cars within the Household.” *Tinbergen Institute Discussion Paper* .
- Busse, Meghan R, Christopher R Knittel, and Florian Zettelmeyer. 2013. “Are Consumers Myopic? Evidence from New and Used Car Purchases.” *The American Economic Review* 103 (1):220–256.

- Chandra, Ambarish, Sumeet Gulati, and Milind Kandlikar. 2010. “Green drivers or free riders? An analysis of tax rebates for hybrid vehicles.” *Journal of Environmental Economics and Management* 60 (2):78 – 93. URL <http://www.sciencedirect.com/science/article/pii/S0095069610000598>.
- Chen, Jiawei, Susanna Esteban, and Matthew Shum. 2010. “How much competition is a secondary market?” Working Papers 2010-06. URL <http://ideas.repec.org/p/imd/wpaper/wp2010-06.html>.
- Clerides, Sofronis and Theodoros Zachariadis. 2008. “The effect of standards and fuel prices on automobile fuel economy: an international analysis.” *Energy Economics* 30 (5):2657–2672.
- De Borger, Bruno and Inge Mayeres. 2007. “Optimal taxation of car ownership, car use and public transport: Insights derived from a discrete choice numerical optimization model.” *European Economic Review* 51 (5):1177–1204.
- D’Haultf  uille, Xavier, Pauline Givord, and Xavier Boutin. 2014. “The Environmental Effect of Green Taxation: The Case of the French Bonus/Malus.” *The Economic Journal* 124 (578):F444–F480. URL <http://dx.doi.org/10.1111/eoj.12089>.
- DTU Transport. 2010. Tech. rep., DTU. URL <http://www.modelcenter.transport.dtu.dk/Publikationer/Transportoekonomiske-Enhedspriser>. Version 1.1.
- Dub  , Jean-Pierre, Jeremy T Fox, and Che-Lin Su. 2012. “Improving the numerical performance of static and dynamic aggregate discrete choice random coefficients demand estimation.” *Econometrica* 80 (5):2231–2267.
- Dubin, J.A. and D.L. McFadden. 1984. “An econometric analysis of residential electric appliance holdings and consumption.” *Econometrica: Journal of the Econometric Society* :345–362.
- Feng, Ye, Don Fullerton, and Li Gan. 2013. “Vehicle choices, miles driven, and pollution policies.” *Journal of Regulatory Economics* 44 (1):4–29. URL <http://dx.doi.org/10.1007/s11149-013-9221-z>.
- Frondel, Manuel, Jorg Peters, and Colin Vance. 2008. “Identifying the Rebound: Evidence from a German Household Panel.” *Energy Journal* 29 (4):154–163.
- Frondel, Manuel, Nolan Ritter, and Colin Vance. 2012. “Heterogeneity in the rebound effect: Further evidence for Germany.” *Energy Economics* 34 (2):461–467.
- Gallagher, Kelly Sims and Erich Muehlegger. 2011. “Giving green to get green? Incentives and consumer adoption of hybrid vehicle technology.” *Journal of Environmental*

- Economics and Management* 61 (1):1 – 15. URL <http://www.sciencedirect.com/science/article/pii/S0095069610000768>.
- Gavazza, Alessandro, Alessandro Lizzeri, and Nikita Rokestkiy. 2014. “A quantitative analysis of the used-car market.” *American Economic Review* URL <http://mpa.ub.uni-muenchen.de/id/eprint/38414>.
- Gillingham, K. 2012. “Selection on Anticipated Driving and the Consumer Response to Changing Gasoline Prices.” *Working paper* .
- Gillingham, Kenneth, Fedor Iskhakov, Anders Munk-Nielsen, John Rust, and Bertel Schjerning. 2013. “A Dynamic Model of Vehicle Ownership, Type Choice, and Usage.” *Working Paper* .
- Gillingham, Kenneth and Anders Munk-Nielsen. 2015. “The Tail Wagging the Dog: Commuting and the Fuel Price Response in Driving.” *Working Paper* .
- Goldberg, P.K. 1998. “The effects of the corporate average fuel efficiency standards in the US.” *The Journal of Industrial Economics* 46 (1):1–33.
- Greene, David L. 2010. “How Consumers Value Fuel Economy: A Literature Review.” Tech. Rep. EPA-420-R-10-008, U.S. Environmental Protection Agency.
- Grigolon, Laura, Mathias Reynaert, and Frank Verboven. 2015. “Consumer Valuation of Fuel Costs and the Effectiveness of Tax Policy: Evidence from the European Car Market.” *Working Paper* .
- Huse, Cristian and Claudio Lucinda. 2013. “The Market Impact and the Cost of Environmental Policy: Evidence from the Swedish Green Car Rebate.” *The Economic Journal* URL <http://dx.doi.org/10.1111/eoj.12060>.
- Hymel, Kent M. and Kenneth A. Small. 2015. “The Rebound Effect for Automobile Travel: Asymmetric Response to Price Changes and Novel Features of the 2000s.” *Energy Economics* forthcoming.
- Iskhakov, Fedor, John Rust, Bertel Schjerning, and Jinhyuk Lee. 2015. “Constrained Optimization Approaches to Estimation of Structural Models: Comment.” *Working Paper* URL <http://ssrn.com/abstract=2583655>.
- Jacobsen, Mark. 2013. “Evaluating U.S. Fuel Economy Standards in a Model with Producer and Household Heterogeneity.” *American Economic Journal: Economic Policy* 5(2):148–187.
- Judd, Kenneth and Che-Lin Su. 2012. “Constrained Optimization Approaches to Estimation of Structural Models.” *Econometrica* 80 (5):2213–2230.

- Judd, Kenneth L and Benjamin S Skrainka. 2011. "High performance quadrature rules: How numerical integration affects a popular model of product differentiation." *Working Paper* URL <http://ssrn.com/abstract=1870703>.
- Li, Shanjun, Joshua Linn, and Erich Muehlegger. 2014. "Gasoline Taxes and Consumer Behavior." *American Economic Journal: Economic Policy* 6 (4):302–42. URL <http://www.aeaweb.org/articles.php?doi=10.1257/pol.6.4.302>.
- Mabit, Stefan L. 2014. "Vehicle type choice under the influence of a tax reform and rising fuel prices." *Transportation Research Part A: Policy and Practice* 64:32–42.
- Marion, Justin and Erich Muehlegger. 2011. "Fuel tax incidence and supply conditions." *Journal of Public Economics* 95 (9–10):1202–1212. URL <http://www.sciencedirect.com/science/article/pii/S0047272711000545>.
- Mayeres, Inge and Stef Proost. 2013. "The taxation of diesel cars in Belgium - revisited." *Energy Policy* 54 (0):33–41. URL <http://www.sciencedirect.com/science/article/pii/S0301421511009670>.
- Miravete, Eugenio J., Maria J. Moral, and Jeff Thurk. 2014. "Protecting the European Automobile Industry through Environmental Regulation: Adoption of Diesel Engines."
- Reynaert, Mathias. 2014. "Abatement strategies and the cost of environmental regulation: Emission standards on the European car market." *KU Leuven Center for Economic Studies Discussion Paper Series DPS14* 31.
- Rust, John. 1987. "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher." *Econometrica* :999–1033.
- Sallee, James M, Sarah E West, and Wei Fan. 2010. "The effect of gasoline prices on the demand for fuel economy in used vehicles: Empirical evidence and policy implications." *Working Paper* URL <http://www.ntanet.org/images/stories/pdf/proceedings/09/032.pdf>.
- Schiraldi, P. 2011. "Automobile replacement: a dynamic structural approach." *The RAND journal of economics* 42 (2):266–291.
- Small, K.A. and H.S. Rosen. 1981. "Applied Welfare Economics with Discrete Choice Models." *Econometrica* 49 (1):105–130.
- Small, K.A. and K. Van Dender. 2007. "Fuel Efficiency And Motor Vehicle Travel: The Declining Rebound Effect." *Energy Journal* 28 (1):25.

- Spiller, Elisheba. 2012. “Household Vehicle Bundle Choice and Gasoline Demand: A Discrete-Continuous Approach.” *Working Paper* .
- Stolyarov, Dmitriy. 2002. “Turnover of used durables in a stationary equilibrium: Are older goods traded more?” *Journal of Political Economy* 110 (6):1390–1413.
- Train, K.E. 2009. *Discrete choice methods with simulation*. Cambridge University Press, second ed.
- Verboven, Frank. 2002. “Quality-based price discrimination and tax incidence: evidence from gasoline and diesel cars.” *Rand Journal of Economics* :275–297.
- Wakamori, Naoki. 2011. “Portfolio considerations in differentiated product purchases: An application to the Japanese automobile market.” *Bank of Canada Working Paper* (27).
- West, S.E. 2004. “Distributional effects of alternative vehicle pollution control policies.” *Journal of Public Economics* 88 (3):735–757.

Chapter 3

A Dynamic Model of Vehicle Ownership, Type
Choice and Usage

A Dynamic Model of Vehicle Ownership, Type Choice, and Usage*

Kenneth Gillingham, *Yale University*[†]

Fedor Iskhakov, *University of New South Wales*[‡]

Anders Munk-Nielsen, *University of Copenhagen*[§]

John Rust, *Georgetown University*[¶]

Bertel Schjerning, *University of Copenhagen*^{||}

Abstract

This paper develops an estimable structural microeconomic model of car choice and usage that features endogenous equilibrium prices on the used-car market. Households buy and sell cars in the market and car owners choose how much to drive their car in a finite-horizon model. Moreover, we explicitly model the choice between scrapping the car or selling it on the used-car market. We estimate the model using full-population Danish register data on car ownership, driving and demographics for the period 1996–2009, covering all Danish households and cars. Simulations show that the equilibrium prices are essential for producing realistic simulations of the car age distribution and scrappage patterns over the macro cycle. We illustrate the usefulness of the model for policy analysis with a counterfactual simulation that reduces new car prices but raises fuel taxes. The simulations show how equilibrium prices imply that the boom in new car sales come at the cost of accelerated scrappage of older cars. Furthermore, the model gives predictions on tax revenue, fuel use, emissions, the lifetime of vehicles as well as the composition of types and ages of cars in the future.

Keywords: Automobiles, emissions, carbon tax, dynamic programming, secondary market

JEL classification: D92, L11, L13, Q38

*The authors would like to thank the participants at the 2014 Dynamic Interactions and Behavior SITE meetings at Stanford for useful comments. We would also like to thank Maria Juul Hansen for excellent research assistance. Finally, we would like to acknowledge funding from the IRUC research project, financed by the Danish Council for Independent Research.

[†]Yale University, Forestry & Environmental Studies, Department of Economics, School of Management, 195 Prospect Street, New Haven, CT 06511, kenneth.gillingham@yale.edu

[‡]University of New South Wales, Centre of Excellence in Population Ageing Research, Kensington NSW 2052, Australia, f.iskhakov@unsw.edu.au

[§]University of Copenhagen, Department of Economics and Centre for Applied Microeconomics (CAM), ster Farimagsgade 5, building 26, DK-1353, Copenhagen K, Denmark, anders.munk-nielsen@econ.ku.dk

[¶]**Corresponding Author:** Georgetown University, Department of Economics, Georgetown University, Washington, DC, jr1393@georgetown.edu

^{||}University of Copenhagen, Department of Economics and Centre for Applied Microeconomics (CAM), ster Farimagsgade 5, building 26, DK-1353 Copenhagen K, Denmark, bertel.schjerning@econ.ku.dk

Contents

1 Introduction

2 Background and Data

- 2.1 Institutional Setting
- 2.2 Data
- 2.3 Descriptives
- 2.4 Previous literature
 - 2.4.1 Discrete-continuous Models of Durables
 - 2.4.2 Models of Equilibrium in Automobile Markets
 - 2.4.3 Estimation of Dynamic Discrete Choice Models

3 The Model

- 3.1 Household Dynamic Vehicle Choice Problem
- 3.2 Utility Specification
- 3.3 Specification of the Transition Densities

4 Estimation of the Model

5 Solving for Equilibrium Prices

- 5.1 Solving for Equilibrium Prices
- 5.2 Simulating Forward in Time
- 5.3 Non-Stationary Expectations

6 Results

- 6.1 Implementation
- 6.2 First-Stage Results
- 6.3 Structural Estimates
- 6.4 Equilibrium Prices: In-Sample
- 6.5 Counterfactual Simulations

7 Conclusion

A Appendix: Income Transitions

B Appendix: Background and Data

- B.1 Institutional Background
- B.2 Additional Descriptives
- B.3 Scrappage
- B.4 Driving

C Appendix: Flexible Price Function Specification

- C.1 Derivatives of the price function with respect to (θ, α)
- C.2 Non-monotonic specification

D Appendix: Test equilibria

E Appendix 5: Additional Results

E.1	Estimates with Fixed Transaction Costs	
E.2	Equilibrium Prices	
E.3	Counterfactual Simulations	

F Notation

1 Introduction

Government policies that affect durable goods inherently influence equilibria in both the new and used markets. The presence of a secondary market may even lead to unintended consequences. This is particularly true in the automobile market. For example, Corporate Average Fuel Economy Standards in the United States can be expected to raise the price of new vehicles and delay scrappage of older and often more polluting vehicles (Jacobsen, 2013). In other countries, this effect is even more evident. In Denmark, the new vehicle registration tax nearly triples the price of vehicles, disincentivizing new vehicle purchases and leading to a much older fleet than would be expected given the high per capita income of the country.

There are important dynamic considerations in consumer decisions that mediate how policies affect the allocation of new and used durable goods. The stock of vehicles is persistent and vehicles depreciate in value over time. Moreover, transaction costs lead to inertia in consumer holdings due factors such as costly search or asymmetric information. These dynamic considerations are particularly important for the welfare consequences of policies addressed to both the primary and secondary markets.

This paper develops a tractable life-cycle model of vehicle ownership, vehicle choice, and usage. The model can for example be used to examine the effects of a proposed reform that reduces the exceptionally high Danish vehicle registration tax and replaces it with road user charging, in which drivers pay a tax based on the number of kilometers driven. We model the dynamic considerations of the consumer in a framework that includes macroeconomic conditions, aging, replacement, and scrappage. Using this framework, we study the non-stationary equilibrium in the secondary market and can replicate the waves of vehicle prices and ownership decisions corresponding to the business cycle that are observed in the data. We estimate our model using detailed data from the Danish registers on all vehicles in Denmark and their odometer readings matched to individual and household-level demographics. These data contain longitudinal information on income, wealth, labour market status, household composition, distance to work, occupation, and family patterns, as well as information on all vehicle transactions and suggested depreciation rates at the make-model-vintage level.

This paper contributes to several strands of the literature. The proposed policy affects the vehicle market, a well-studied market in the economics literature, with significant work on product differentiation and consumer choice of new vehicles (Bresnahan, 1981; Berry, Levinsohn and Pakes, 1995; Goldberg, 1995; Petrin, 2002). These seminal papers allow for general patterns of substitution across differentiated products, but do not model secondary markets or the dynamics of the consumer decision process. Economists have demonstrated the importance of secondary markets for the allocation of new and used durable goods (Rust, 1985c; Anderson and Ginsburgh, 1994; Hendel and Lizzeri, 1999a,b; Stolyarov,

2002; Gavazza, Lizzeri and Roketskiy, 2014), as well as the influence of durability on the dynamics of vehicle demand (Adda and Cooper, 2000a; Stolyarov, 2002; Esteban and Shum, 2007; Chen, Esteban and Shum, 2013). This paper models secondary markets and the dynamics of consumer decisions in the context of a major proposed policy reform using impressively detailed household-level data. Schiraldi (2011) models the consumer’s dynamic decision process to estimate transaction costs and the effects of a counterfactual scrappage subsidy in Italy, but does not model counterfactual equilibrium prices in new and used vehicle markets.

Since Berkovec (1985) economists have estimated numerical equilibria in new and used vehicle markets. Rust (1985c) estimates a stationary equilibrium in new and used vehicle markets with an equilibrium price function that matches the distribution of supply with the distribution of demand. Konishi and Sandfort (2002) prove the existence of a stationary equilibrium in the presence of transaction and trading costs. Stolyarov (2002) and Gavazza, Lizzeri and Roketskiy (2014) estimate stationary equilibria with transaction costs that match several key features of the U.S. automobile market. One key assumption in these papers is a discrete uniform distribution of vehicles in each age cohort. Adda and Cooper (2000b) demonstrate that the age distribution is non-stationary: macroeconomic shocks and gasoline price shocks create “echo effects” or “waves” in the age distribution. We model equilibria in the automobile market that is a function of both macroeconomic conditions and gasoline prices, allowing us to capture these waves in the age distribution of vehicles.

By examining the welfare effects of a key policy reform, our paper also contributes to the literature examining environmental policies in vehicle market. For example, Bento, Goulder, Jacobsen and von Haefen (2009) use a static model of consumer demand and a Bertrand oligopoly model for automobile supply to examine the welfare and distribution effects of vehicle taxes in the United States. Jacobsen (2013) builds on this modeling framework to examine the effects of Corporate Average Fuel Economy Standards in the United States. These papers model both vehicle choice and usage decisions to provide useful policy insight, but abstract from the intertemporal dependence of consumer decisions. Our paper also uses more comprehensive data that allows us to model the impact of macroeconomic conditions on the vehicle purchase decision. Gillingham (2012) develops a two-period vehicle choice and usage model to examine the effects of gasoline taxes and policies that change the price of new vehicles. The focus in Gillingham (2012) is on estimating the rebound effect, i.e., the additional driving in response to a policy that raises fuel economy. A major contribution of our paper is that it develops a tractable model of dynamic consumer choice to estimate primitives that allow us to simulate the counterfactual equilibrium and accordingly, the effects of an important policy reform that is actually being considered.

There are a number of attractive features of our approach for examining the effects

of the proposed reform. First and most importantly, the structural parameters have a clear interpretation from the theoretical model, allowing for counterfactual simulations to examine the welfare effects of the proposed reform. Our data allow us to obtain aggregate demand for vehicle investments, fuel consumption, and usage by aggregating individual demands resulting from consumer dynamic optimizing behavior. Furthermore, our empirical setting and data contain several reforms that provide plausibly exogenous variation to identify our structural parameters.

We find that our model can not only replicate waves in the observed data due to business cycles, but can rationalize the vehicle choice and usage behavior in Denmark. We conduct a simple counterfactual experiment of a reform that reduces the new car prices and raises fuel prices. The simulations show that both the model with and without equilibrium prices predict a shift towards younger cars. However, in the equilibrium-version, this shift occurs at the cost of accelerated scrappage of the older cars. This behavior is driven by the equilibrium prices; without equilibrium prices, the reform increases shifts demand towards newer cars for all households, regardless of which car they currently own. When prices adjust to equate demand and supply, demand will drop relatively more for cars ages that are abundant. Thus, the counter-movements of equilibrium prices imply that the demand-response to the reform will depend on the individual household’s car state as well as the aggregate car stock. In the simulation we see a large group of old vintages where the reform depresses prices for those car ages so much that it leads to a spike in scrappage. This is the type of behavior that is documented empirically by e.g. Jacobsen (2013). The ability to study the interplay between car taxation, the car stock and the macro cycle is a primary innovation of this model.

The remainder of the paper is structured as follows. The next section provides background on the institutional setting and discusses our dataset. Section 3 develops our dynamic model of consumer purchase, vehicle type, replacement, and usage choices. Section 4 discusses our estimation approach and the data. Section 5 describes how we solve for the non-stationary equilibrium. Section 6 presents our results and Section 7 concludes.

2 Background and Data

This section provides background on the relevant policy questions that this model was designed to address, and describes the data we use to estimate the model, and provides a deeper review of the literature we built upon, highlighting the new contributions in this thesis. Section 2.1 summarizes the institutional setting in Denmark and several significant policy changes that occurred during our sample period. Sections 2.2 and 2.3 discuss the data sources used to estimate the model and provides a descriptive summary of the main features of the data we hope to capture in our model. Finally section 2.4 provides a fuller review of four separate literatures our model builds upon and was inspired by, and

summarizes the areas where we contribute to each of them.

2.1 Institutional Setting

Denmark provides a very useful empirical setting for examining policies that affect the new vehicle registration tax and the operating cost per kilometer driven. Vehicle taxation in Denmark currently is made up three components: a one-time registration tax when the vehicle first enters the Danish fleet, an annual tax, and fuel taxes. The registration tax is a very large proportional tax with a kink, where various deductions apply.¹ For example, in 2010 the tax was 105 percent of the first DKK 79,000 (about \$14,500) and 180 percent of the portion of the price exceeding the kink at DKK 79,000. The kink changes over time but the rates of 105 percent and 180 percent have remained stable.

There have been numerous changes over time in the registration tax, that provide exogenous variation to help us identify our structural primitives. There have been three reforms from 1992 to the present with an increasing focus on creating incentives for households to purchase more fuel efficient vehicles. Data on the fuel efficiency of new vehicles is available from the first reform in 1997. This reform set the annual tax for all vehicles first registered prior to July 1, 1997 according to the weight of the vehicle. At the same time, it set the annual tax for all vehicles registered after July 1, 1997 according to the fuel economy of the vehicle (in kilometers per liter). The motivation behind this reform was to tax older vehicles for wear and tear on the road and incentivize households to purchase more fuel-efficient new cars.

In 2000, deductions in the registration tax were introduced for vehicles in the higher end of the fuel efficiency scale (above 25 km/l). Therefore, only a very limited fraction of the vehicles sold in that year were actually affected by the reform. In the 2007 reform, these deductions were expanded so that all vehicles have their registration tax depend on fuel efficiency according to a piecewise linear schedule. If the vehicle has a fuel efficiency (FE) of more than 16 km/l, it receives a deduction of $4,000(\text{FE}-16)$, and if it has a fuel economy less than 16 km/l, the tax is increased by $1,000(16-\text{FE})$. Not surprisingly we see a very strong response at the extremes: The market share of the most fuel efficient cars increased from 8.1 percent prior to the reform to 50.4 percent at the end of the period in 2011 whereas for cars driving 16.6 km/l or less it decreased from 71.3 percent to 19.4 percent.

The Danish Ministry of the Environment pays out a scrappage subsidy for cars that are scrapped in an environmentally sound way by an authorized scrap yard. The subsidy was put in place on July 1st, 2000, and amounts to 1,500 DKK.

¹Examples of deductions include a reduction of the taxable value of the vehicle of DKK 3,750 if ABS brakes are installed and a reduction of DKK 12,000 from the final tax if the vehicle drives 19 to 20 km per litre of gas.

2.2 Data

The dataset used in this paper draws on many different Danish sources. At the core of the dataset is information on the fleet of vehicles registered in Denmark is available from Statistics Denmark in the database *bildata*. The main source for the database is the Central Register of Motor Vehicles. The database keeps track of nearly all vehicles in Denmark and in particular all private personal vehicles.² For each vehicle we have the motor register’s vehicle identification number (VIN) and the owner’s CPR number, which uniquely identifies all individuals in Denmark.³ This register not only contains basic vehicle information, but also allows us to track ownership over individual vehicles over time.

Socioeconomic data for the owners of vehicles comes from various Danish registers. These contain the full Danish population in each year with the exception of Danes living abroad. The CPR number is given to any individual taking residence for longer than 3 months in Denmark (6 months for Nordic or EU citizens) and is used in nearly all dealings with official authorities from education and taxation to the purchase of medicine and criminal records. Thus, the dataset includes detailed educational information, place of residence and time of movements, income and wealth information from the tax report (which for most employees is 3rd party reported). We merge in information on spouses and children to give an adequate picture of the household.

Another important vehicle register dataset contains information on the vehicle tests performed by the Danish Ministry of Transportation (MOT). There are three main types of tests, with the goal of ensuring that vehicles in Denmark are safe to drive. A registration test is performed when the vehicle is registered. Periodic tests are performed bi-annually from the fourth year since the car was registered and the rest of its lifespan. Customs tests are performed on imported used vehicles prior to their registration test when they are registered in Denmark. The most important variable from the MOT tests is the odometer reading, which allows us to track the usage of individual vehicles. Using the VIN, these odometer readings are merged with the vehicle register database. Note that for the first observation of a given VIN at a test, we assume that the odometer was at zero when the car was originally purchased. There are two possible exceptions to this; if the car was taken for test drives prior to the purchase, then that will have taken prior to the first registration, which occurs when the car is purchased from the dealer and registered to the consumer. The second is if the car was imported, which relates to the following data

²Exceptions that are not included in the register include for example company cars and military vehicles. For company cars, we instead observe a tax variable indicating whether an individual has access to a company car that can be used privately. This is the case for 3.4% of Danish households.

³Note that the VIN found on American vehicles differs from our VIN; in the US, the VIN can be used to back out much information about the car manufacturer etc. We also have access to the first 11 characters of the VIN number but we have found this variable to be unreliable in our dataset, inconsistent over time and many observations having VINs we cannot justify based on online databases.

issue.

One shortcoming of the vehicle data is that we do not observe the make year of the vehicle. Instead, we only observe the date of the first registration in Denmark. This means that if a used car has been imported, we are incorrectly classifying it as a zero year old car. However, imported used cars must also pay the Danish registration tax, which means that the net-of-tax new and used car prices in Denmark are generally lower than in other European countries (see Figure B.4). Therefore, importing of cars is not

Finally, the Danish Automobile Association (DAF) maintains a database of prices of vehicles by make, model, variant, year and vintage, allowing us to follow the value of used cars as well. The main limitation of these data is that we do not observe what additional equipment was purchased with the car. However, DAF does provide an informed guess of the typical price, as well as a high and low price, bounding the price range for that specific vehicle. DAF also provides the price a professional car dealer would pay and the price he would demand for a given vehicle, giving a proposed margin. The prices are highly reliable and are used by professional car dealers in setting the price of a used vehicle.

We define scrappage in our data as having occurred when a car's ownership spell ends and we do not observe a new one starting. The car may have been exported out of the country although exports are generally not a large concern because the high taxes in Denmark mean that used-car prices are fairly high internationally. We observe quite low scrappage rates in the first two sample years, 1996 and 1997, so to validate our data in terms of scrappage, we can compare the scrap rates to data on the number of scrappage subsidies paid out.⁴ We will discuss this issue in greater detail later.

2.3 Descriptives

We will now present some key descriptives for our estimation sample. We will focus on the main variables to be incorporated in the model, namely car characteristics, fuel prices, car ownership by household age and income and the discrete choices made by households.

The most important piece of descriptive evidence for this paper is the “waves” in the car age distribution shown in Figure 2.1. The waves appear as newly purchased cars travel through the age distribution of cars over time as they age. It is well-known that new car sales is one of the most volatile components of GDP, clearly showing the business cycle. Along the axis of calendar time for car age zero, we see the new car sales increasing in the boom in the late 90s, staying low during the brief recession in 2001–2003 before then again increasing in the following boom up towards the financial crisis. Then, as time moves forward these purchases travel through the age distribution along the diagonal, until they begin to die out as the car age approaches 20 and cars start to be scrapped.

While Figure 2.1 shows the car age distribution, this is not necessarily informative

⁴The data is available on the website www.bilordning.dk (accessed March, 2015).




Figure 2.1: Car Age Distribution Over Time: “Waves”




Figure 2.2: Purchases by Car Age Over Time

about how much trading takes place along the waves; it might be that the same owner holds on to a given car for its entire lifetime or that they are traded. Figure 2.2 show the number of purchases for a car of a given age in a given year. Firstly, we see that new car sales dwarf any of the other age groups, as would be expected. Secondly, we see the macro state clearly in the new car purchases since car sales are highly pro-cyclical. This fact is a key motivation for our modeling strategy; the macro shocks drive the new car sales which then travel through the age distribution as “waves”. Thirdly, we see that the waves can also be seen in the transactions, meaning that we see more trading for cars that are more abundant. This becomes more clear if we remove the new car sales from Figure 2.2, which we have done in Figure B.5.

Table F.2 provides summary statistics for key variables in the full dataset. In our empirical analysis, we will be aggregating to only two car types: gasoline and diesel cars. To construct the choice set, we aggregate the characteristics of the underlying cars by taking un-weighted averages within each of the two car types. Figure 2.3 shows the new




Figure 2.3: Car Characteristics over Time

car price in 2005 DKK and fuel efficiency in km/liter for the two types over the sample period. The figure shows that the new car prices have converged; a diesel car cost 15.5% more than a gasoline car in 1996, which had fallen to 1.6% by 2009. At the same time, the average fuel efficiency has increased relatively more for diesel cars than for gasoline cars.

Figure B.1 shows the real price of gasoline and diesel over time. Prices have been increasing for both types of fuel but we also note that the two prices appear to have converged over time. Figures B.2 and B.3 demonstrate how the composition of the fuel prices have changed over time, which shows that the changes have mainly been driven by the product prices; fuel taxes were increased slightly in 1996 and 2000 but were otherwise kept constant (i.e. the fixed part was kept constant and the proportional tax rate was not changed).

Our dataset allows us to paint a very complete picture of car ownership over the life cycle and for the full household. We will focus the number of cars owned and how it relates to household age and income. First, note that only 12.1% of the households in our sample owns more than 1 car (Table B.1). This is very low compared to the US but makes sense in light of the very high car prices (see Figure B.4). From 1996 to 2009, the share of no-car households has decreased from 49.1% to 37.2%, and the share of two-car households has also increased (from 6.3% to 14.4%). Like most of the famous models of car choice, our model will be a single-car model, which does not seem to be as critical given the fairly low share of multi-car households. However, since a major focus of this paper is to model the equilibrium of the used-car market, we do not wish to simply drop all these observations. Instead, we choose to treat multi-car households as independent decision-making units; when a household purchases an extra car, we create two observations for that year, where one keeps the original car and the other is counted as a household entering from the no-car




Figure 2.4: Number of Cars Owned by Household Age

state. The two observations will split the household income equally to ensure that the total amount of resources in the economy remains stable.⁵

Figure 2.4 shows the number of cars owned by the household age (defined as the male's age for couples). The figure shows that the ownership rate increases rapidly up through the 20s and then flattens by the late 30s where around 70% of households owning at least one car. As the household approaches retirement age, the share of no-car households increases somewhat and it appears that some 2-car households sell of one of their cars.

Next, we consider how car ownership varies with the income of the household. Figure 2.5 shows for each income decile, the percent of households owning zero, one, two or more than two cars. As expected, higher income is associated with a higher probability, and for incomes above the median, the share in the one-car category decreases as households start to be able to afford having more than one car.

We now consider the discrete car ownership choices that will be relevant to our model. If households own no car, they can choose to remain in the no-car state. If they have a car, they can either keep it, sell it or replace it. Recall that if they choose to buy an additional car, we will treat them as an additional household coming into the sample. Figure 2.6 shows for each income decile, the fraction of households choosing each of these discrete choices. Firstly, we see that over 80% of households in lowest income decile choose to remain in the no-car state and that this decreases to less than 10% for the highest income decile. Similarly, the probabilities of keeping and replacing the existing car increases. The probability of selling the car and going to the no-car state remains low throughout. Given that household income rises over the the life cycle, it is not possible from Figure 2.4 and 2.5 alone to determine whether the most important drivers are related to household age

⁵If the household once again becomes a one-car household, then the extra observation will count in the final year as having sold to go to the no-car state and will be deleted from future time periods.




Figure 2.5: Number of Cars Owned by Income Decile




Figure 2.6: Discrete Choice by Income Decile

(e.g. the presence of children) or income (e.g. leisure activities or work).

We now take the perspective of the cars being purchased. Figure 2.7 illustrates how long households hold their purchased car conditional on the car age at the time of purchase. For each ownership spell where the car was a years old at the time of purchase, we show the distribution of the lengths of the ownership spells. As expected, the figure shows that when a household purchases a young car, they tend to hold it for longer. Interestingly, the holding times go up after age 22; this is most likely due to the selection effect of vintage or specialty cars not being scrapped.

Finally, we show the scrappage in the data over time. Figure 2.8 shows scrappage by car age; we have pooled the sample and computed for each car age, the pct. of all cars at that age that are scrapped. Note that we truncate car age at 24, which is the maximum age used in the model. The figure shows that the mode of car scrappage occurs at car age




Figure 2.7: Years of Ownership by Car Age at Purchase




Figure 2.8: Scrappage by Vehicle Age

22, after which scrappage declines somewhat. This is most likely due to a selection effect where specialty or vintage cars are kept very long while normal cars are scrapped earlier. We also note that scrappage is markedly higher in even years; this coincides with the *test years*. In other words, the pattern is consistent with an individual taking his car to the inspection test and deciding to scrap the car if it fails the inspection and is deemed unfit to drive. Figure ?? further shows that there is still considerable trading activity for the higher age groups; around a third of all ending ownerships are terminated in a transaction rather than a scrappage for the highest age groups.

Figure 2.9 shows the number of cars being scrapped in each year by the car age. When compared to the waves in Figure 2.1, we can see the scrappage spike in 2000–2005 as being explained by cars from the boom in the 1980s being scrapped and that wave dying out in the car age distribution. An important feature of the data that becomes clear from Figure




Figure 2.9: Scrappage by Year and Car Age

2.9 is that the age distribution changes; so while Table B.2 indicates that the number of cars being scrapped each year is relatively stable over time, this masks the fact that the age composition of cars being scrapped in the late 2000s is quite different from the ones being scrapped in the early 2000s, with younger cars being scrapped later.

In Appendix B.3, we go into more details about our scrappage data. Most importantly, we find too low scrap rates for 1996 and 1997 for that data to be believed (Table B.2). This means that we are only seeing a very small number of ownership periods ending prior to this. However, from 1999 the rate appears to be on par with the remainder of the period. We have been unable to discover the cause of this oddity in the data but we choose to use the data from 1999 and onwards.

We conclude the descriptive section by discussing some correlations between VKT and the state variables. Figure B.10 shows the VKT by the age of the car. The graph shows that driving is highest for four year old cars (just over 55 km per day on average) and then declines almost linearly towards the 20 year old cars, that are driven just over 35 km per day on average. This unconditional relationship might reflect a number of other factors correlated with the car age. Figure B.11 shows the corresponding graph with real income instead of the car age. The figure indicates that driving increases in income for the largest part of the data but decreases for very high income levels. Table B.3 shows regressions of VKT on different sets of controls for the full sample. We find very large price sensitivities, unless we control for a diesel dummy in the driving equation. Gillingham and Munk-Nielsen (2015) provide evidence that diesel drivers tend to drive much more and be more price responsive than gasoline car drivers.

2.4 Previous literature

This paper builds on and contributes to four different literatures: 1) a literature on discrete/continuous choice of durable goods, where there is a discrete choice of type of durable (including attributes such as the durable's energy efficiency) and where the continuous

choice represents usage of the durable (such as driving in the case of automobiles), 2) a literature on numerical and theoretical models of equilibrium in automobile markets, and 3) a literature on structural estimation of dynamic choice models, including dynamic discrete choice models applied to choice of automobiles. We provide reviews of each of these literatures below. Most of these literatures emerged after the oil price shocks and concern about permanently higher fuel prices in the late 1970s. Since that time fuel prices have increased but not as dramatically as once feared. Instead attention has refocused more recently on concerns about the effects of vehicle emissions on the environment, with particular concern about CO₂ emissions and its impact on global warming.

2.4.1 Discrete-continuous Models of Durables

This literature goes back to Dubin and McFadden (1984), where households choose electrical appliances taking into account their future usage of the durable. The key insight is that the usage falls out of Roy’s identity. Models of this type place strict cross-equation restrictions on the parameters of the model in the sense that they force the consumer to be time-consistent in treating money in the same way when making the purchase decision and the usage decision.⁶ Earlier work on discrete-continuous choice models tended to use two-step approaches (Manning and Winston, 1985; Goldberg, 1998; West, 2004). More recently, applications to car choice and usage have featured simultaneous estimation of both choice margins (Feng, Fullerton and Gan, 2005; Bento, Goulder, Jacobsen and von Haefen, 2009; Jacobsen, 2013). For example, Bento, Goulder, Jacobsen and von Haefen (2009) use their model to analyze the distributional impacts of fuel taxes in the US. In their model, the discrete choice is the car choice and the continuous choice is how much to drive the car. Gillingham (2012) also uses a discrete-continuous model applied to car choice and use and focuses on the selection of consumers based on anticipated driving and allowing for selection on observed and unobserved factors. Munk-Nielsen (2015) applies a similar model to new car sales in Denmark to study the costs of environmental taxation. The model admits an estimate of the so-called “rebound effect”, the effect on driving from an exogenous increase in fuel efficiency. This important policy parameter has been widely discussed and estimated (e.g. Small and Van Dender, 2007; Hymel and Small, 2015).

Engers, Hartmann and Stern (2009) study the interrelationship between vehicle usage and price depreciation in the used car market. They argue that “changes in a vehicles net benefits, proxied by annual miles, explain the observed decline in used car prices over the vehicle’s life.” (p. 29). They find that households drive fewer miles per year the older their car is, and estimate a structural model of household choice of driving and vehicle type that differs from the literature surveyed above. They conclude that “the structural model of household mileage decisions better explains the observed price decline in used car

⁶Whether consumers accurately take into account future savings in fuel costs is widely discussed in recent empirical work (Allcott and Wozny, 2012; Busse, Knittel and Zettelmeyer, 2013).

prices.” and “the observed decline in used car prices as a vehicle ages is best explained by decomposing the age effect into three components: the direct aging effect, the household portfolio effect, and the household demographics (or car turnover) effect.” (p. 30).

2.4.2 Models of Equilibrium in Automobile Markets

This paper builds on a theoretical and empirical literature for modeling equilibrium in the market for automobiles. The earliest work that we are aware of was by Manski (see Manski (1980), Manski and Sherman (1980) and Manski (1983)). We believe Manski’s original work stimulated the subsequent chain of research on micro-econometrically estimable equilibrium models of the automobile market, and his work provided both theoretical models of equilibrium in new and secondhand auto markets, and numerical calculation of equilibrium prices and quantities that demonstrated how these models could be used for policy forecasting of a wide range of policies of interest. Manski and Sherman (1980) did their pioneering work in an environment around the first large oil price shocks in the late 1970s when it first became clear that gasoline prices would inevitably rise and there would be a demand for increasingly fuel efficient vehicles. They concluded that “our initial research on developing and applying a disaggregate modeling approach to forecasting future motor-vehicle sales and holdings has proved highly encouraging. Our results are really the beginning of an ongoing need to analyze and monitor the motor vehicle market through the 1980s. . . . With an eye toward improvement of our models, future work should seek to further illuminate the linkages that connect household behavior in choosing motor vehicles and other vehicle-related decisions. In particular, a joint analysis of ownership level, the composition of holdings, and vehicle use would be a valuable contribution.” (p. 103).

The contributions of Manski and coauthors inspired further work such as the 1983 PhD thesis research by James A. Berkovec at MIT (subsequently published as Berkovec (1985)) who followed the footsteps of Manski and Sherman (1980) and developed the second microeconomically estimated and numerically solved large scale equilibrium model of the new and used car markets that we are aware of. The contributions of Manski and coauthors, and Berkovec was extremely advanced given the limits of computing power at the time, and still represents the closest point of departure and template for our own work in this area.

Berkovec described his model as a “short run” equilibrium model as it was based on a model where consumer *expectations* about depreciation rates of their vehicles was estimated econometrically using data on new and used car prices in 1978. Berkovec assumed that consumers choose vehicles based on a quasi-linear utility function that is an additively separable sum of a utility for car attributes (with declining utility for cars of older ages) less the disutility of the “expected capital cost” of owning the vehicle. The expected capital cost is essentially the expected depreciation of holding the vehicle plus

maintenance costs, using the econometrically estimated depreciation rates.

Berkovec assumed that consumers choose a vehicle that maximizes their utility where the price of the vehicle enters via the expected capital cost. He used a nested logit discrete choice model that allow for patterns of correlation in the unobserved components of the utility of a vehicle that capture patterns of similarity in the unobserved characteristics of vehicles in 13 different car classes that he used in his analysis (e.g. luxury cars, compact cars, vans, pickups, etc). He developed and estimated separate microeconomic nested logit discrete choice models for households that own 1, 2 and 3 cars, respectively.

Using the microeconomically estimated choice model, Berkovec constructed an “expected demand function” for vehicles of different ages and classes by summing the estimated discrete choice probabilities for cars of each age and class. He defined an equilibrium to be a vector of prices (with one price for each possible age and price of car) that equates the expected demand for vehicles of each car age and type to the actual supply of such vehicles, net of scrappage. Berkovec used a probabilistic model of vehicle scrappage due to Manski and Goldin (1983) where the probability a vehicle is scrapped is a decreasing function of the difference between the second-hand price of the car (net of any repair costs) and an exogenously specified scrap value for the vehicle. This implies that, except for random accidents, there is very little chance that new cars are scrapped, but the probability a used car is scrapped increases monotonically with the age of the car.

Berkovec used Newton’s method to compute the equilibrium prices in the market. For the problem he analyzed there were 131 vehicle class/age price categories. At the time Berkovec did his work, inversion of the 131×131 Jacobian matrix of excess demands necessary to implement Newton’s method was a much bigger computational challenge than it is today. Berkovec showed that the Jacobian matrix had special structure he called “identity outer product” that enabled him to invert the Jacobian via inverting a smaller 48×48 matrix and doing some additional matrix vector multiplications. Though Berkovec’s paper did not discuss the equilibrium prices implied by his model, he concluded that “Overall, the simulation model forecasts appear to do reasonably well for the 1978-1982 period. Although there are discrepancies in specific areas (as would be expected because of underlying macroeconomic fluctuations), the general trends evident in the data would seem to be captured in the forecasts.” (Berkovec (1985), p. 213).

Subsequent work on empirical equilibrium models of the automobile market includes Bento, Goulder, Jacobsen and von Haefen (2009) who estimated a micro level discrete/continuous model of automobile driving and model/age choice using a sample of 20,429 U.S. households from the 2001 National Transportation Survey. Using microaggregated demands from the estimated discrete choice model, they numerically solved for equilibrium in the new and used car markets for a total of 284 composite age/model vehicle classes. They used their model to predict the impact of a 25 cent increase in the U.S. gasoline tax. Their model predicts that most of the response to this tax increase is via reduced driving: they

found negligible longer run substitution to more fuel efficient vehicles or to non-automobile modes of transportation: “the size of the vehicle fleet falls about 0.5 percent” but “The impacts on new and used car ownership differ substantially over time. In the first year of the policy, the reduction in vehicle ownership comes largely by way of a decline in new car purchases. However, the ratio of fuel economy of new to old vehicles increases over time, and the increased gasoline tax gives greater importance to fuel economy. As a result, the decline in new car ownership is attenuated over time, and by year 10 the reduction in car ownership applies nearly uniformly to new and used vehicles.” (p. 697).

A separate more theoretically oriented thread of the literature focused on modeling the role and benefits of the secondary market for automobiles (or more generally for other durable goods) in frameworks where the dynamics of trading were more explicitly modeled relative to the work surveyed above. Rust (1985c) established the existence of a stationary equilibrium in a market for new and used durable assets that provided a theoretical rationale for the conditions under which consumer choice of a stochastically deteriorating durable good (e.g. an automobile) involves a trade-off between the utility provided by the durable and its expected price depreciation. He showed that the key condition for this to hold is that there are *zero transactions costs* in the market. When this holds, the optimal strategy for each consumer in a stationary equilibrium (i.e. one where there are no macro shocks or other time-varying factors altering the prices or quantities of vehicles in the market) involves trading each period for the preferred age/condition of car $x^*(\tau)$ where τ is a parameter that indexes heterogeneity among consumers, e.g. differential preferences for “newness” or different degrees of wealth that affect consumer willingness to pay for newer/better condition durables. Similar to Berkovec, Rust assumed that per period preferences for durables are quasi-linear in the attributes of the durable and in income, which is a simple way of representing preferences for all other goods without explicitly modeling them.

Unlike Berkovec, who considered a discrete set of car classes and ages, Rust modeled a durable as having a *state* x , where $x = 0$ corresponds to a brand new durable good and higher values of x correspond to more deteriorated, less desirable older durable goods. For example in the case of automobiles, x might be the odometer on the car, and consumers may be more concerned about the level of wear/tear on a car as represented by the odometer value x than the discrete age of the car. In this framework an equilibrium requires finding a *price function* $P(x)$ that clears the market (i.e. sets the demand for durables of each condition x equal to the supply). The supply of durables is represented by another function $S(x)$ that Rust called a *holdings distribution* — it is the fraction of durables in the economy with condition less than or equal to x . If each vehicle deteriorates stochastically, according to a Markov transition probability $f(x'|x)$ (where x' is the condition of the durable next period given that its condition is x this period), then in a stationary equilibrium with a continuum of agents, Rust showed that there will be a

stationary distribution $S(x)$ that is related to the *invariant distribution* of the Markov transition probability $f(x'|x)$.

Rust assumed that there was an infinitely elastic supply of new durables at an exogenously fixed price $\bar{P}$ and an infinitely elastic demand for scrap at an exogenously fixed price $\underline{P} < \bar{P}$, and provided sufficient conditions for a *stationary equilibrium* $(P(x), S(x))$ that satisfies the conditions 1) $P(0) = \bar{P}$, 2) there is a *scrap threshold* $\gamma > 0$ such that $P(x) = \underline{P}$ if $x \geq \gamma$, and 3) there is equilibrium for all conditions $x \in (0, \gamma)$, i.e. the fraction of consumers who wish to hold a durable with condition less than or equal to x equals the stationary holdings distribution $S(x)$. Rust called condition 3) *stock equilibrium* i.e. it amounts to the usual condition that the demand for every condition of car equals the supply in the case of a continuum of goods x . Rust also showed that a stationary equilibrium also implies a condition he called *flow equilibrium* i.e. the fraction of used durables that are scrapped each period equals the fraction of the population that buys new durables. This implies that the overall stock of durables in the economy is not changing over time.

In subsequent work Rust (1985a) showed that when applied to the automobile market, a calibrated version of the stationary holdings distribution implied by Rust (1985c) provides a good approximation to the joint distribution of ages and odometer values (x) in the US economy using data from the 1970s. He also showed that for a range of plausible utility functions for consumers, the stationary equilibrium resulted in *convex price functions* $P(x)$, which implies the rapid early depreciation for new cars and the slower depreciation for older cars that we observe in most auto markets. However the assumption of zero transactions costs is an unrealistic feature of his model as it implies that it is optimal for consumers to trade every period for their preferred condition $x^*(\tau)$ and this is something we definitely do not observe in real world automobile markets. When there are transactions costs (which are separate from *trading costs*, i.e. the difference between the list price of a car x a consumer wishes to buy, $P(x)$ and the list price $P(x')$ of an older car x' that the consumer wishes to trade in exchange for the newer car x), Rust (1985c) showed that the optimal strategy generally involves keeping the current durable for multiple periods. In a stationary market the optimal trading strategy in the presence of transactions costs consists of two thresholds $(\underline{x}^*(\tau), \bar{x}^*(\tau))$ where $\underline{x}^*(\tau) < \bar{x}^*(\tau)$ and $\underline{x}^*(\tau)$ is the condition of the optimal replacement durable that a consumer will choose whenever he/she replaces their current durable x , but $\bar{x}^*(\tau)$ is a *selling threshold* and it is not optimal to replace the current durable x until x exceeds the selling threshold $\bar{x}^*(\tau)$. When $x > \bar{x}^*(\tau)$ the consumer of type τ sells their current durable x for $P(x)$ and buys a replacement durable of condition $\underline{x}^*(\tau)$ for price $P(\underline{x}^*(\tau))$. Notice that generally $\underline{x}^*(\tau) > 0$, so the replacement durable is generally not a brand new durable good $\underline{x}^*(\tau) = 0$. However consumers who are sufficiently rich or who have a sufficiently strong preference for “newness” will replace their used durable with a brand new one.

Establishing the existence of a stationary equilibrium in the presence of transactions costs is a much more daunting undertaking due to the possibility that there may be consumer types τ who desire to buy a slightly used but not completely brand new durable $\underline{x}^*(\tau)$ yet, there may not be any consumer type τ' whose optimal strategy involves buying a brand new durable whenever they replace their old one who has a selling threshold $\bar{x}^*(\tau') < \underline{x}^*(\tau)$. That is, there is no automatic guarantee that there will be someone willing to sell a sufficiently new durable good to another consumer who wishes to buy a very new but not brand new durable good — perhaps to try to take advantage of the the rapid early depreciation in durables and buy an “almost new” durable good for a price that is much lower than the price of a new durable good $\bar{P}$. However Konishi and Sandfort (2002) did prove that a stationary equilibrium can exist in the presence of transactions costs under certain conditions. Their proof shows that it is possible for the equilibrium price function $P(x)$ to adjust to prevent any of the coordination failures of the type discussed above, i.e. where some consumer type τ' wishes to buy some sufficiently new durable good $\underline{x}^*(\tau')$ but no other consumer type τ is willing to sell their used durable to that consumer.

There is also a growing literature on the interaction between the market for new and used durable goods. Generally, the secondary market increases the lifetime of a durable good by facilitating a string of trades from customers who prefer newer durables and sell to a sequence of customers who are either poorer or who have weaker preferences for new durable goods relative to older ones. If the secondary market does not exist, each consumer can of course simply “buy and hold” — that is, buy brand new durable goods and hold them until they decide to scrap the old one and then buy another brand new replacement. Rust (1985c) showed that if consumers are homogeneous, they are indifferent between following such a buy and hold strategy, or trading each period for a preferred durable good in the secondary market. Thus, the existence of a secondary market does not produce any net welfare gain when consumers are homogeneous. However if consumers are heterogeneous, there is a welfare gain from the existence of a secondary market, and durable goods will have a longer lifespan on average when there is a secondary market than when it does not exist. Intuitively, the secondary market enables a chain of “hand me downs” of an aging durable good that would not be possible in the absence of a secondary market, and hence durables will live longer before they are scrapped when a secondary market exists, and consumers will be strictly better off compared to situation where there is no secondary market. In fact, Figure B.6 indicates that the most common a 15 year old car is to have had five owners.

However a secondary market is not necessarily desired by producers of new durable goods, because it allows consumers to keep used durable goods longer. Since used durable goods serve as a substitute for new durable goods, the existence of a secondary market limits a firm’s ability to extract rents from consumers via sales of new durable goods,

i.e. in the “primary market.” A standard solution to this problem is, in the case of a monopolist producer of durable goods, for the monopolist to rent rather than sell durable goods. Then the monopolist has the ability to control when durables are scrapped and extract rents from consumers without the distortions caused when durables are sold. When a monopolist sells new durables but attempts to set a high price, consumers react by keeping their used durables longer to reduce how frequently they have to replace their durables and thus pay the high price to the monopolist. However if a limitation to rental contracts only is not feasible, Rust (1985d) showed that a monopolist producer of new durable goods has an incentive to limit competition provided by the existence of a market for used durable goods by engaging in “planned obsolescence” — i.e. selling new durable goods that deteriorate more quickly than would be optimal under a social planning solution. In extreme cases the monopolist might even find it optimal to kill off the secondary market by producing goods with zero durability.

Esteban and Shum (2007) and Chen, Esteban and Shum (2013) developed empirically implementable models of equilibrium in new and used automobile markets and used these models to study the effect of the existence of a secondary market on oligopoly competition between new car producers in the primary market. Esteban and Shum (2007) formulated a model of oligopoly competition in new car markets under the assumption that a secondary market exists and there are zero transactions costs. Under their assumptions, demand for various new car models are linear functions of price, and the firms’ profit functions are quadratic functions of current and future production levels, which implies that a Markov perfect equilibrium exists in strategies that specify the auto companies’ production quantity decisions that are linear functions of a vector of the stock of cars produced prior to the current period that are still traded in secondary markets. Though the authors reported “difficulty of the theoretical model in generating price patterns similar to those observed in the data” (p. 345), they are able to use their model to study the effect of a temporary elimination of the secondary market on production decisions in the new car market. “Overall, we find that aggregate new-car production would increase by 12.08% for the 1987–1990 time frame were the secondary market to disappear temporarily.” (p. 349).

Chen, Esteban and Shum (2013) estimated a model of dynamic oligopolistic competition in the new car market allowing for the existence of a secondary market in each car brand (make/model) sold in the primary market, and allowing for the possibility of positive transactions costs. Their econometrically estimated transaction cost was \$4,400, and they note that “This is corroborated by the Kelley Blue Book, which indicates that, typically, the difference between the trade-in value of a used car (seller’s price for consumers) and its suggested retail value (buyer’s price) — which may serve as a proxy for the transactions cost — is in the \$3,000 to \$4,000 range.” (p. 2922). They conduct counterfactual experiments by varying the transactions cost parameter from a value large

enough to result in the closure of all secondary markets (for transactions costs larger than \$8,000), to a value of \$0, which corresponds to the case of a “frictionless” and active secondary market with no transactions costs. They find that relative to the equilibrium where there are no secondary markets, the case where there are active secondary markets with no transactions costs lowers firms’ profits in the primary market by 35 percent. They also find that “when the secondary market becomes more active, firms have a stronger incentive to make their cars less durable.” (p. 2929).

The final study on equilibrium in automobile markets that is most relevant to this paper is Gavazza, Lizzeri and Roketskiy (2014). The focus of their analysis is to quantify the welfare benefit of the secondary market and to investigate the effect of transactions costs on consumer trading and welfare. They formulated and numerically solved a dynamic model of vehicle holding that allows for the presence of transactions costs, and similar to the Chen, Esteban and Shum (2013) study, they used a discrete state model where automobiles are distinguished by their discrete age t rather than the continuous state framework that Rust (1985c) and Konishi and Sandfort (2002) used. Rather than focusing on the effects on profits of firms in the primary market, the focus of Gavazza, Lizzeri and Roketskiy (2014) was on welfare of consumers in the secondary market, and how welfare is affected by changes in transactions costs. They find that a calibrated version of their model “successfully matches several aggregate features of the US and French used car markets.” and that “Counterfactual analyses show that transactions costs have a large effect on the volume of trade, allocations, and the primary market. Aggregate effects on consumers surplus and welfare are relatively small, but the effect on lower-valuation households can be large.” (p. 3668).

While our review of the literature shows that there has been tremendous progress in both theoretical and empirical modeling of equilibrium in automobile markets, one of the gaps in the literature is the absence of work on modeling the effects of macroeconomic shocks. Since automobiles are among the most expensive durable goods outside of housing, it should not be surprising that macroeconomic fluctuations can have a huge effect on the timing of household purchases of new cars. In particular, when the economy is in a recession or about to go into recession, households worry about heightened risks of unemployment if they have not experienced unemployment already. Precautionary motives as well as tightened budget constraints appear to induce customers to hold onto their existing durables longer and wait to replace them until better times when they start to have more optimistic expectations about their employment prospects and earnings potential. We have already seen evidence of this in the descriptive graphs in section 2.3. Other analyses that have found similar effects include Adda and Cooper (2000a) and Adda and Cooper (2000b). However these studies have not modeled equilibrium in the primary and secondary markets in the presence of macro shocks. The cyclical variations in purchases of new cars generate slowly evolving “waves” in the stock of used cars as we illustrated

in our descriptive analysis of the Danish data in section 2.3. Prices in the secondary market must adjust dynamically to enable the wave in the “supply” of used cars from a previous macroeconomic boom period to match the demand. Thus, both quantities and prices in an automobile market that is subject to macroeconomic shocks do not satisfy the conditions for “stationary equilibrium” that has been the focus of analysis in virtually all of the existing literature that we are aware of.

A major reason why there has been little work on modeling equilibrium in a non-stationary environment with macro shocks and other time-varying factors affecting consumer demand for automobiles is due to the complexity in modeling the dynamics of equilibrium prices in the presence of a dynamically evolving stock of vehicles in the economy. Since the stock of vehicles that have not been scrapped that have been inherited from the previous period affects the supply of various ages of vehicles that will be supplied to the market, it follows that potentially consumers would need to know the entire age distribution of the vehicle stock to help predict market prices and how they will co-evolve over time along with the macro economic shocks and other time varying variables such as fuel prices that affect new car purchase, scrappage of old cars, and decisions on whether to sell or keep existing used cars. In principle, a high dimensional object — the entire age distribution of the automobile stock — needs to be one of the “state variables” that individuals need to keep track of to improve their forecasts of future auto prices. However due to the well known “curse of dimensionality” of dynamic programming, it becomes computationally infeasible to incorporate such high dimensional state variables in consumers’ optimization problems.

In this paper we follow an approach of Krusell and Smith (1998) that avoids the curse of dimensionality of carrying the entire age distribution of cars as a state variable in the model and instead using “summary statistics” to capture movements in this distribution over time. In the problem Krusell and Smith (1998) studied, consumer heterogeneity implies that it is generally necessary to know the entire distribution of wealth in the economy to determine interest rates, which in turn affect individual consumers’ savings decisions. However they showed that consumers can make highly accurate forecasts of future interest rates if they only keep track of the *mean value of wealth* (i.e. the mean of the distribution of wealth). Specifically, they found that the R^2 of regressions of current interest rate on mean wealth holdings in the economy was very high: typically over 97% in the numerical solutions and simulations of their model. This suggests that it is not necessary to confront the huge computational burden of carrying the entire distribution of wealth as a state variable to provide good forecasts of future interest rates.

In our paper we follow their insight and do not attempt to carry the entire distribution of car types and ages as a state variable in our dynamic programming model that we assume consumers solve to determine their holdings and trading decisions for vehicles. In fact, we go a step further and do not even attempt to use the mean ages of different vehicle

types (in analogy to what Krusell and Smith did) as state variables that consumers use to forecast future automobile prices. Instead we assume that sufficiently good forecasts can be obtained using a flexibly parameterized price function of the form $P(\tau, a, p, m)$ where τ is the type of car, a is the age of a vehicle, and (p, m) capture the current fuel price and macro state (which are assumed to evolve as an exogenous Markov process). We use a flexibly parameterized price forecasting function and find that it enables consumers to provide very good forecasts of future auto prices for different ages and types. It appears that there are high substitution elasticities for demands for vehicles of different ages of a given type, as well as high substitution elasticities for the decision to sell existing used cars, so even when there are pronounced “waves” in the stock of vehicles caused by macro shocks, these waves do not result in pronounced waves in the prices of vehicles due to the high substitution elasticities. That is, it is not necessary for prices to adjust dramatically over time to equate supply and demand for cars of different ages in response to various shocks and dynamic factors that lead to bunching and waves in the stock of vehicles.

2.4.3 Estimation of Dynamic Discrete Choice Models

This paper extends the literature by using fully dynamic models of individual households’ decisions about which vehicles to hold and to trade. As we noted above, most of the previous models in this literature ignored the fact that consumer decisions about automobiles are inherently dynamic choices. Previous empirical models of household choices such as as Berry, Levinsohn and Pakes (1995), Goldberg (1998), or Petrin (2002) focused on household choice of new vehicles only, and did so using a static discrete choice modeling approach. As we noted, the earliest empirical, disaggregate discrete choice models of equilibrium in the automobile market such as Manski and Sherman (1980) and Berkovec (1985) did estimate discrete choice models of holdings that allowed consumers to choose both new or used cars, but they also adopted a static choice perspective that treated consumers as making these choices every period, which would potentially result in excessive amount of trading of cars relative to what actually occurs.

As we noted above, when there are zero transactions costs, the assumption that consumers trade their existing cars for another new or used car every period can be rigorously justified, but this is clearly not an empirically realistic assumption. In the presence of transactions costs, households face a decision of whether to keep their current vehicle versus to trade for another new or used one. A literature on dynamic discrete choice, originating in the late 1980s (see, e.g. Rust (1985b)) provided the econometric methods for structural estimation of dynamic discrete choice models. This is a very flexible class of models that model probabilistic discrete dynamic choice models where the values or discounted utilities of choosing various discrete alternatives in each period are computed from the solution to a dynamic programming problem. These models can readily accommodate transactions costs and result in predicted behavior that is much closer to what we

actually observe, specially with regard to the frequency at which households trade their existing vehicle for another one.

Schiraldi (2011) is an example of the application of a micro-based dynamic discrete choice modeling approach to study holding and trading decisions of Italian households, but using *aggregate data*. Schiraldi takes prices of new and used cars in the Italian market as exogenously determined from the standpoint of individual households, and formulates and solves an individual households' optimal holding and trading strategy for vehicles to maximize their discounted expected lifetime utility. Using microaggregation of the individual consumer decision rules implied by the dynamic programming problem, Schiraldi was able to predict the aggregate vehicle holdings and trading patterns for the Italian economy as a whole, and he estimated the parameters of model using a simulated method of moments estimation strategy that finds parameter values for household preferences that enable the predicted, simulated moments to best match a set of actual moments characterizing aggregate holdings and trading of different types of vehicles over the period 1994 to 2004.

A novel feature of Schiraldi's analysis is to allow households to be "uncertain about future product attributes but rationally expect them to evolve, based on the current market structure." He captures this uncertainty using a variable he calls the "mean net augmented utility flow" arguing that "In a durable-goods setting, where the quality of the goods changes over time and there is the possibility of reselling, consumers maximize the utility derived from the good in any particular period net of the implicit rental price paid in that period to keep the good. Hence, the net augmented utility flow seems a natural index that captures the per-period quality adjusted by the price that consumers take into account to make their decisions." (p. 274). Schiraldi estimates significant transaction costs, with mean transactions cost equal to about €3200 in 1994 that slowly decline over time. It is interesting that these estimates are in the same ballpark as those provided by Chen, Esteban and Shum (2013) for the U.S. market.

We are not aware of any dynamic discrete choice model of household-level holdings and trading of vehicles that has been estimated using disaggregate household-level choice data, and believe this is one of the contributions of this paper. With aggregate data, it is impossible to observe how long individual households keep their vehicles before they are traded, nor it is possible to say much about the heterogeneity in vehicle choices, such as which types of households choose to hold newer cars and which choose older ones.

Cho and Rust (2010) provide an analysis of the vehicle trading behavior of a large rental car company. Unlike most households, rental car companies typically buy brand new vehicles and sell them very quickly, typically once the car is one or two years old. Due to the rapid initial price depreciation of vehicles that we observe in most car markets, this strategy would prove to be very expensive one and this is why we see few households except the very wealthiest ones following this type of trading strategy. Another interesting feature

of the rental car market that Cho and Rust (2010) point out is that rental car prices are typically *flat* as a function of age or odometer value, whereas they argue that predictions of most models of equilibrium in a competitive auto market are that rental prices should be declining functions of age or odometer value, reflecting the decline in prices and price depreciation rates in the used car market as a function of these variables. Cho and Rust (2010) argue that the trading strategy of the rental car company they analyze is also “too expensive” in the sense that it is suboptimal from a profit maximization perspective. Cho and Rust (2010) perform counterfactual analyses using a microeconometrically estimated dynamic programming model that show that the car rental company could significantly increase its profits by keeping its rental cars longer and discounting the rental prices of older rental vehicles to induce its customers to rent them. Their findings caused the rental car company to undertake a controlled experiment to verify the predictions of their model and the company did indeed find that profits did increase significantly from shifting to the recommended policy of discounting rental prices of older cars and keeping rental cars roughly twice as long as the company keeps its cars under its *status quo* operating policy.

There has been comparatively little work on solution and estimation of dynamic models of discrete and continuous choice beyond some recent work in this area such as Iskhakov, Jorgensen, Rust and Schjerning (2015) that is not directly applicable to our problem. A final contribution of this paper is to provide an estimable dynamic model with both discrete and continuous choices, where households make an optimal short run continuous choice of how much to drive their vehicle each period in response to their characteristics, the type of car they own, and the price of fuel, as well as a longer run dynamic choice of the type of car to own, which takes into account expectations of future driving and fuel prices, the household’s future income and age-varying life cycle needs for driving (e.g. the presence of children, retirement, etc) as well as future macro shocks that can affect both car prices and the household’s income.

3 The Model

In this section, we present the model. We first explain the state variables and decision variables as well as the model fundamentals. Then, in Section 3.1, we explain the household’s dynamic optimization problem, how we handle scrappage in the model and derive the Bellman equation. Finally, in Section 3.2, we present the utility specification and the optimal driving equation.

We estimate a finite horizon lifecycle model of automobile holdings, driving and trading decisions that features both vertical and horizontal product differentiation. Let τ denote the “type” of vehicle. We will assume there are a finite number of possible types, $\tau \in \{1, \dots, \bar{\tau}\}$. These can be thought of as a make-model combination or simply a vehicle class (e.g., “luxury,” “compact,” “economy,” “SUV,” “sport,” and “minivan”). In our

estimation, we use two car types according to the fuel types: gasoline and diesel.

To capture vertical product differentiation, we also distinguish the *age* of the vehicle, $a \in \{0, 1, \dots, \bar{a}\}$, where $a = 0$ denotes a brand new vehicle, and $a = 1$ a one year old vehicle, and $\bar{a}$ is the oldest vehicle in the market. For simplicity, we let $\bar{a}$ be a catchall class of all cars that are of age $\bar{a}$ or older. Thus, we index the set of cars that consumers in Denmark can choose from by (τ, a) where τ specifies a particular type of car and a denotes its age.

This formulation is very useful for the tractability of the model, but does abstract from changes in technology.⁷ We can note that changes in the real prices of cars are likely to be more attributable to macroeconomic conditions than a particular technology innovation, but this is an area for future work. There may also be a considerable degree of unobserved heterogeneity in used vehicles of a given age and type. For example, some have been driven more than others, and some are in better condition than others. However, Cho and Rust (2010) show that vehicle age and odometer readings are highly correlated and that once age is included as a predictor of car prices, the incremental predictive value of including the odometer is small.

We assume there is a secondary market where consumers can buy and sell used vehicles. The vast majority of trade in the secondary market in Denmark (about 90% according to bilbasen.dk, the largest used car website in Denmark) is intermediated by auto dealers rather than done as direct exchanges between individual consumers. Dealers refurbish/repair the used cars they buy and are legally required to guarantee the quality of the used cars they sell to consumers. We assume that as frequent traders in the used car market, dealers have a comparative advantage in inspecting and determining the physical condition of the used cars they buy from consumers. This lessens the problem of asymmetric information about the condition of a used car traded in Denmark, and thus we do not deem the Akerlof (1970) “lemons problem” to be a significant barrier to trade of used cars in Denmark.⁸ In addition, this tends to reduce the degree of idiosyncratic

⁷We decided not to adopt the modeling approach of Schiraldi (2011) of using a device similar to his “mean augmented net utility” since this is an endogenous stochastic process that is not firmly rooted in first principles in the sense that there is no way we can see to derive the form of this stochastic process from more primitive assumptions about consumer beliefs about the arrival of new technologies and models of vehicles to the market over time. It was not clear to us that making a somewhat arbitrary assumption about beliefs of “endogenous objects” (such as how consumers’ value functions change over time in response to new technological innovations in the vehicle market) result in more trustworthy forecasts than the simpler assumption of “stationary expectations” — i.e. the assumption that consumers do not expect any future technological innovations. Note that while we maintain an assumption of stationary expectations with respect to technology, we do allow non-stationarity due to the effects of macroeconomic shocks on the market, and we have chosen to focus on modeling how these factors affect consumer beliefs and trading since it is far more obvious from our analysis of the data how such shocks affect new car purchases and used car scrappage over time. We will attempt to investigate how our stationary expectations assumptions regarding technology can be relaxed in future work.

⁸Despite the wide attention to the “lemons problem” that Akerlof article raised, there is not clear empirical evidence that it is a serious problem in actual automobile markets. For example Bond (1982) found that pickup trucks that were “purchased used required no more maintenance than trucks of similar

variation in the unobserved quality of cars that consumers can buy, which helps to justify our assumption of a common price $P(\tau, a, p, m)$ for used cars of type τ and age a in Denmark.

Of course there will be idiosyncratic variation in the quality of cars that are sold to dealers, but we assume that by repairing/refurbishing used cars to be resold to other consumers, dealers help to homogenize the condition of used cars that are sold. We assume that dealers have a comparative advantage in estimating the costs of repairing and reconditioning a used car they buy from a consumer and this repair cost is borne by the consumer who sells their used car to a dealer. The idiosyncratic variability in this repair cost is captured by a random component in the transactions cost that a consumer incurs when they sell their used car to a dealer. This leads to the possibility that if a consumer has a used car that is in sufficiently poor condition, the amount they would receive from selling this car to a dealer net of the cost of repairing/refurbishing the vehicle could exceed the scrap price, $\underline{P}(\tau, p, m)$. In this case we assume that the car would be scrapped rather than sold to the dealer. We will describe this *scrappage decision* in further detail in Section 3.1, but we will show that it constitutes a *static subproblem* that a consumer faces whenever they decide to sell their existing car.

Our model allows for idiosyncratic factors such as the condition of the current car owned, and other unobserved factors to affect decisions about keeping a vehicle or trading it for another one. We account for these unobserved factors with random variables that capture the net effect of unobserved variables that pertain both to the consumer and to different cars they might consider buying, and other factors that may vary over time. For computational tractability of the model, we assume these unobserved factors have *IID* (over time) multivariate Type 3 generalized extreme value distributions that result in a “nested logit” structure for car choices. The nested logit specification allows for correlation in the unobserved transactions costs faced by a consumer who chooses to replace their current car. This enables the model to capture *endogenous scrappage decisions*, i.e., the consumer’s choice of whether to scrap their current car, or sell it in the used car market.

Besides the variables (τ, a) that index the type and age of car the consumer may currently own as well as all vehicles they can choose from at any given point in time, we introduce the key *macro variables* that we believe are relevant both for individual choices

age and lifetime mileage that had not been traded.” leading him to conclude that “This leads to a rejection that the market for pickup trucks is a market for lemons” (p. 839). However other studies, such as Engers, Hartmann and Stern (2008) conclude that “Our empirical results strongly suggest that there is a lemons effect because there is significant unobserved heterogeneity.” However we do not see sufficiently strong evidence for a lemons problem that would justify the added complexity in trying to explicitly account for it in our model. Certainly the most extreme prediction of asymmetric information does not hold: namely, the ‘lemons problem’ if it exists, is clearly not severe enough to kill off trading in secondhand markets for autos. Around the world, we see active secondhand markets for cars, which suggests to us that concerns about problems of asymmetric information and unobserved vehicle quality are of second order of importance relative to the primary benefit of the gains to trade that come from having an active secondary market.

and for the equilibrium of the market as a whole, (p, m) where p is the current price of fuel (we assume that diesel fuel is a fixed fraction of the price of gasoline, which is reasonably justified from the evidence presented in section 2) and m is an indicator of the “macro state” of the Danish economy. We model m as a binary variable where $m = 0$ indicates that the economy is in a recession period, and $m = 1$ indicates a non-recession period.

Consumer expectations of the price of a *typical car* of type and age (τ, a) when the economy is in state (p, m) are given by the function $P(\tau, a, p, m)$. These expectations affect individual agents’ choices of vehicles in an important way as we describe in more detail below. However we do not assume that agents have *perfect expectations* of vehicle prices in the sense that their beliefs about car prices coincide exactly with the actual future prices of new and used cars, that may change over time due to the effects of unforeseen macroeconomic or fuel price shocks. We define a notion of *temporary equilibrium* in Section 5 where *realized prices* of vehicles are computed that clear the market in the sense of setting expected excess demand to zero. We place no restrictions on the form of these realized or temporary equilibrium prices and allow them to vary freely over time to clear the market period by period. While consumers may not be able to *exactly predict* future prices of vehicles, they can form very good predictions of future prices using flexibly parameterized price functions $P(\tau, a, p, m)$ that depend on the type of each car τ , the age of the car a , and (p, m) the current fuel price and macro state. In fact, in our initial work, we find we are able to provide good approximations to future prices using expectation functions of the form $P(\tau, a)$ that do not even depend on the variables (p, m) at all. We will discuss the distinction between consumer expectations of prices and the prices that actually clear the market in more detail in section 5.

Since Denmark has no domestic car production, we make a “small open economy” assumption that there is an infinitely elastic supply of new cars in Denmark at fixed “world prices”. That is, we assume that the prices of all new cars are exogenously fixed at values $\bar{P}(\tau, p, m) \equiv P(\tau, 0, p, m)$ that represent auto producers’ profit maximizing pricing decisions under the assumption that demand for new cars from Denmark is a negligible component of their overall worldwide sales. Similarly, we assume there is an infinitely elastic demand for vehicles for their scrap value at an exogenously fixed price $\underline{P}(\tau, p, m) = P(\tau, \bar{a}(\tau), p, m)$, where $\bar{a}(\tau)$ is the oldest age of a vehicle of type τ in our model.⁹ We will present our model of the scrappage decision below, but it is helpful to point out that this model incorporates idiosyncratic shocks to the choice of scrapping and the choice of selling a used car in the secondary market. Sometimes it is possible that a consumer would choose to scrap a car (τ, a) even though the scrap price is lower than the prevailing secondary market price of that vehicle $P(\tau, a)$. The idiosyncratic shocks capture unobserved costs associated with scrapping versus selling, such as repair costs that an owner would have to undertake to put their car in “sellable condition.” Net of

⁹In the estimation, we will set $\bar{a}(\tau) = 24$ for $\tau = 1, 2$ corresponding to gasoline and diesel.

these repair costs the amount a household could receive from selling their car could be less than what they would receive from scrapping it, so these shocks can explain situations where households scrap cars for an amount that appears less than the amount they could receive from selling the car. While the temporary equilibrium prices we compute are generally monotonically decreasing from the exogenously fixed new car price $\bar{P}(\tau)$ to the exogenously specified scrap price $\underline{P}(\tau)$, due to the presence of idiosyncratic shocks and the effects of sufficient concentrations of older cars on market prices, it can sometimes be the case that there will be slight non-monotonicities in the prices we calculate, including a possibility that some used car prices of sufficiently old vehicles could be slightly below the scrap price.

The scrappage decision is important for helping our model to capture the age distribution the vehicle stock in Denmark, which has an upper tail that declines with age. If we made an alternative assumption that no car is scrapped until it reaches the oldest age $\bar{a}(\tau)$, then in the absence of macro shocks the model would imply a uniform stationary distribution of vehicle ages which is contrary to what we observe. Further, our model allows for *accidents* that result in a total loss of the vehicle. We model this as a probability $\alpha(\tau, a, x)$ that a car of type τ and age a owned by a consumer with characteristics x will experience an accident during the one year period of our model that is so severe that it is uneconomic to repair the vehicle. When such an accident occurs, the consumer is assumed to lose the vehicle, and thus the consumer enters the next period $t + 1$ as a household that does not own a car. In this way, accidents constitute an “involuntary” component of vehicle scrappage in our model that will help the model to fit the non-zero fraction of young cars being scrapped as shown in Figure 2.8.

We assume that households cannot purchase a car of the highest age $a = \bar{a}(\tau)$ in the used car market. Nevertheless, our model does allow consumers to *own* cars that are of this age. They can do this simply by keeping their current car until it reaches age $\bar{a}(\tau)$. Once the car reaches this age, we assume that it no longer important to keep track of its exact age. Thus to keep the age variable a bounded, we simply assume that all cars that are age $\bar{a}(\tau)$ and older are in the oldest age “equivalence class.” We do observe a slight upwards shift in the car age distribution for $a = 24$ in Figure 2.1. When a consumer holding one of these cars wishes to get rid of it, the only option is to scrap it and receive the scrap price $\underline{P}(\tau, p, m)$. The model can easily be extended to allow for trading in cars of the oldest age.

The prices of cars at all ages below the maximum, $a \in \{1, \dots, \bar{a} - 1\}$ are determined endogenously in the secondary market for vehicles in Denmark, i.e. as the prices that equate the supply and demand for vehicles of each type τ and each age $a \in \{1, \dots, \bar{a}\}$ when the macro state is (p, m) . These prices will generally exceed the scrap price $\underline{P}(\tau, p, m)$, and there will generally be supply of cars of these ages, but we do not make any restrictions on equilibrium prices yet. We return to this when we discuss the scrappage problem.

Let x denote a vector of *household-specific variables* the most important of which include a) age of household head, b) household income, and c) other observed and unobserved time-invariant factors. Age and income are treated as time-varying state variables. In the empirical application, we do not currently include any variables under c) but we include it in the exposition for completeness. An example of c) would be to allow for unobserved heterogeneity in households in their preferences for cars. Other types of observable heterogeneity can be allowed such as estimating separate models for urban and rural households. In future work we plan to explore various specifications that allow for richer types of unobserved and observed heterogeneity, but our approach is to start with the simplest specification that already allows for a good deal of heterogeneity via avenues a) and b) above.

We focus on households that own at most one car, which accounts for 87.9% of Danish households. We assume decisions are updated on an annual basis. At the start of each year a household makes a decision about whether to buy a new vehicle and/or sell their existing vehicle, but our model does not allow a household to purchase more than one vehicle in any period, and if a household has an existing vehicle, it cannot purchase another one unless it simultaneously sells the existing one. We assume that if a transaction decision is made, it occurs at the beginning of the period, i.e. if the customer trades for a new car, they will be able to use the new car immediately and for the rest of the one year time period. Let $d' = (\tau, a)$ denote the car choice decision, where $d' = (\emptyset, \emptyset)$ denotes the decision not to have any car.

It is important to realize that the last year's car choice constitutes part of the *current state* of the household at the start of time t when we assume it updates its decision about its automobile holdings. Thus we let $d = (\tau, a)$ denote the household's *car state* where we use the state $d = (\emptyset, \emptyset)$ to denote a household that does not currently own any car. If a household has no car, at the start of each (one year) period in the model we assume that the household makes a *car purchase decision* $d' = (\tau', a')$ where τ' is the type and a' is the age of car it chooses to buy. If the household chooses not to buy any car, this corresponds to the decision $d' = (\emptyset, \emptyset)$.

Now consider a household that has an existing car $d = (\tau, a) \neq (\emptyset, \emptyset)$. This household actually faces two simultaneous discrete decisions: 1) a *sell decision* and 2) a *buy decision*. In order to reflect the sell decision, we add a third component d_s to the vector $d' = (\tau', a', d_s)$ where the sell decision d_s takes three possible values, $d_s \in \{-1, 0, 1\}$ where $d_s = -1$ denotes a decision to sell the car for scrap, i.e., to receive $\underline{P}(\tau, p, m)$ for it, $d_s = 0$ denotes the decision not to sell the car (i.e. keep the current car $d = (\tau, a)$), and $d_s = 1$ denotes the decision to sell the car in the secondary market, i.e. to receive an expected price of $P(\tau, a, p, m)$. As we noted above, there are random shocks to utility (to be described in more detail shortly) that capture a number of factors that are observed by the household and unobserved by the econometrician, including any deviation between

the actual selling price of the existing vehicle and its expected value $P(\tau, a, p, m)$.

The sell decision provides the notational distinction we need to reflect the fact that a household who owns a car $d = (\tau, a)$ may either want to keep that car ($d_s = 0$), scrap that car ($d_s = -1$) or trade that car ($d_s = 1$) and purchase another car $d' = (\tau', a')$ of the *same type and age*. Notice that when a household chooses to keep the current car, $d_s = 0$, then the only possible value for the (τ', a') components of d' are $(\tau', a') = (\tau, a)$ where $d = (\tau, a)$ is the type and age of the currently owned vehicle. However if the household chooses to scrap or trade the current car, then they are free to choose any type of replacement vehicle, including a vehicle with the same type and age (τ, a) as their currently owned vehicle.

Thus, the choice set of a household that owns a car $d = (\tau, a) \neq (\emptyset, \emptyset)$ is

$$D(d) = \left\{ (\tau, a, 0), \{(\emptyset, \emptyset, d_s), d_s \in \{-1, 1\}\}, \{(\tau, a, d_s), \tau \in \{1, \dots, \bar{\tau}\}, a \in \{0, \dots, \bar{a} - 1\}, d_s \in \{-1, 1\}\} \right\} \quad (3.1)$$

corresponding to the options of 1) keeping the current car, or 2) selling or scrapping the current car and *not* buying another one to replace it (where $(\tau', a') = (\emptyset, \emptyset)$ denotes this choice), or 3) choosing to buy some other car $d' = (\tau', a')$.

The choice set for a household that does not have a car $d = (\emptyset, \emptyset)$ is

$$D(d) = \left\{ (\emptyset, \emptyset), \{(\tau, a), \tau = 1, \dots, \bar{\tau}, a = 0, \dots, \bar{a}\} \right\} \quad (3.2)$$

corresponding to the options of 1) continuing to not have any car, or 2) buying some car $d' = (\tau', a')$.

We use the notation $v_s(d', d, p, m, x)$ to denote the generic *indirect utility* that a household whose head is aged s and has observed characteristics x receives from the vehicle choice d' at the start of period t if it starts that period with a current car state d , and the fuel price is p and macro state is m . The reason we use the term “indirect utility” is that for households who choose to own a car $v_s(d', d, p, m, x)$ reflects the household’s expected utility from the use of that car during the coming year. We will introduce additional notation and a more detailed model of vehicle driving decisions in the next section, and show how we derive tractable functional forms for the indirect utility function from flexibly specified regression models of household driving decisions. For households who choose not to own a vehicle, $v_s(d', d, p, m, x)$ reflects the indirect utility from use of alternative non-car modes of transportation, such as bicycles, walking, and public transportation.

3.1 Household Dynamic Vehicle Choice Problem

We now describe the household’s dynamic optimization problem. The household lives for a finite time (with stochastic mortality of the household head, at which point we treat

the household as dissolved) and makes a sequence of car ownership decisions at annual intervals over the lifetime of the household. We assume the youngest age of any household head is $s = 20$ and the oldest possible age of a household head is $s = 85$. In addition, the households who own a car have an additional continuous decision on the number of kilometers to drive their car over the year, and the details of this decision will be described in the next section.

Just as in much of the relevant literature on vehicle choice, we do not solve a complete life-cycle optimization problem for the household. That is, we ignore the overall consumption-savings problem and do not carry household wealth as a state variable of the decision problem. Instead, we ignore borrowing constraints and assume that the household has enough cash on hand to buy a car when it wants to. Further we assume that the indirect utility function $v_s(d', d, p, m, x)$ is a “quasi-quasi-linear” function of the after tax household income y (a component of the vector of observed household characteristics x). That is, we assume that y enters $v_d(d', d, p, m, x)$ in an additively separable fashion but we allow y to enter into a coefficient $\theta(y, m)$ representing the “marginal utility of income” to reflect the effects of shifts in income on car usage, holding and purchase decisions. Low income households will have high marginal utilities of income, and thus a high “opportunity cost” for use of income for consumption other than automobiles. This will cause low income households to buy cheaper new cars, or used cars and perhaps to drive less compared to higher income households. Also expectations of future income and macro shocks will affect car purchases, and if a household expects to be in a period where their income will be persistently low (e.g. during a recession) they will expect their marginal utility of income to be high during this period and this could cause them to delay a purchase of a new car until better times when the economy is out of recession and their income is higher.

Though we do not model liquidity constraints explicitly, variations in the marginal utility of income can also indirectly reflect liquidity effects. A liquidity constrained household is likely to have a high marginal utility of income, and thus is less likely to purchase a new car. The cost of trading vehicles is captured by a *trading cost* function $T(d', d, p, m)$. This function captures the cost of buying a new car d' net of the proceeds received from selling the existing car d , plus a *transactions costs* and *taxes* associated with the purchase of a new car. More over, and perhaps more importantly, it covers non-monetary factors that result in higher holding times such as search costs, information frictions and psychological

attachment to an old car. The trading cost function is given by

$$T(d', d, p, m) = \begin{cases} 0 & \text{if } d' = (\tau, a, 0) \text{ or } d, d' = (\emptyset, \emptyset) \\ P(\tau', a', p, m) - P(\tau, a, p, m) + c_T(\tau', a', p, m) & \text{if } d' = (\tau', a', 1) \text{ and } d = (\tau, a) \\ P(\tau', a', p, m) - \underline{P}(\tau, p, m) + c_T(\tau', a', p, m) & \text{if } d' = (\tau', a', -1) \text{ and } d = (\tau, a) \\ -P(\tau, a, p, m) & \text{if } d' = (\emptyset, \emptyset, 1) \text{ and } d = (\tau, a) \\ -\underline{P}(\tau, p, m) & \text{if } d' = (\emptyset, \emptyset, -1) \text{ and } d = (\tau, a) \\ P(\tau', a', p, m) + c_T(\tau', a', p, m) & \text{if } d' = (\tau', a') \neq (\emptyset, \emptyset) \text{ and } d = (\emptyset, \emptyset) \end{cases} \quad (3.3)$$

Thus, there are no trading costs if the household keeps its current car, or does not have a car and chooses not to buy one. Trading costs are incurred when a household trades in their current car (τ, a) and buys a new one (τ', a') . The function $c_T(\tau', a', p, m)$ represents the *transactions cost* that a household incurs when *purchasing* a car (τ', a') . We assume that there are no transactions costs for *selling an existing car* (d, a) so that $P(\tau, a, p, m)$ represents the net amount a consumer would receive from an auto dealer if they were to sell their current car, whereas if they were to *buy* the same car (τ, a) from the dealer, the total price would be $P(\tau, a, p, m) + c_T(\tau, a, p, m)$. Thus, $c_T(\tau, a, p, m)$ can be regarded as a “bid-ask spread” that reflects both the repair and cleaning costs the dealer incurs to put a used car into “selling condition” as well as a profit margin for the dealer.

We assume that total transactions costs consist of a part that is proportional to the cost of the car plus an additive, fixed component

$$c_T(\tau', a', p, m) = P(\tau', a', p, m)b_1(\tau', a', p, m) + b_2(\tau', a', p, m) \quad (3.4)$$

where b_1 is the part of transactions costs that is proportional to the price of the car (τ', a') the consumer buys. In our initial estimation we use a simple specification where transaction costs are independent of the type and age of the vehicle, which amounts to the restriction $b_1(\tau', a', p, m) = b_1$ and $b_2(\tau', a', p, m) = b_2$.

We also assume that the new car registration tax is included in the (exogenously determined) prices of new cars, $P(\tau, 0, p, m)$. There is no tax on purchases of used cars in Denmark. Thus, a household that does not currently own any vehicle but decides to buy a car (τ, a) will incur a buy transactions cost that is incorporated in the gross (bid) price $P(\tau, a, p, m) + c_T(\tau, a, p, m)$, but a household who wants to sell a car (τ, a) does not incur any transaction costs, but instead receives the net of transaction cost (ask) price $P(\tau, a, p, m)$.

Note that the indirect utility function $v_s(d', d, p, m, x)$ will depend on the trading cost function $T(d', d, p, m)$ and the precise way it depends on T will be detailed in the next section. In the remainder of this section we present the Bellman recursion equations that define the household’s optimal dynamic vehicle holding and trading strategy. As

is the traditional practice in dynamic discrete choice models, we augment the set of state variables to allow for *IID* extreme value distributed *unobserved state variables* ϵ the enable us to derive convenient multinomial conditional choice probabilities for the events of whether a household keeps their car, buys a new car, etc. Thus, in addition to the indirect utility function v_s there is an additive error term $\epsilon(d')$ representing the impact of idiosyncratic unobserved factors that affect the consumer's choice, so the total current period utility becomes $v_s(d', d, p, m, x) + \epsilon(d')$. Let $\epsilon = \{\epsilon(d') | d' \in D(d)\}$ be the vector of these unobserved terms for all possible choices d' in the consumer's choice set $D(d)$. The choice set depends on the current car choice d so that only choices relevant to the consumer's current state are available.

Let $V_s(d, p, m, x, \epsilon)$ be the value function for a household of age s that owns a car $d = (\tau, a)$ (or no car if $d = (\emptyset, \emptyset)$) when the macro state is m , the fuel price is p , and the household has observed characteristics x and unobserved characteristic (state) (ϵ) . Our specification treats ϵ is a vector-valued *IID* extreme value process with a number of components equal to the number of elements in the household's state-dependent choice set $D(d)$ described in section 3.1 above. Note that the Type 3 extreme value distribution involves both *contemporaneous independence* between different components $\epsilon(d)$ and $\epsilon(d')$ for $d \neq d'$ as well as serial independence in the overall vector stochastic process $\{\epsilon_t\}$. These assumptions are mostly for computational convenience, though it is far easier to relax the assumption of contemporaneous independence, whereas relaxing the serial independence assumption is significantly harder and appears to be computationally infeasible given currently known econometric methods and computer technology.¹⁰

In future work we intend to relax the assumption of contemporaneous independence between the components $\epsilon_t(d)$ and $\epsilon_t(d')$ for $d \neq d'$ for any fixed time period t . A natural specification is the *generalized extreme value* (GEV) distribution for the vector ϵ_t that allows for contemporaneous correlation in the components of ϵ_t corresponding to a partition of the choice set of cars into car *classes* such as commonly used marketing categories such as “compact” “luxury” “sport utility vehicle” (SUV) and so forth. This partition of the car types τ can reflect unobserved characteristics of cars that are not easy to capture using traditional observable variables such as car weight or wheel base, that reflect characteristics of cars that consumers can observe that constitute patterns of “similarity” in these characteristics. The resulting model is the well known *nested logit* model that has been frequently used in discrete choice models of auto choice. In our initial model since we only allow for two different car types, diesel and gasoline, we feel that the types themselves capture the relevant unobserved characteristics of these two broad groups of vehicle types. The nested logit model is more revelant for future specifications

¹⁰Reich (2013) provides a promising new method for structural maximum likelihood estimation of dynamic discrete choice models with serial correlated unobservables, but so far the method has been only demonstrated for binary choice models and it is not clear that this method will continue to be tractable for high dimensional choice sets such as in the auto choice problem.

where we might add more type of vehicles in the model, such as different model or brands within the two broad categories “gas” and “diesel”.

The Bellman equation for V_s is given by

$$V_s(d, p, m, x, \varepsilon) = \max_{d' \in D(d)} [v_s(d', d, p, m, x) + \varepsilon(d') + \beta EV_s(d', d, p, m, x, \varepsilon)] \quad (3.5)$$

where $EV_s(d', d, p, m, x, \varepsilon)$ is the conditional expectation of $V_{s+1}(\tilde{d}, \tilde{p}, \tilde{m}, \tilde{x}, \tilde{\varepsilon})$ given the current state $(d, p, m, x, \varepsilon)$ and decision d' , where the tildes over the variables $(d, p, m, x, \varepsilon)$ entering V_{s+1} indicate the expectation is taken over the uncertain time $t + 1$ variables of these time-varying state variable. Since there are no wealth effects in our model, any decision that involves selling the current current car d (such as whether it should be sold on the secondary market or scrapped) does not affect the expected value of future utility conditional on the current choice d' , and thus EV_s depends only on d' , not d . Further, due to the fact that $\{\varepsilon_t\}$ is serially independent, EV_s depends on ε only via the current choice d' and thus EV_s does not depend directly on ε given d' , and we can write it as $EV_s(d', p, m, x)$. This implies that we can write the Bellman equation as

$$V_s(d, p, m, x, \varepsilon) = \max_{d' \in D(d)} [v_s(d', d, p, m, x) + \varepsilon(d') + \beta EV_s(d', p, m, x)]. \quad (3.6)$$

Let $V_s(d', d, p, m, x)$ denote the *choice-specific value function*

$$V_s(d', d, p, m, x) = v_s(d', d, p, m, x) + \beta EV_s(d', p, m, x). \quad (3.7)$$

Then following Rust (1985b) we can rewrite the Bellman equation (3.5) in terms of the choice-specific value functions (3.7) as

$$V_s(d, p, m, x, \varepsilon) = \max_{d' \in D(d)} [V_s(d', d, p, m, x) + \varepsilon(d')]. \quad (3.8)$$

Equation (3.8) simply says that the value function $V_s(d, m, p, x, \varepsilon)$ is the maximum over all alternatives $d' \in D(d)$ of the choice-specific value functions $V_s(d', d, p, m, x)$ accounting also for the effects of the *IID* extreme value shocks $\varepsilon(d')$ which represent transient, idiosyncratic unobserved components of utility that affect consumers' choices.

We now discuss an assumption on the distribution of the shocks $\varepsilon(d')$ that allows us to model endogenous scrappage decisions in a particularly simple manner. Note that for any alternative d' that involves trading an existing car for another one, the consumer has two possible options: 1) scrap the existing car, or 2) sell it in the secondary market. The assumptions we place on the utility function (quasi-linearity in the utility of driving from consumption of other goods) imply that the decision of how best to dispose of the existing vehicle is separable from the decision of which new car to buy. The consumer will sell the existing car on the secondary market if the net proceeds from doing this is

greater than the net proceeds the consumer would receive from scrapping it. Recall that for decisions involving trading the existing vehicle, the decision is represented by three components, $d' = (\tau', a', d_s)$ where $d_s = 1$ if the consumer sells the car in the secondary market, and $d_s = -1$ if the consumer chooses to scrap the car.

We assume a nested logit structure for the distribution of the unobservable components of cost/utility $\varepsilon(\tau', a', d_s)$ associated with each of the two possible decisions d_s for any decision $d' = (\tau', a', d_s)$ involving trading the current vehicle (i.e. where $d \neq (\emptyset, \emptyset)$ and $d_s \neq 0$). We assume that the unobservable components $(\varepsilon(\tau', a', -1), \varepsilon(\tau', a', 1))$ corresponding to the choice of whether to sell or scrap the currently held vehicle have a bivariate marginal distribution given by

$$F(\varepsilon(\tau', a', -1), \varepsilon(\tau', a', 1)) = \exp \left\{ - [\exp\{-\varepsilon(\tau', a', -1)/\lambda\} + \exp\{-\varepsilon(\tau', a', 1)/\lambda\}]^\lambda \right\} \quad (3.9)$$

where $\lambda \in [0, 1]$ is a parameter indexing the degree of correlation in $(\varepsilon(\tau', a', -1), \varepsilon(\tau', a', 1))$. These are independent Type 3 extreme value random variables when $\lambda = 1$ and they become increasingly correlated as $\lambda \rightarrow 0$. It is not hard to show that $\max(\varepsilon(\tau', a', -1), \varepsilon(\tau', a', 1))$ has a Type 3 extreme value distribution with a scale parameter $\lambda = 1$ which is the scaling parameter we assume (as a normalization) for the Type 3 extreme value distributions we assume for all of the distributions of all of the unobserved components of utility $\varepsilon(d')$ for the “upper level” decisions $d'(\tau', a')$ (i.e. all decisions except the decision about whether to scrap or sell the current car).

For each decision d' that involves trading the existing vehicle $d = (\tau, a)$, the consumer will prefer to sell the vehicle in the secondary market if

$$P(\tau, a, p, m) + \varepsilon(\tau', a', 1) \geq \underline{P}(\tau, p, m) + \varepsilon(\tau', a', -1). \quad (3.10)$$

Note that the unobserved components in the decision of whether to scrap the current vehicle or sell it in the secondary market depend on (τ', a') , which is the consumer's choice of *new car*. The third component, which takes the values $\{-1, 1\}$, corresponds to the decision to scrap or sell the *current car* $d = (\tau, a)$. We assume that the pairs $(\varepsilon(d', -1), \varepsilon(d', 1))$ and $(\varepsilon(d, -1), \varepsilon(d, 1))$ are independently distributed for any pair of upper level choices $d' = (\tau', a') \neq d = (\tau, a)$. This implies that conditional on making the “upper level” choice to trade the current car for a car $d' = (\tau', a')$ the consumer decides to sell their current car with probability

$$\Pr \{d_s = 1 | d, d', p, m, x\} = \frac{\exp \{P(\tau, a, p, m)/\lambda\}}{\exp \{P(\tau, a, p, m)/\lambda\} + \exp \{\underline{P}(\tau, a, p, m)/\lambda\}}. \quad (3.11)$$

The conditional probability of scrapping the car is just $1 - \Pr\{d_s = 1 | d, d', p, m, x\}$, and these choice probabilities can be calculated independently of the overall solution of the

dynamic programming problem given in equation (3.6) since the sell/scrap “subproblem” involve the simple choice of whether the net proceeds of selling the car in the secondary market exceed the scrap value $\underline{P}(\tau, p, m)$, accounting for unobservable components of the transactions costs associated with selling the car to a dealer, $\varepsilon(\tau, a, 1)$, and scrapping it, $\varepsilon(\tau, a, -1)$, respectively.

Letting $d' = (\tau', a')$, then we can write

$$\begin{aligned} \max [v_s((d', -1), d, p, m, x) + \varepsilon(d', -1), v_s((d', 1), d, p, m, x) + \varepsilon(d', 1)] = \\ \lambda \log (\exp\{v_s((d', -1), d, p, m, x)/\lambda\} + \exp\{v_s((d', 1), d, p, m, x)/\lambda\}) + \varepsilon(d') \end{aligned} \quad (3.12)$$

where $\varepsilon(d')$ is a Type 3 Extreme value random variable with scale parameter $\lambda = 1$ that is distributed independently of $\varepsilon(d)$ for $d' \neq d$. What we mean by the representation given in equation (3.12) is that the left and right hand sides have the same probability distribution, and the right hand side is equivalent to a “regression equation” that expresses the maximum utility of whether to scrap or sell the current car in terms of expected value (the log-sum term on the right hand side of (3.12)) and a single error term $\varepsilon(d')$ that has as Type 3 Extreme value distribution with scale parameter $\lambda = 1$.

Using equation (3.12) we can redefine the indirect utility function $v_s(d', d, p, m, x)$ as the *expected maximum* over the two decisions $d_s \in \{-1, 1\}$ for any upper level choice $d' = (\tau', a')$ that involves trading the current car $d = (\tau, a)$ for a new one. This allows us to abstract from the “lower level” scrap versus sell decision d_s and treat $d' = (\tau', a')$ as just the upper level decision of whether to keep the current car (or continue to have no car if $d = (\emptyset, \emptyset)$), or choose one of the available vehicles $d' = (\tau', a')$. For this “upper level choice problem” over $d' = (\tau', a')$ we redefine the indirect utility as

$$\begin{aligned} v_s(d', d, p, m, x) = \\ \lambda \log (\exp\{v_s((d', -1), d, p, m, x)/\lambda\} + \exp\{v_s((d', 1), d, p, m, x)/\lambda\}) + \varepsilon(d') \end{aligned} \quad (3.13)$$

Then with this redefinition/reduction, the Bellman equation (3.6) applies to the “upper level” choices $d' = (\tau', a')$. The probability that a consumer will choose to trade their existing car $d = (\tau, a)$ for another car $d' = (\tau', a')$ is then given by the standard multinomial logit model

$$P(d'|d, p, m, x) = \frac{\exp\{V_s(d', d, p, m, x)\}}{\sum_{d'' \in D(d)} \exp\{V_s(d'', d, p, m, x)\}}. \quad (3.14)$$

where $V_s(d', d, p, m, x)$ is the choice-specific value function (3.7) except that the indirect utility function $v_s(d', d, p, m, x)$ is given by the redefined log-sum value given in equation (3.13) above. Then given the choice to trade the current car $d = (\tau, a)$ for another car $d' = (\tau', a')$, the conditional probability that the consumer chooses to scrap the current car is given by equation (3.11) and the conditional probability that the consumer chooses

to sell the current car is just 1 minus this probability.

As usual in nested logit models, it is important to remember that the decisions of which car to trade for $d' = (\tau', a')$ and whether or not to scrap or sell the current car d_s are made *simultaneously* at each time period t even though the nested logit conditional choice probabilities create a strong temptation to view them as *sequential* decisions. The only sequential choices are those made at *different time periods*: all of the choices made at any given time period are made simultaneously at each time t .

Now we can further simplify the Bellman equation by writing it in terms of an “upper level log-sum”, where the choices are now $d' = (\tau', a')$ and we have subsumed the lower level choice of whether to scrap or sell the current car as described above. Let $f(d')$ denote the state of the chosen car d' *next period* $t + 1$. This is simply a reflection that if the consumer either chooses to keep their current car or trade for another one, that car $d' = (\tau', a')$ will be one year older next year (except at $a = \bar{a}$). Using primes to denote next period values of the time varying state variables, (p, m, x, ε) , we can use the properties of the independent Type 3 extreme value shocks $\varepsilon(d')$ to write the expectation of V_{s+1} with respect to ε' as follows:

$$\begin{aligned} \int_{\varepsilon'} V_{s+1}(f(d'), p', m', x', \varepsilon') q(d\varepsilon') &= \int_{\varepsilon'} \max_{d'' \in D(f(d'))} [V_{s+1}(d'', f(d'), p', m', x') + \varepsilon'(d'')] q(d\varepsilon') \\ &= \log \left(\sum_{d'' \in D(f(d'))} \exp \{V_{s+1}(d'', f(d'), p', m', x')\} \right) \\ &\equiv \varphi(f(d'), m', p', x'). \end{aligned} \quad (3.15)$$

Following Rust (1987) we can write the following recursion equation for the choice-specific value functions

$$\begin{aligned} V_s(d', d, m, p, x) &= v_s(d', d, m, p, x) + \\ &\quad \beta \sum_{m'} \int_{p'} \int_{x'} \varphi(f(d'), m', p, x') g(x'|x, m', p', m, p) h(p', m'|m, p) dx' dp' \end{aligned} \quad (3.16)$$

where $f(d')$ is given by

$$f(d') = \begin{cases} (\emptyset, \emptyset) & \text{if } d' = (\emptyset, \emptyset) \text{ or } d' = (\emptyset, \emptyset, d_s), d_s \in \{-1, 1\} \\ (\tau', \min[\bar{a}, a' + 1]) & \text{if } d' = (\tau', a') \text{ or } d' = (\tau', a', d_s) d_s \in \{-1, 0, 1\}. \end{cases} \quad (3.17)$$

As mentioned earlier, the continuation value (i.e. the expected discounted value of future utility, given by the expression multiplied by β in equation (3.17) above) depends only on d' and not on d . This is what equation (3.17) formalizes; for households who buy a new or used car, the continuation value is independent of whether the previous car was sold on the secondary market or scrapped. The expected utility only depends on the type and

age of the replacement car, $d' = (\tau', a')$. In addition to ignoring whether the previously held car was sold or scrapped, the f function ages the car that the household chose (or continued to hold, if $d_s = 0$) by one year, incrementing its age from a' at the start of period t to $a' + 1$ at the start of period $t + 1$.

As we noted previously, to keep the state space bounded we only track the age of vehicles of type τ up to some maximum age $\bar{a}(\tau)$, and we lump all cars of that type that are older than $\bar{a}(\tau)$ into an equivalence class of “very old cars”. Note that the Bellman equations do allow consumers to keep cars that are age $\bar{a}$ and older. This is what makes it possible for the model to predict “mass points” in the age distribution of cars in the cell representing very old cars that are age $\bar{a}$ and older. This mass point reflects consumers who decide to hold these cars rather than scrap them.

Comparing the two versions of the Bellman equations (3.8) and (3.17) we see that

$$EV_{s+1}(d', p, m, x) = \sum_{m'} \int_{p'} \int_{x'} \varphi(f(d'), m', p, x') g(x'|x, m', p', m, p) h(p', m'|m, p) dx' dp' \quad (3.18)$$

Note that the expected value function is only a function of the *chosen car* $d' = (\tau', a')$ but not the current car $d = (\tau, a)$ or the decision d_s of whether to scrap, or sell the current car, except in the case where the consumer chooses to keep the current car another year. Furthermore, the indirect utility functions we consider will have the property of additive-separability in the d' and d decision variables. This implies a substantial reduction in the dimensionality and we exploit this property to dramatically reduce the time required to solve the model by backward induction: instead of computing and storing the full set of choice-specific value functions $V_s(d', d, m, p, x)$ for all ages s and all values of the state variables, it is sufficient to compute and store only the expected values $EV_s(d', p, m, x)$. This computational reduction can be substantial even at fairly coarse discretization.

A small adjustment to the recursion equations is necessary to account for accidents that “total” a car (i.e., completely destroy it, beyond all chance of repair). In such cases, we assume that the car involved in the accident must be replaced at the start of the next period, but that insurance covers part of the cost of the car involved in the accident, but with some coinsurance rate ψ . So if the household chose a car $d = (\tau, a)$ at the start of the period, and this car was involved in an accident that totaled it, the household would receive a payment of $(1 - \psi)P(\tau, a, m, p)$. Then at the start of the next period the household would have no car $d = (\emptyset, \emptyset)$, but could use the insurance payment towards the purchase of a replacement vehicle of its choice. Let $\alpha_s(\tau, a, x)$ denote the probability that a household of age s with characteristics x that owns a car (τ, a) will have an accident that totals the car sometime during the period. Then the equation for the expected value

of future utility (3.18) above needs to be modified as follows

$$\begin{aligned}
EV_{s+1}(d', p, m, x) = & \quad (3.19) \\
(1 - \alpha_s(d', x)) \sum_{m'} \int_{p'} \int_{x'} & \varphi(f(d'), m', p, x') g(x'|x, m', p', m, p) h(p', m'|m, p) dx' dp' \\
+ \alpha_s(d', x) \sum_{m'} \int_{p'} \int_{x'} & \varphi_R(f(d'), m', p, x') g(x'|x, m', p', m, p) h(p', m'|m, p) dx' dp'.
\end{aligned}$$

where φ_R is the expected maximum utility over a restricted choice set $D_R(d)$ that requires the consumer to scrap their current car choice d that was involved in the accident:

$$D_R(d) = \{(\emptyset, \emptyset, -1), (\tau, a, -1), \tau \in \{1, \dots, \bar{\tau}\}, a \in \{0, \dots, \bar{a} - 1\}\} \quad (3.20)$$

corresponding to the options of 1) scrapping the current car and *not* buying another one to replace it (where $(\tau', a') = (\emptyset, \emptyset)$ denotes this choice), or 2) choosing to buy some other car $d' = (\tau', a')$, possibly including another car $d' = d = (\tau, a)$ of the same type and age as the current car that was involved in the accident. The definition of φ_R is similar to the definition of φ in equation (3.17) above except that the expectation is taken over the restricted set of alternatives $D_R(d)$ and the value functions entering into φ_R reflect a modified version of the trading cost function $T(d', d, p, m)$ given in equation (3.3) that reflects the insurance reimbursement net of coinsurance. Specifically, the modified trading cost function for a household who owns a car $d = (\tau, a)$ that is totalled in an accident, denoted $T_R(d', d, p, m)$, is given by

$$\begin{aligned}
T_R(d', d, p, m) = & \quad (3.21) \\
\begin{cases} -P(\tau, a, m, p)(1 - \psi) & \text{if } d' = (\emptyset, \emptyset) \\ [P(\tau', a', p, m) - P(\tau, a, p, m)(1 - \psi) + c_T(\tau', a', p, m)] & \text{if } d' = (\tau', a', -1) \text{ and } d = (\tau, a) \end{cases}
\end{aligned}$$

The Danish register data do not allow us to distinguish between “involuntary scrapping” caused by accidents that result in a total loss (unrepairable loss) to the vehicle, and “voluntary scrapping” where the customer makes a decision to scrap in connection with a trade, as discussed above.

3.2 Utility Specification

The approach here loosely follows that in Gillingham (2012) and Munk-Nielsen (2015). Let k be the total planned kilometers traveled by car over the coming year, and let $p^k(\tau, a, p, c^o)$ be the cost per kilometer traveled, defined as $p^k(\tau, a, p, c^o) \equiv \frac{p}{e(\tau, a)} + c^o$, where e denotes the fuel efficiency of the vehicle in kilometers per liter and c^o contains additional per-kilometer driving costs such as operating and maintenance costs but could also contain road tolls. Thus, the total costs of driving k kilometers is $p^k(\tau, a, p, c^o)k$.

Let $u(vkt, \tau, a, p, m)$ be the *conditional direct utility* a household expects from owning a vehicle of type τ and driving a planned k kilometers, given by

$$u(k, \tau, a, s, p, m) = \theta(y, m)[y - p^k(\tau, a, p, c^o)k - T] + \gamma(y, s, a, m)k + \phi k^2 - q(a) + \delta_n \mathbb{1}(a = 0) + \delta_\tau. \quad (3.22)$$

where $\theta(y, m)$ the marginal utility of money. We let $\theta(y, m)$ be a function of income, y , and the macro shock, m to capture the idea that households are less inclined to spend their money on cars during downturns and when income is low. The utility of driving is a 2nd-order polynomial in k , allowing for heterogeneity in the marginal utility of driving through $\gamma(y, s, a', m)$ and a concave relationship, with a diminishing marginal utility of driving, i.e. $\phi < 0$.¹¹ The coefficient δ_τ is a car-type fixed effect, δ_n is a coefficient on a new car dummy, and $q(a')$ is a 2nd-order polynomial in car age, capturing the rising maintenance costs with car age and ensuring scrappage. This helps to both fit the share of the no-car state as well as fitting the relative shares of the different car types in the data. Finally, recall $T(d'; d; p; m)$ is the trading cost function defined above.

We assume that driving does not affect the value of a car once we condition on it's age and type, such that the driving decision is separable from then car ownership decisions. The next period value function is therefore independent of k , such that the consumer's optimal *planned* driving is a fully static problem

$$k^* = \arg \max_k u(k, \tau, a, p, m).$$

The first-order condition for the optimal driving implies that

$$k^* = \frac{\theta(y, m)p^{km}(a, \tau) - \gamma(y, s, a, m)}{2\phi}.$$

We specify the heterogeneous parameter affecting the utility of driving as

$$\gamma(y, s, a, m) = \gamma_0 + \gamma_1 a + \gamma_2 a^2 + \gamma_3 s + \gamma_4 s^2 + \gamma_5 m + \gamma_6 y + \gamma_7 y^2,$$

Note that the optimal driving equation has no error term since we are considering the *planned* driving by the consumer. To take the driving equation to the data, we will think of the driving variable to be observed with measurement error. Finally, to capture that households are less inclined to spend their money on cars during downturns and when income is low, we allow dependence on the macro conditions, m , and for a diminishing marginal utility of household income, y ,

$$\theta(y, m) = \theta_0 + \theta_1 y + \theta_2 y^2 + \theta_3 m.$$

¹¹In the estimation, the function is monotone everywhere and predicts only strictly positive driving.

Inserting $\gamma(y, s, a, m)$ and $\theta(y, m)$, in the equation for the optimal k , we obtain the following linear equation

$$k^* = \frac{1}{2\phi}(\theta_0 + \theta_1 y + \theta_2 y^2 + \theta_3 m)p^k(a, \tau) - \frac{1}{2\phi}(\gamma_0 + \gamma_1 a + \gamma_2 a^2 + \gamma_3 s + \gamma_4 s^2 + \gamma_5 m + \gamma_6 y + \gamma_7 y^2) \quad (3.23)$$

$$= \kappa_0 + \kappa_1 a + \kappa_2 a^2 + \kappa_3 s + \kappa_4 s^2 + \kappa_5 m + \kappa_6 y + \kappa_7 y^2 + (\kappa_8 + \kappa_9 y + \kappa_{10} y^2 + \kappa_{11} m)p^{km}(a, \tau), \quad (3.24)$$

where $\kappa_j = -0.5\gamma_j/\phi$ for $j = 0, \dots, 7$ and $(\kappa_8, \kappa_9, \kappa_{10}, \kappa_{11}) = 0.5(\theta_0, \theta_1, \theta_2, \theta_3)/\phi$. The κ parameters are identified from this equation alone, implying that the structural parameters in $\theta(\cdot)$ and $\gamma(\cdot)$ are identified up to a normalization by ϕ . However, in the full model, all parameters are identified. We return to this in section 4.

3.3 Specification of the Transition Densities

In this section, we specify the stochastic structure of household income, y_{it} , fuel prices, p_t , and the macro state, m_t . We introduce the subscript i for households to emphasize that income varies across households and over time while the macro state and fuel prices are common to all households. We introduce the subscript t to more quickly clarify the time dimension of transition. We will also use this notation in the remainder of the paper.

For the income transition density, $g_s(y_{it}|y_{it-1}, p_t, m_t, p_{t-1}, m_{t-1})$, we assume that income follows a log-normal AR(1) process with an age profile,

$$\log y_{it} \sim \mathcal{N}(\mu_y, \sigma_y^2). \quad (3.25)$$

where μ_y is given by

$$\mu_y = \rho_1 \log y_{it-1} + \rho_2 s_{it} + \rho_3 s_{it}^2 + \rho_4 m_t + \rho_5 m_{t-1} + \rho_6 \mathbb{1}_{\{m_t=1 \wedge m_{t-1}=0\}} + \rho_7 \mathbb{1}_{\{m_t=0 \wedge m_{t-1}=1\}} \quad (3.26)$$

The coefficients ρ_6, ρ_7 allow for flexibility in the first year of a boom or a bust which will allow us to accommodate some of the sluggishness in the income processes that we observe in the data.

We next assume that log fuel prices follow a random walk. Anderson, Kellogg, Sallee and Curtin (2011) provide evidence using the Michigan Survey of Consumers that this is consistent with consumer expectations about the evolution of fuel prices. More precisely, we assume that

$$\log p_t \sim \mathcal{N}(\log p_{t-1}, \sigma_p^2). \quad (3.27)$$

Finally, we assume that the binary macro state, $m \in 0, 1$ follows a Markov process

with transition probabilities $\Pr(m_t = j | m_{t-1} = l)$ for $j, l \in 0, 1$.¹²

4 Estimation of the Model

In this section, we outline our strategy for estimating the proposed model using the Danish register data. We first explain some details before we get to the full likelihood function. After this, we outline a “two-stage” estimation strategy to simplify the estimation.

The detailed Danish register data enable us to identify the type of car and its age (τ, a) for every Danish household that owns a car, and the type and age (τ', a') of a replacement vehicle for any household that trades a vehicle. So we construct a panel dataset $\{d_{i,t}, x_{i,t}, k_{i,t}\}$ based on a large random sample from our data which contains all Danish households, $i = 1, \dots, N$ over time periods t where $d_{i,t}$ is the car holding/trading decision by household i during year t (including the scrappage decision), $k_{i,t}$ is the vehicle kilometers traveled for households owning a car, and $x_{i,t}$ are other household level variables we include in our dynamic programming model, the most important of which are the age of the household head $s_{i,t}$ and the household’s income $y_{i,t}$. We do not observe scrap prices in the data.¹³ Instead we assume that they are equal to the used car price at the maximum age as indicated by the scrappage rates we have from DAF (the Danish Car Dealer Association). That is, we assume that

$$\underline{P}(\tau, p, m) = \zeta_\tau^a P_0(\tau),$$

where $P_0(\tau)$ denotes the new car price we observe in the data (merchant suggested retail price, MSRP), and ζ_τ is the depreciation factor.¹⁴

Given the one year decision time intervals in our model, we fix a particular time at which decisions are assumed to take place for purposes of matching the model to the data. Specifically, we assume decisions are made on January 1 of each year. We also assume that income $y_{i,t}$ represents total income (after tax) in the present year and the age variable s_{it} is the age of the household head as of January 1.¹⁵ For the decision variable, we assume that a decision pertains to the *coming year* and so a household is recorded as trading its vehicle if we observe a sale between January 1st of the year and December 31st of that

¹²To extend this further, we could allow the transition probabilities for the macro indicator to be conditional on fuel prices, since fuel prices might be informative about the Danish macro state. The mechanism is that fuel prices proxy for oil prices which proxy for world demand.

¹³The scrappage subsidy paid out by the Danish Ministry of the Environment equals 1,500 DKK.

¹⁴On average in the data, ζ_τ is around 0.88. Unfortunately, we do not have variation over time but from correspondence with DAF, the depreciation rates are rarely updated over time. This is why we view them as unrealistic for the actual average transaction prices in a given year. However, the rates are only suggestive, so dealers will most likely be varying their margins around these, which are only available to dealers that are members of DAF and pay for the data.

¹⁵Alternatively, one could use income data for the previous year to make sure that car decisions are made conditional on income already earned.

year.

We solve the dynamic discrete choice model using backward induction. There is no bequest-motive in the final period but we solve the model with a maximum age of 85 even though we truncate our dataset, setting all household aged above 80 to be 80. For the continuous state variables, we use Chebychev-polynomials to approximate the expected value function, which is a very smooth object. The integrals in the transitions are solved using Gauss-Hermite quadrature, which we have found to be superior to simulation based integration given that they are basically univariate integrals.

In order to solve the model we need to evaluate it at a set of used-car prices. So far, when we have talked about the used car price system, $P(\tau, a, p, m)$, we have loosely discussed this as the consumers *belief* about used-car prices in the single-agent model. However, when we zoom out and look at the market as a whole, we may start to think about what prices will equilibrate the market in a given year t . We will therefore distinguish between the household-level beliefs about prices, $P(\tau, a, p, m)$, and the market-level prices, $P(\tau, a, t)$. Instead of diving directly into a joint estimation of both structural parameters and equilibrium prices, our strategy for taking the model to the data proceeds in two steps; in the first step, we will read in a set of initial used car price functions based on the suggested depreciation rates, ζ_τ . Our approach for solving for equilibrium prices is outlined in Section 5. Therefore, we start by solving using the price system

$$P(\tau, a, p, m) = \zeta_\tau^a P(\tau, 0), \quad \forall p, m,$$

where the used-car prices do not vary over the business cycle.

Another part of estimation involves estimating an income process for households to create the transition probability $g_s(y'|y, p', m', p, m)$ and the process $h(p', m'|m, p)$ for the macro shock and the gasoline prices described in section 3.3. We follow Rust (1985b) and estimate the transition densities separately in a first stage.

We also wish to include the data on driving and scrappage in our estimation. In order to leverage the driving information, we assume that driving in the data is contaminated by a Gaussian IID measurement error. Thus, the partial likelihood contribution from the driving equation is given by

$$f_k(k_{i,t}|x_{i,t}; \vartheta) = \Phi \left[\frac{k_{i,t} - k^*(x_{i,t}; \vartheta)}{\sigma_k} \right],$$

where Φ denotes the standard normal density. The partial likelihood contribution from the scrappage decision has already been derived and is given by the logit formula in equation (3.11). It will only apply for households choosing to scrap. This component will be key to identifying the scaling parameter in the scrappage decision, λ . In fact, for a given full set of car prices and scrap prices, we can estimate λ offline in a first stage.

These frequencies are shown in Figure B.7. However, including the scrappage probability in the full likelihood may prove once we start to change the other used-car prices since that will affect scrappage.¹⁶

Let ϑ contain all parameters jointly. The log-likelihood for the full sample is

$$L(\vartheta) = \sum_{i=1}^N \sum_{t \in \mathcal{T}_i} \log \left\{ \Pr(d_{i,t}|x_{i,t}; \vartheta) f_k(k|x_{i,t}; \vartheta) [\Pr(d_{i,t,s}|x_{i,t})]^{\mathbb{1}\{d_{i,t,s} \neq 0\}} \right\}, \quad (4.1)$$

where the conditional choice probability for the car decision, $d_{i,t}$, is given by (3.14) and $\mathcal{T}_i$ denotes the years where we observe household i . Recall that $d_{i,t,s}$ denotes the decision whether to sell the in the secondary market ($d_{i,t,s} = 1$, get rid of the vehicle ($d_{i,t,s} = 0$) or scrapping the car ($d_{i,t,s} = -1$) for household i and time t . Hence, $\mathbb{1}\{d_{i,t,s} \neq 0\}$ is an indicator for the decision involving the consumer getting rid of a vehicle where the household must make a decision about whether to scrap or sell at the used car market.

We then maximize the log-likelihood using analytical gradients and a range of common optimization algorithms, including BHHH and several quasi-Newton algorithms. We have also used the gradient-free optimizer, Nelder-Mead, which has proven helpful whenever the gradient-based methods got “stuck” in the sense that they could not improve the likelihood along the gradient.

To simplify estimation, we start out with a “two-stage approach”; in the first stage, we estimate the κ parameters in the driving equation (3.23). Let $k_{i,t}^*(\kappa)$ denote the predicted driving for household i at time t . We can now solve the model, inserting this predicted driving from the first stage wherever we need the driving and keeping the κ -parameters fixed while searching over the remaining parameters. Formally, we solve the model replacing the flow utility with:

$$\begin{aligned} u[k^*(\kappa), \tau, a, s, p, m] &= \theta(y, m)[y - p^k(\tau, a, p, c^o)k^*(\kappa) - T] \\ &\quad + \gamma(y, s, a, m)k^*(\kappa) + \phi[k^*(\kappa)]^2 - q(a) + \delta_n \mathbb{1}(a = 0) + \delta_\tau. \end{aligned} \quad (4.2)$$

Then we use the following 2nd stage likelihood function,

$$L^{\text{2step}}(\vartheta) = \sum_{i=1}^N \sum_{t \in \mathcal{T}_i} \log \left\{ \Pr(d_{i,t}|x_{i,t}; \vartheta, k = k^*(\kappa)) [\Pr(d_{i,t,s}|x_{i,t}, k = k^*(\kappa))]^{\mathbb{1}\{d_{i,t,s} \neq 0\}} \right\}, \quad (4.3)$$

where the conditioning on $k = k^*(\kappa)$ is to indicate that the model should be solved using the flow utility given in (4.2). Note that we are still searching over the same parameters as when we use the full likelihood from equation (4.1); the γ - and θ -parameters are identified by the discrete choice alone. In a sense, this two-stage approach is similar to thinking of the predicted driving, $k_{i,t}^*(\kappa)$ as a characteristics of the chosen car, $d_{i,t}'$. The two-stage

¹⁶In the empirical application, we have kept λ fixed at 0.9.

approach breaks the otherwise very strict cross-equation restriction that the consumer should care equally much about money spent on buying and selling the car and money spent on driving the car. However, we can check if the estimated γ - and θ -parameters divided by ϕ correspond in magnitude to the respective κ -parameters as a test of the cross-equation restrictions.

5 Solving for Equilibrium Prices

In this section, we present our strategy for modeling used car prices. We first describe the consumer expectations, which we simplify here. We then outline first how we solve for equilibrium prices in-sample and then out-of-sample (for simulating forward in time). In Section 5.3 we will discuss an alternative approach using different assumptions.

5.1 Solving for Equilibrium Prices

We follow a literature stretching back to Rust (1985c) that estimates equilibria in both primary and secondary markets using an equilibrium price function. A key feature of our approach is that we relax the stationarity assumption. Specifically, we allow for the effects of macroeconomic shocks and changes in fuel prices, which was shown to play an important role in the U.S. vehicle fleet in Adda and Cooper (2000b) and our data suggests in the case in Denmark as well. We combine this with equilibrium price adjustments, which was shown to be important by Gavazza, Lizzeri and Roketskiy (2014).

To do this, we will allow used car prices to vary freely over time, but assume that consumers have *stationary expectations* regarding the prices of cars in the sense that consumers expect that the used car price system they observe today will be the same tomorrow. While this for example neglects how equilibrium vehicle prices in the future could depend on future macro conditions, gasoline prices and the distribution of vehicles in the used car market, we think this is a reasonable approximation that will, in principle, allow us to precisely equate supply and demand for all car types and ages in every single year. We will discuss the alternative approach of solving for a price system as a function of (p, m) , where consumer expectations are non-stationary, but do not solve exactly for equilibrium in a given year.

To explain our strategy for finding equilibrium, we will first go through an approach based on simulated *realized excess demand* to fix the intuition. We then outline our preferred approach, based on what we call the *expected excess demand*. The first strategy for finding equilibrium prices $P(\tau, a, t)$ proceeds as follows. Just as in previous literature, we search for a vector of prices $P(\tau, a, t)$ that will set excess demand to zero for vehicles for all vehicle types and ages and in each time period t . These excess demand functions arise as aggregations of the individuals' actions; since households are simultaneously the

supply and demand side of the used car market, we can find excess demand for car (τ, a) by taking the sum of individuals purchasing the car and subtracting the sum of individuals selling it. We will not work with this “realized excess demand”, however, because it will be an unwieldy criterion function to work with numerically since it will be locally flat. The reason for this is that, holding uniform draws fixed, a small change in a given parameter value might not induce any consumer to change their discrete choice and when it does, the change will be discontinuous for the same reason. For this reason, we will instead work with the “expected excess demand”, $ED(\tau, a, S, P)$, where S denotes the matrix containing all cars and households and P is the price system. We define ED as

$$ED(\tau, a, S, P) = \sum_{i=1}^N \Pr[d' = (\tau, a) | d_{i,t}, x_{i,t}; P] - \sum_{i=1}^N \{1 - \Pr[d'_s = 0 | d_{i,t}, x_{i,t}; P]\} \mathbb{1}(d_{i,t} = (\tau, a)), \quad (5.1)$$

for $a \in \{1, \dots, \bar{a} - 1\}$ and all τ .

Note that $ED(\tau, 0, S, P) = ED(\tau, \bar{a}, S, P) = 0$ by assumptions discussed earlier.¹⁷ The first term in equation (5.1) is the expected demand for (τ, a) -cars. The second term is the expected supply of these cars, given by the sum of probabilities not to keep the car for the households that own a car of type (τ, a) . This is an important distinction between expected demand and supply; all households contribute to the demand for all cars but they only contribute to the expected supply of a car if they own that car. Since the choice probabilities are continuous in prices, ED will be continuous.

Our algorithm therefore proceeds on a year-by-year basis. Consider year t in our sample; let P_t denote a vector of prices. We then calculate $ED(\tau, a, S_t, P_t)$ for each τ and for $a \in \{1, \dots, \bar{a} - 1\}$ and stack them in the vector $ED(S_t, P_t)$. If we have one price for each car category, the price system $P(\tau, a, t)$ is fully non-parametric and we can solve the non-linear system of equations,

$$ED(S_t, P_t) = 0. \quad (5.2)$$

If we have fewer prices than there are car classes, then we will not generally be able to solve the system and we can instead choose the prices that “do best” in the sense of minimizing $\|ED(S_t, P_t)\|$, where we might for example use the regular L_2 norm. This paper makes no claim as to uniqueness, but we have successfully solved (5.2) using real data, so existence is proven constructively for our situation. The conditions under which equilibrium prices exist are left for future work.

¹⁷As discussed earlier, Denmark is a small country without domestic car production, so $ED(\tau, 0, S, P) = 0$ is the “small open economy” assumption that Danish demand does not move the world car prices. The assumption that households cannot trade cars of the oldest age implies that these cars will always be scrapped at the exogenous scrap price, which means that $ED(\tau, \bar{a}, S, P) = 0$.

When working with the non-parametric specification in (5.2), we found that it was essential to use analytic gradients and an advanced root-finding algorithm. This is because cars (τ, a) and $(\tau, a + 1)$ are very close substitutes, so the cross-derivatives are extremely important to account for. Once that was done, the algorithm converged nicely with excess demands on the order of 10^{-10} .

5.2 Simulating Forward in Time

We simulate forward in time by recursively solving for equilibrium and simulating one step ahead. Let $P^*(S_t)$ denote the equilibrium prices that set excess demand to zero in (5.1) given the car distribution S_t . Let $\Gamma(S_{t+1}|S_t, P_t)$ denote the density of the next-period state variables. It is a sequential density in the sense that to draw from it, we first draw the discrete choice from the conditional choice probability (3.14). Next, we can take draws of the remaining state variables from their respective transition densities. Finally, if an accident occurs, the household's car is destroyed and their simulated car state becomes the no-car state. Note that the fuel price and the macro state are synchronized across households; since these are exogenous, we can draw them without regard to the individuals' car choices.¹⁸

The recursive simulation proceeds as follows in the r th step: Given S_r , find the equilibrium price vector $P^*(S_r)$. Then simulate next-period states from $\Gamma[\cdot|S_r, P^*(S_r)]$. Proceed until the desired number of simulated periods has been reached.

By simulating this way, we ensure that cars do not appear out of nowhere and do not disappear, except for scrappage or accidents. Without equilibrium prices adjusting, the number of cars of type (τ, a) may be higher or lower than $(\tau, a - 1)$ in the previous year. Note, however, that we do not impose this; the equilibrium prices guarantee that it will be the result. The exception is of course simulating noise in drawing from Γ .

Since we want to simulate data from the model forward in time, we need to think about households reaching the maximum age. We handle this by letting a new household enter the sample at the youngest age whenever a household reaches the oldest age and dies. This new household will be born with the dying household's car endowment to make sure that cars do not disappear out of the economy and cause a mismatch of supply and demand over time. This will ensure that the population and the car stock remains representative.

¹⁸The macro state could in principle be allowed to depend on the car purchases since new car sales are well-known to precede upswings. However, we cannot allow households to form expectations about this since that would require knowledge about not only their own actions but the actions of everyone else, which requires knowledge about the full cross section, S_t .

5.3 Non-Stationary Expectations

The approach outlined in the previous section can be expanded to relax the assumption of stationary expectations. However, the problem is that the equilibrium prices, defined as setting excess demand to zero, will in general be a function of the full age distribution in addition to the other state variables. To solve a model where households form expectations based on this would ordinarily require carrying the entire age distribution of vehicles as part of our vector of state variables. This may be possible in a dynamic programming model with a very limited number of types of vehicles, but quickly becomes infeasible due to the curse of dimensionality. An alternative and more pragmatic approach is to follow Krusell and Smith (1998) and assume that equilibrium vehicle prices can be well predicted using a much smaller-dimensional set of “sufficient statistics,” for example the price of gasoline and the macro state (p, m) .¹⁹

To implement this approach, we would choose some parameterization, $P(\tau, a, p, m) = P(\tau, a, p, m; \vartheta^P)$. For any trial value of the parameters indexing ϑ^P , we can calculate the excess demand for all our sample years. From this starting point, we can search for the value of ϑ^P that yields the smallest excess demand across years. Consumers would have “correct” expectations about future used-car prices but in any given year, the market might be out of equilibrium. This would imply that the model would do worse in terms of matching the waves in the car stock that we observe in Figure 2.1. For this reason, we choose to maintain the assumption of stationary expectations and leave non-stationarity for future work.

6 Results

This section presents the results from estimating the model. We start with a discussion of the practical implementation and the choices and simplifying assumptions we have made. We then present the results from the first-stage estimation of the driving equation and then the full set of structural parameters. We present a range of results illustrating the fit of the model and finally show a simulation of the car stock forward in time. After this, we turn to solving for equilibrium prices in all the sample years and analyze the in-sample fit under equilibrium prices. We then present a forward simulation with equilibrium prices and compare the waves in the car stock to those generated by the non-equilibrium model. Finally, conduct a counter-factual policy experiment, comparing the predicted response with and without equilibrium prices.

¹⁹Krusell and Smith also include the average value of the individual specific savings as a sufficient statistic. We could similarly add the average vehicle age to the households’ state variables but we choose the simpler route and see how far we can get in replicating the fleet dynamics by using only gasoline prices and the macro state.

6.1 Implementation

The results presented below are carried out for a 1% random subsample of the households in our data where nothing else is noted.²⁰ This is done to ease the computational burden of estimating the model and solving for equilibria where the primary constraint is the number of observations.

For the fuel price process, we assume them to follow a random walk according to equation (3.27) and estimate the variance on the innovations, σ_p , as the standard deviation of the change in real log fuel prices from 1972–2013 to be $\hat{\sigma}_p = 0.0693$. We have estimated different versions of the AR(1) income process and the estimated coefficients are shown in Table A.1. While we can reproduce the life-cycle path in income very clearly, we found the surprising result that the coefficient on the macro dummy (ρ_4) got a negative sign. In Appendix A, we furthermore show estimates from an AR process for labor income only, which also produces a negative macro dummy (Tables A.2). We believe that the problem is related to the very mild recession in 2001–2003, which actually saw higher growth rates than in most of the years of the boom in the 1990s (Table A.2). To avoid the problems that these counterintuitive transition rates might introduce, we have chosen to estimate a model where households expect that their income will never change (i.e. $\rho_1 = 1$, $\sigma_y = 0$ and $\rho_j = 0$ for $j > 1$). This will shut down the life-cycle perspectives that there might otherwise be in the model with regard to for example young households expecting to earn more in the future.²¹ However, we still utilize the cross-sectional distribution in income, which will generate gains from trade as richer household buy newer cars and hand them down to households with lower incomes.

To solve and estimate the model, we must make choices on discretization. We choose to have 25 age categories, making the maximum car age $\bar{a} = 24$. This is because by age 24, we have seen the larger part of the waves in Figure 2.1 die out due to scrappage. For the household age, we solve the model with a maximum household age of 85. When a household in the model becomes 85 years old, it dies and there is no bequest motive in the model, so households close to this age may choose to sell their cars and eat all they have since there is no continuation value to owning a car. To avoid this behavior, we top code all households aged 80 and above as being of age 80.

Regarding prices, we take the MSRPs and take the unweighted average within each of the two car types in each year to construct the new car prices. We do the same with all car characteristics as well as the DAF suggested depreciation rates, ζ_τ . We fix the scrap price so that it equals the price of a 24 year old car, i.e. $\underline{P}(\tau, p, m) = \zeta_\tau^{\bar{a}} P(\tau, 0, p, m)$.

²⁰The subsampling is over households, so we select all observations for a given household if it is selected. This is to ensure that we have a panel. Since we do not exploit the explicit matching between the buyers and sellers in the market, the random subsampling will not affect our results beyond precision.

²¹Recall that the most common ownership length is 5 years (Figure 2.7). If households had held on to the same car from new until scrappage, assuming away the life-cycle aspects of income growth would have been a considerably worse assumption.

Table 6.1: First-stage Driving Estimates

Main effects			
	Variable	Estimate	std.err.
κ_0	Const	35.74***	0.9359
κ_1	Car age	-0.3444***	0.0197
κ_2	Car age sq.	0.002467**	0.0009
κ_3	m	-19.63***	0.7554
κ_4	Inc	-0.0004118***	0.0013
κ_5	Inc sq.	3.826e-07***	0.0000
κ_6	HH age	0.3178***	0.0147
κ_7	HH age sq.	-0.004956***	0.0001
Price interactions			
	Variable	Estimate	std.err.
κ_8	PPK	-26.46***	1.2110
κ_9	PPK*inc	0.0007932***	0.0019
κ_{10}	PPK*inc sq.	-5.81e-07***	0.0000
κ_{11}	PPK*m	27.34***	1.0711
	Avg. PPK-elasticity		-0.6652
	R^2		0.1030
	N		111231

Data for all years 1996–2009 is used.

We choose to use the “two stage” estimation procedure outlined in Section 4: we start by estimating the κ -parameters in the driving equation (3.23) in a first stage. Then we use the κ -parameters to predict driving and use that in the flow utility as shown in equation (4.2), and find the structural parameters by maximizing (4.3).

6.2 First-Stage Results

To make matters simpler, we estimate the parameters from the driving equation in a first step and keep those fixed in the estimation of the remaining structural parameters. This greatly limits the number of parameters to be estimated. We estimate these parameters on the 1% subsample, where we pool all the driving observations from households who have a car (111,231 households). For the estimation, we have used individual-level variation in fuel prices, matching the daily fuel prices to the driving period at the daily level. This considerably increases the variation and we found that only relying on annual fuel prices gave insufficient identifying power to adequately identify the price parameter and, in particular, the interaction effects. The results are shown in Table 6.1.

The driving results imply an elasticity of the Price Per Kilometer (PPK) of -0.67 . This elasticity is not out of bounds from what has been found elsewhere but perhaps a bit

on the high side, compared to the findings of Munk-Nielsen (2015). However, if we were to include a more flexible functional form, accounting for more observable heterogeneity, this elasticity does go down. Since our model limits us by the state variables, we go with the results in 6.1. In Table B.3, we show regressions corresponding to the first stage specification in Table 6.1, but adding the heterogeneity sequentially and on the full dataset. The results differ somewhat for the full sample, resulting in higher PPK elasticities. We discuss this more in Appendix B.4 but choose, for consistency, to use the κ -estimates coming from the same sample that we use for the estimating the full structural model.

While a simultaneous estimation of the driving parameters and the remaining structural parameters is superior to this two-stage approach, it is not completely unrealistic. This approach breaks the tight cross-equational restriction imposed in most discrete-continuous models, yielding more flexibility for fitting the data but at the cost of internal model consistency.

6.3 Structural Estimates

The estimates shown below are based on the 1% subsample and only the cross sections for the years $t = 97, 99, 01, 03, 05, 06$ are used; we have used only a subset of the periods to reduce the computational burden required for estimation and we found that adding more years did not substantially change our estimates. We only include the intercept in the utility of driving (γ_0) and fix $\gamma_j = 0$ for $j > 0$, since heterogeneity in the realized driving is already accommodated by the reduced-form driving parameters (the κ s).²² Standard errors are estimated based on the inverse of the Hessian at the estimated parameters.

In estimating the model, we found that the transaction cost parameter deserved extra attention. In the literature, this has often been estimated to have relatively high values (e.g. Schiraldi, 2011) but we have found much higher estimates than what we have seen in the literature. Therefore, we estimate two versions of the main specification; one where we estimate the transaction costs and another specification where we keep it fixed at an a priori sensible level. For the latter, we choose a fixed cost of 10,000 DKK and a proportional cost of 20% of the traded car's value. If anything, we feel that these are somewhat high. However, we found that by increasing transaction costs and lowering the utility of money (θ_0), the likelihood did in fact increase. Our preferred estimates are from the model where we estimate fixed transaction costs and fix the proportional transaction costs to zero because it provides a superior fit of the data.

Our preferred estimates are shown in Table E.1. Most notably, the fixed transaction

²²Including the γ -heterogeneity parameters seems futile since a more fruitful long-term goal would be to jointly estimate the driving parameters and the rest of the structural discrete choice parameters. Then, the κ s would not be used and the driving equation would help give identification power to the γ -heterogeneity terms.

cost parameter (b_2) is estimated to be 233.33. Since money is measured in 1,000 2005-DKK, this corresponds to 233,330 DKK or the equivalent of two-thirds of a new car's price. We fix the proportional transaction cost (b_1) to zero.²³ We find this estimate too high to be reasonable but acknowledge that given the rest of the model, households are behaving *as if* transactions costs were so high. We note that transactions costs proxy for any source of frictions that might exist in the market, including psychological costs, asymmetrical information costs (lemons premia), etc., so they may of course be higher than the purely monetary cost of buying a car. Nevertheless, the high transaction cost parameter can also be seen as a sign of misspecification somewhere in the model. One possible explanation is related to curvature in income; it might be that the utility of money relevant for making driving decisions is much lower than the utility of money that applies when making car purchase decisions.²⁴ We think that extensions of the model in these directions might prove valuable for getting more reasonable transaction cost estimates.

The remaining parameter estimates are sensible; $\hat{\gamma}_0 > 0$ so that households tend to prefer the types of cars that are also associated with high driving (coming from $k^*(\kappa)$). We also find that the utility of money is positive, $\hat{\theta}_0 > 0$, and that the interaction with the macro state is negative, $\hat{\theta}_3 < 0$; this indicates that in bad macro times, money becomes more dear to households. This effect can be thought of as proxying for the changing shadow value of money as risk increases or as credit becomes tighter.

Next, we turn to the fit of the model for these parameter values. Figure 6.1 shows the model fit in terms of the choice probabilities (observed and predicted), here shown for the 2002 cross-section.²⁵ We note that in particular the age profile in demand tracks the observed transaction frequencies quite closely. There are, however, deviations; for car ages 3 and 4 and for 14–18, we under-predict. These are examples where the fixed depreciation rates appear to be unrealistic. Nevertheless, the model appears to get the overall functional form of the keep probability over the car age right on average. The figure also shows that we are under-predicting used-car purchases for car-owning households and over-predicting the purge decision. Similarly, we under-predict the number of no-car households staying in the no-car state. This might be because a lot of the heterogeneity in the keep decision appears to be related to life-cycle patterns (cf. Figures 2.4 and 2.5). We

²³We have tried estimating both the fixed and proportional transaction costs, b_1, b_2 , but found that the likelihood function was maximized for *negative* proportional transaction costs ($b_1 < 0$). This is theoretically impossible, so we chose to just fix $b_1 = 0$ and estimate b_2 .


²⁴Specifically, we have in mind a model where households are liquidity constrained. Then the choice to purchase a new car might push the household down into a region where the utility of money is much higher. Fuel costs, on the other hand, are not really paid up front such as it is indicated by the flow utility, but are paid weekly. In a quasi-linear model, this makes no difference, but in a model with curvature in the utility of money it can make a big difference. In fact, the macro term shifting up and down the utility of money is already something we think of as an approximation to the shadow value of money changing as the household's risk of becoming unemployed changes.

²⁵We have chosen to consider model fit for a single year because pooling the years is complicated by the fact that the choice set changes over time (and in principle, policy parameters might change over time although we have not pursued this).

Table 6.2: Structural Estimates — Estimated Transaction Costs

	Variable	Estimate	Std.err.
Model setup			
	Min. Hh. age	20	
	Max. Hh. age	85	
	# of car ages	25	
	# of car types	2	
	Clunkers in choicest	1	
β	Discount factor	0.95	
ρ	Inc. AR(1) term	1	
σ_y	Inc. s.d.	0	
ρ_p	Fuel price AR(1) term	1	
σ_y	Fuel price s.d.	0.0699	
$\text{Pr}(0 0)$	Macro transition	0.75	
$\text{Pr}(1 1)$	Macro transition	0.8	
	Accident prob.	0.0004	
λ	Logit error var.	1	
λ^{scrap}	Scrappage error var.	0.9	
Monetary Utility			
θ_0	Intercept	0.032508***	0.0001248
θ_1	Inc.	-2.664e-05***	2.038e-07
θ_2	Inc. sq.	2.7409e-08***	2.063e-10
θ_3	Macro	-0.0011238***	2.307e-05
Driving Utility			
γ_0	Intercept	0.046713***	0.0004668
γ_1	Car age	0	
γ_2	Car age sq.	0	
γ_3	Hh. age	0	
γ_4	Hh. age	0	
γ_5	Macro	0	
γ_6	Macro	0	
γ_7	Macro	0	
ϕ	Squared VKT	0	
Car Utility			
$q(a)$	Car age, linear	0.073057***	0.000819
$q(a)$	Car age, squared	5.7638e-05	3.249e-05
δ_1	Car type dummy	0.64764***	0.01111
δ_2	Car type dummy	0.14377***	0.01184
Transaction costs			
	Fixed cost	223.33***	0.8837
	Proportional cost	0	
	N	169,733	


Figure 6.1: Model Fit: Conditional Choice Probabilities (CCPs)



conjecture that the fit would be improved if the heterogeneity parameters in the driving utility (γ_j , $j > 0$) were estimated. Alternatively, it might be that the fact that households expect their incomes to be constant is causing this; when young households believe that it will increase shortly, it will make sense for them to postpone purchasing a car to a period where the utility of money is lower because their income is higher.

To explore the fit of the model by state variables, Figure 6.2 shows the predicted choice probabilities by four of the state variables for the 2002 data. To do this, we must choose one particular discrete choice, so we choose to focus on the “keep” decision, since it captures much of the dynamics in the model. The top left panel shows the fit for income. First, income is divided into bins according to quantiles of the income distribution. Within each of these bins, the figure shows the average probability of choosing keep according to the model predictions (evaluated at the state variables in the 2002 data) and observed in the data. The figure does not condition on car ownership so “keep” may mean to keep a car or to remain in the no-car state (which probably explains some of the heterogeneity over the income distribution). The figure shows that the model predicts a strong U-shape over the income distribution but that the data has a much flatter distribution. This indicates that while high-income households in the data do care less about money, the predictions of the model have an even stronger relationship (working through θ_1, θ_2). In the top right panel, the fit over household age is shown. For each age, the average probability of keeping is shown for the model prediction and the data. Here, the reverse is the case; the data shows a much stronger U-shape than the model prediction. This is probably because household age affects neither $\theta(\cdot)$ nor $\gamma(\cdot)$ but only works through the

Figure 6.2: Model Fit by State Variables: The Keep-decision




first-stage predicted driving (κ_6, κ_7). In the lower left panel, the fit is evaluated by car age. This panel is only based on car-owning households and for each car age group, the average predicted and observed probability of keeping is matched up. The figure shows that the model captures the keep probability over car age very well. Finally, in the lower right graph, we show the fuel price. Since there is only one fuel price per year, this just shows the average probability. This serves as a reminder that it is hard to compare the model fit in terms of the fuel price because there is just one fuel price per year. The panel also indicates that we are on average under-predicting the keep decision.²⁶

Finally, we present a simulation forward in time from the model to illustrate how the car age distribution of cars develops for these estimated parameters. To do this, we take the dataset in 2002 as the baseline. Then we iteratively compute choice probabilities and simulate choices and subsequently simulate the next-period-states, i.e. drawing from the density, $\Gamma(\cdot|S_t, P_t)$ from Section 5.2, using the DAF used-car prices for P_t . We choose to keep car and fuel prices fixed at the 2002 values in the simulation but simulate the macro process, which is synchronized across all agents. The resulting simulated car age distribution is shown in Figure 6.3 and the simulated macro process is shown in Figure E.5.

²⁶For the households that choose to own a car, we have the fuel price matched to the realized driving period. However, it is not given that the household will keep the car for the entirety of the driving period, which may be two or four years. Thus, if we were to use the cross-sectional variation in fuel prices due to the precise start date of the driving period, we would be conditioning on past and/or future decisions in addition to the current and the variable would in particular not be available for households choosing not to own a car.


Figure 6.3: Forward Simulation from the Non-equilibrium Model



First off, we do not see the clear macro waves in the car age distribution in 6.3 that we observe in the actual data (Figure 2.1). There is a wave at the beginning of the simulation, coming from the large number of 2–6 year old cars in the initial car stock in 2002. This wave gradually dies out and does not proceed all the way to the age where scrappage starts to kick in. This is because the only thing coordinating the agents’ trading behavior is the macro dummy, which shifts up and down the utility of money. We do see some tiny waves in new car purchases and some ridges of these cars being held but they die out in a few years. This is because used-car prices are fixed and do not adjust to match demand and supply of car vintages. Eventually, if the macro state became degenerate, the car age distribution would just be a standing wave, reflecting the choice probabilities by car age that was indicated in Figure 6.1. Finally, note that the only thing creating the small “waves” present in Figure 6.3 is transactions costs forcing the same people to hold on to the same cars over time. As we shall see later, with equilibrium prices re-adjusting, we can have lots of trading along the ridges of the car age distribution.

Appendix E.1 presents results from the model where transaction costs are kept fixed at lower values. Table E.1 show parameter estimates where we have kept transaction costs fixed. Here, the utility of money, θ_0 , is estimated to be much higher (0.140 vs. 0.033). With these estimates, the model under-predicts keep probabilities substantially, leading to too much trading. Simulations from the model produces a car age distribution that looks very unrealistic (see Figure E.3).

Figure 6.4: Equilibrium Prices: Gasoline Cars



6.4 Equilibrium Prices: In-Sample

We now start to solve for equilibrium prices. Holding fixed the structural parameters, we loop over each of the years from 1996 to 2009 and search for the equilibrium prices that set expected excess demand equal to zero. These prices are shown in Figure 6.4. A few things are worth noticing; firstly, the price schedule is nicely behaved, downward sloping and convex, as expected. Secondly, we see that the first-year depreciation increases over time. This large first-year depreciation will tend to lower the demand for new cars, but note that the equilibrium solver is unaffected by what happens to the demand for new cars since there is zero excess demand there by assumption. Secondly, we note a dip in prices in 2008, which appears to be proportional across age groups. From closer inspections, we found that the model fits quite poorly in 2008 and predicts that too many car-owning households should sell their cars. The explanation may well be the spike in real fuel prices for both gasoline and diesel in 2008 (cf. Figure B.1); if this causes all households to want to sell their cars, the equilibrium prices will adjust downwards to counteract that and keep the market in equilibrium. Finally, we note that we do not see major waves traveling down the price schedule. As we shall see later, however, the waves are quite clear when we look at the annual *depreciation rates* rather than the actual prices in levels.

Next, we turn to comparing the predicted market activity to the realized one under the equilibrium prices. Figure 6.5 shows 4 panels; the first panel shows the negative log differences of figure 6.4, giving the annual depreciation rates implied by the equilibrium

prices (i.e. the % the used car price falls by when it ages one year). In this graph, it is easier to see waves traveling through — two such “waves” are noticeable as small dents traveling along the diagonal of the xy-plane (i.e. tracking a particular cohort of cars). The depreciation rates look somewhat jittery for the higher ages, which is mainly because there are few of those cars.²⁷ The upper right panel shows the car age distribution over time and here we notice, that the waves in the car age distribution (coming from past macro shocks traveling through in time) coincide with the waves in the depreciations (upper left panel).

The two lower panels show the number of transactions occurring from the data and predicted from the model using the equilibrium prices.²⁸ First off, we note that both the predicted and the actual number of transactions by age-category clearly mirrors the car age distribution. In particular, we see more transactions for the abundant car age groups. The predicted and observed transactions disagree in terms of how the number of transactions generally changes with the car age; in the prediction, there are clearly more transactions for car age categories around 10–15, while there are plenty of trades even for fairly young cars in the data (see also Figure ??). The reason why this can happen is that the equilibrium prices’ only purpose is to set excess demand to zero; this may happen either at high or low volumes of trade for a given car age.²⁹


We conclude the analysis of the equilibrium model by presenting a simulation forward in time, keeping fuel prices and the choice set constant and equal to their 2002 values but simulating household behavior moving forward (similarly to 6.3). Figure 6.6 shows the car age distribution in this simulation. When compared with Figure 6.3, we see the key difference between the equilibrium and non-equilibrium models; in the non-equilibrium simulation, the car stock converges to an approximately stationary distribution. For the equilibrium model, there are clear waves in the age distribution, consistent with the data (Figure 2.1). The primary difference between the simulated car stock from the equilibrium model and the real-world data is that the booms in new car sales induced by the macro state in the simulation appear to only last for the first period of the upswing; in the real data in Figure 2.1, new car sales are persistently higher throughout the booms and persistently low throughout the busts. Figure E.5 shows the macro state process and the

²⁷When there are only few of a given car in the dataset, the equilibrium prices may become very high because when the cars are rare, they will most likely be in short supply, which will push up the price to set excess demand to zero. For the diesel segment in 1996 and 1997, there are virtually no owners of the 5 highest age groups; this means that the price must be very high for those groups to remove excess demand. The optimizer got excess demand to the order of 10^{-5} and then kept increasing the prices for higher car ages without ever converging. We consider this a problem related to the 1% subsample.

²⁸To predict the number of transactions, we use the fact that expected supply and demand match up to the order of 10^{-5} . There were a small number of car ages in a few years (particularly for the rare diesel cars) where supply and demand were further apart than 10^{-3} ; in those cases, we took the average.

²⁹In particular, if it is possible to find prices so that no one wants to trade (e.g. infinitely large transactions costs), then that will constitute an equilibrium. We have not found such behavior to be an issue when working with the equilibrium solver.

Figure 6.5: Simulations Under Equilibrium Prices: Gasoline Cars




(constant) fuel prices for this simulation. The macro process is the same as was used for Figure 6.3. Figure E.6 shows additional details about the equilibrium simulation, including new-car purchases, the equilibrium prices and the scrappage pattern. The most important feature is that scrappages are highly coordinated in the model. This happens when a large cohort of cars reach the higher ages and a boom starts. The boom starting induces everyone to want to buy a new car. However, by definition this means that the supply of the cars held by those households increases. So if there are disproportionately more of a given old car age, then the price of that car will have to drop a lot. Once it approaches the scrap price, the households will start to scrap their car instead of selling it at the market. This helps to bring excess demand to zero and therefore, the equilibrium prices will try to incentivize the scrappage.

6.5 Counterfactual Simulations

In this section, we study a concrete counterfactual policy. We simulate the effects of the policy both with and without equilibrium prices. The policy we are interested in is one that changes the relative costs of ownership and usage. We therefore propose a reform that lowers the registration tax and simultaneously increases fuel prices (for example through higher fuel taxes). Such a reform changes the relative values of different car types and ages, making it less costly (in terms of depreciation) to own a newer car but more costly

Figure 6.6: Forward Simulation with Equilibrium Prices



to use cars in general, and in particular fuel-inefficient ones.³⁰ With the model, we can analyze the effects on type choice, car fleet age and driving. With equilibrium prices, we are additionally able to study the immediate and longer term effects of such a reform on scrappage. To simplify the analysis, we keep fuel prices constant except in 2012, where we increase them exogenously. Similarly, agents expect fuel prices to remain constant both before and after the unexpected policy intervention.


We choose to study a reform that lowers the price of new cars by 20% of the baseline price and raises fuel prices by 50%.³¹ We take the 2002 data as the base data and then we simulate 10 years ahead before we implement the counter-factual reform and simulate an additional 10 years under the new policy scheme. That is, in all graphs the reform is implemented in 2012.

First, we analyze the counterfactual using the non-equilibrium model. Figure 6.7 shows the outcomes of the simulation. The upper left panel shows the price schedule over time; the new car price is constant up until 2012 where it drops by 20%. The scrap price is unchanged and we use the DAF suggested depreciation rates with no change. The upper right panel shows the car age distribution in the simulation. The first 10 years of the simulation look like 6.3, with the initial wave quickly dying out and the car age distribution converging to a “standing wave” in the period up to the policy shock. After the reform, we see a shift to newer cars; in particular, there appears to be many more

³⁰Figure B.10 showed that older cars were driven more intensively. The household pays a constant utility cost of $\theta(y, m)p^k(\tau, a, p, c^o)$ per km and receives the constant utility bonus of γ_0 (since $\gamma_j = 0$ for $j > 0$ and $\phi = 0$). Comparing these two indicates whether driving is a net benefit or inconvenience to the consumer.

³¹These values were chosen so that the reform mainly changes the optimal car age and type without drastically changing the number of households in the no-car state.

Figure 6.7: Counterfactual Simulation: Non-equilibrium Prices




1–5 year old cars in the fleet. The lower right panel shows purchases of used cars in the simulated data. We see that the transactions do not track the age distribution, as expected. Similarly, the scrappage shown in the lower left panel displays no signs of waves or coordination. Figure E.7 shows the simulated paths of the macro state and the exogenous fuel price process to aid the interpretation of Figure 6.7.

Now, we turn to simulating the counterfactual policy using the equilibrium model. We do this using the approach explained in Section 5.2. We use the same macro sequence for the equilibrium as the non-equilibrium simulations and can be seen in figure E.7 (and fuel prices are constant except for the exogenous increase of 50% in period 2012).

Figure 6.8 shows the equilibrium simulation. These simulations differ markedly from the non-equilibrium counterparts. The car age distribution displays clear waves traveling through the distribution that look very much like the waves we see in the real data. The equilibrium price schedule is shown over time in the upper left panel. The prices display “ripples” traveling diagonally through the graph, coinciding with the peaks in the car age distribution: one ripple starts for 1-year old cars in 2002 and one for 7-year old cars in 2002. Both of these originate right at the end of a boom in new car sales, indicating that the cars were scrapped and replaced with new cars (or there was a chain of trades).

Note how the prices used of used cars adjust in equilibrium when exogenously changing the new car prices. In the non-equilibrium model this happened by construction since depreciation rates were kept fixed at the constant DAF depreciation rates. However, in the equilibrium model prices of used cars are allowed to vary freely and adjust endogenously to prevent any excess demand of used cars. When prices of new cars decrease exogenously,


Figure 6.8: Counterfactual Simulation: Equilibrium Prices



so does the prices used cars - but as an equilibrium outcome of the model.

In the lower left panel, there is a spike in the scrappage in the reform year for the wave of 15–18 year old cars. The intuition is the following; the reform makes cars cheaper to buy but more expensive to own so it no longer makes sense to hold on to very old cars. This shift in incentives is the same for the non-equilibrium model but the response is remarkably different due to the equilibrium prices; all households have a higher probability of buying a new car but therefore also a higher probability of supplying their currently held used car. This means that if there are waves — i.e. a higher stock of cars of particular ages — then there will be a disproportionate increase in the supply of cars of those ages. The equilibrium prices will therefore have to drop further for those age groups to set excess demand to zero. This brings prices closer to the scrap value, which results in the large, synchronized spike in scrappage in that year. Since scrapped cars do not contribute to excess demand, the prices of the oldest car categories can drop very far down without increasing excess demand. If households did not have stationary expectations about future used car prices, this effect might be dampened somewhat. Currently, when they see the equilibrium prices dropping close to the scrap price, they never expect them to become better again and thus, they might as well scrap their cars sooner rather than later. Figure 6.9 shows a rotated view of the car age distribution in Figure 6.8, which makes it easier to see the new car sales replacing the old wave of cars being scrapped.

Figure 6.9: Counterfactual Simulation of the Car Stock under Equilibrium Prices



7 Conclusion

This paper develops a novel dynamic model of vehicle choice and utilization that includes endogenous scrappage decisions and macroeconomic shocks. We estimate this model on detailed Danish data, and find that we can replicate the observed “waves” in the Danish vehicle fleet caused by macroeconomic recessions and upturns. Moreover, the model can replicate the observed patterns in scrappage and transactions over the business cycle. Our simulations clearly illustrate the importance of accounting for equilibrium price adjustments for creating realistic simulations of the car age distribution into the future. We find the resulting equilibrium price functions to generally be nicely behaved, downwards sloping and convex in age.

We illustrate the usefulness of the model by implementing a counterfactual reform that changes the balance between fixed and variable costs of cars. In the simulation, the reform induces a shift towards new car purchases but comes at the cost of accelerated scrappage of older cars. This scrappage pattern cannot be replicated by the corresponding model without equilibrium prices; it is generated by the combination of the equilibrium prices and the waves in the car fleet that comes from past macro shocks.

The model is uniquely well-suited for analyzing the long-run effects of car tax policies on the age of the vehicle fleet. Moreover, the model gives predictions on household driving and type choice decisions, which allows for a full analysis of the policy implications for tax revenue, driving, emissions, car fleet age and scrappage. Most models in the literature tend to emphasize the short or medium run effects of tax policies.

A lot of important tasks remain for future research; most importantly, we find that transactions costs need to be very high to rationalize the data. We conjecture that a

more realistic modeling of the marginal utility of money may remedy this. Secondly, while the theoretical model admits more realism, we simplify the model in our estimation by assuming that consumers have “stationary” expectations about future equilibrium prices. We propose a simple way of relaxing this assumption by allowing consumers to base their expectations on the macro state and fuel prices, but this is certainly an area with interesting prospects for future research.

A Appendix: Income Transitions

Table A.1 shows the results from the estimation of the equation

$$\log y_{it} = \rho_0 + \rho_1 \log y_{it-1} + \rho_2 s_{it} + \rho_3 s_{it}^2 + \rho_4 m_{it} + \text{error}_{it}. \quad (\text{A.1})$$

We find that controlling for the age profile of income, the AR coefficient is $\hat{\rho}_1 = 0.853$ (Table A.1). We note, however, that the effect of the macro state, m_t , is significant but only implies minor changes in average income “growth” of about -0.4% p.a. The negative sign is very puzzling. We have in and we note that this is presumably driven by the large dummy of 5.3% in 2002 (recesseion) and perhaps also the low dummies in 1999 and 2000.

One explanation for the unexpected sign of the macro dummy is that unemployment insurance is almost universal in Denmark. Since our income measure also captures transfers, the income does not drop to zero for unemployed households. To get around this, we have tried running the AR regression using only wage-based income. We also expand the horizon. The results are shown in A.2 and A.3. We still find the puzzling negative sign on the macro dummy for the wage process as well, but we note that again, real wage growth was not that low during the 2001–2003 mild recession and actually higher than during the boom in the 1990s (which started in 1994). Our problems with finding a clear relationship between incomes at the micro level and the macro state defined based on real GDP growth indicates that the link between the macro cycle in the traditional binary understanding and the micro level is perhaps not that clear cut.

Table A.1: AR regressions for income

	(1)	(2)	(3)	(4)
L.log_real_inc	0.878*** (6766.21)	0.878*** (6765.94)	0.798*** (4888.10)	0.799*** (4898.15)
$\mathbf{1}_{m=1}$		0.000105 (0.47)		-0.00416*** (-18.94)
agem			0.0116*** (554.51)	0.0115*** (549.73)
agemsq			-0.000145*** (-366.77)	-0.000144*** (-363.00)
<i>Year dummies^a</i>				
1998			0.0172***	
1999			0.00635***	
2000			-0.00754***	
2001			0.0160***	
2002			0.0530***	
2003			0.0136***	
2004			0.0237***	
2005			0.0332***	
2006			0.0445***	
2007			0.0534***	
2008			0.0296***	
2009			0.0107***	
_cons	1.564*** (948.47)	1.564*** (944.36)	2.452*** (1214.31)	2.469*** (1233.30)
<i>N</i>	17,053,312	17,053,312	17,053,312	17,053,312

^a: We omit standard errors for year dummies for easier overview.

Selection: Includes only couples and years strictly between 1997 and 2007.

Note: The income measure also includes transfers.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.2: AR regressions for log real wages

	(1)	(2)	(3)
lagged log wage	0.875*** (0.00)	0.875*** (0.00)	0.847*** (0.00)
$m_t = 1$		0.008*** (0.00)	-0.001*** (0.00)
Age (male)			0.044*** (0.00)
Age squared			-0.001*** (0.00)
Constant	1.603*** (0.00)	1.593*** (0.00)	1.205*** (0.00)
N	14,988,295	14,988,295	14,988,295
r ²	0.650	0.650	0.663

Selection: Only couples with male aged 18 to 65 and all years [1992;2009].

The explained variable only measures wage income.

B Appendix: Background and Data

In this appendix, we go into details about our dataset and institutional background that have been omitted from the main text in Section 2.

B.1 Institutional Background

Figure B.1 shows the fuel prices for gasoline and diesel cars respectively over the sample period. Both have increased and diesel prices have converged towards gasoline prices. The fuel price composition over time in the sample period is shown for gasoline in Figure B.2 and for diesel in Figure B.3. The figures show that the main variation in fuel prices in our sample period 1996–2009 comes from the product price.

To shed light on the Danish car taxation in a European perspective, Figure B.4 shows the price of the same car, a Toyota Avensis, in different European countries. First off, the figure shows that the Danish price including taxes is the highest, approximately 50% larger than the second-highest (Portugal). Secondly, the price net of tax is the lowest in Denmark, consistent with the intuition that car dealers reduce their markups in the higher tax environment.

B.2 Additional Descriptives

Figure B.5 shows the number of transactions by car age and over time. When compared to the car age distribution in Figure 2.1, we clearly see that the “waves” appear in both graphs. This indicates that transactions tend to follow the age distribution.

Table A.3: AR regressions for log real wage

	(1)	(2)	(3)
Lagged log wage	0.842*** (0.00)	0.847*** (0.00)	0.845*** (0.00)
$m_t = 1$	-0.003*** (0.00)	-0.001*** (0.00)	
Age (male)		0.044*** (0.00)	0.045*** (0.00)
Age squared		-0.001*** (0.00)	-0.001*** (0.00)
<i>Year dummies^a</i>			
1994			0.019***
1995			0.033***
1996			0.035***
1997			0.053***
1998			0.061***
1999			0.059***
2000			0.049***
2001			0.064***
2002			0.049***
2003			0.035***
2004			0.068***
2005			0.074***
2006			0.092***
2007			0.101***
2008			0.086***
2009			0.048***
Constant	2.053*** (0.03)	1.205*** (0.00)	1.178*** (0.00)
Age dummies	Yes	No	No
N	14,988,295	14,988,295	14,988,295
r ²	0.667	0.663	0.664

^a: We omit standard errors for year dummies for easier overview.

Selection: Only couples with male aged 18 to 65 and all years [1992;2009].




Figure B.1: Real Fuel Prices Over Time




Figure B.2: Composition of the Gasoline Price (Octane 95)




Figure B.3: Composition of the Gasoline Price (Octane 95)




Figure B.4: MRSP For a Toyota Avensis: Differences in Europe




Figure B.5: Purchases by Car Age Over Time (new cars omitted)

Table B.1: Number of Cars Owned per Household Over Time

	0 cars	1 car	2 cars	> 2 cars	<i>N</i>
1996	.4910	.4418	.06294	.004279	1,985,421
1997	.4469	.4644	.08211	.006531	2,001,998
1998	.4198	.4757	.09608	.008402	1,995,553
1999	.4139	.4803	.09792	.007862	1,973,977
2000	.4113	.4832	.09827	.007278	1,947,799
2001	.4053	.4849	.1025	.007278	1,950,103
2002	.3975	.4868	.1082	.007586	1,965,165
2003	.3930	.4856	.1134	.008026	1,975,094
2004	.3914	.4823	.1178	.008448	1,980,979
2005	.3820	.4812	.1273	.009586	1,988,611
2006	.3744	.4793	.1357	.01060	1,989,600
2007	.3675	.4775	.1433	.01168	2,003,445
2008	.3668	.4770	.1448	.01143	2,016,840
2009	.3721	.4723	.1441	.01153	2,022,166
Total	.4023	.4765	.1126	.008622	27,796,751

Table B.1 shows the shares of households owning zero, one, two or more than two cars for each year in our sample.

Figure B.6 shows that the cars in Denmark are typically handed down through a long chain of owners with a mode of 5 owners for a 15 year old car. The figure takes all cars in 2009 that are 15 years old (i.e. first registered in Denmark in 1994) and where we observe the first owners of the car. The first owners is observed for about two thirds of the cars. The reason for restricting to 15 year old cars in 2009 is to avoid mixing car ages together, which will produce a mixed picture due to scrappage and missing data. The figure indicates that the most common is for a car to have switched owners once every third year.

B.3 Scrappage

We do not observe scrappage per se in our dataset. Instead, we define scrappage as occurring when a car ownership ends and we never see a new one starting for that car. This measure is not perfect because an individual may choose to de-register his car and leave it in his garage for a while. This may be particularly important for specialty cars and vintage cars but since these are outside the scope of our paper, we are not too concerned with behavior of that sort.

We first consider the scrappage together with transactions; this highlights that when an individual decides to sell a car in the model, he may either sell it on the used-car market or at the scrap price. Figure B.7 shows for each car age the number of transactions in the data and the number of scrappages. Firstly, the figure shows that the number of




Figure B.6: Number of owners for 15 year old cars in 2009

transactions increases up to a car age of 3, after which it is relatively constant up until car ages of 14, whereafter it falls linearly until age 23. The number is slightly higher for 24, but that is because we have truncated the car age distribution. The scrappage frequency increases up to age 16 after which it falls (because there are not that many cars left to scrap). Recall that the annual scrappage in percent of the car stock increases over age categories in Figure 2.8. We see the same spikes in scrappage in even years that correspond with the inspection years.

Table B.2 shows the number of scrappages in our data for all the sample years. We note that we have exceptionally few scrappage observations in 1996 and 1997 while 1998 appears to be half-way to the average that persists thereafter. To validate the number of scrappages, we also show the number of scrappage subsidies paid out for environmentally friendly scrappage of older cars. The data comes from the website bilordning.dk, which is maintained by *Sekretariatet for Miljøordning for Biler*, a government office under the Ministry of the Environment overseeing vehicle scrappage and the scrappage subsidy. The subsidy was introduced in July 2000 and has been fixed at 1,500 DKK throughout our sample period (it was changed in 2014). Given the introduction half-way through the year, the 30,439 subsidies corresponds closely with the increasing trend from 60,000 up to just under 100,000 subsidies paid out annually. The number is lower than the number of scrappages by our definition of scrappage, which is to be expected for a number of reasons; firstly, some cars are de-registered for a few years and then re-register again later. This may explain the higher number of scrappages later in our sample and perhaps particularly some of the younger scrapped cars, we see in Figure 2.9 for the latest years. Our dataset was drawn from the license plate registers in September 2011, so we do not observe cars that have since then been re-registered. Secondly, some cars are exported, which we do not observe. However, given the higher used car prices in Denmark, we expect this to be




Figure B.7: Number of Scrappages and Number of Transactions by Car Age

a minor issue. Thirdly, some cars are kept as collectibles (e.g. vintage or specialty cars). These cars are outside the focus of this paper so we do not worry about not being able to fit those cars.

Figure B.8 shows the subsidies paid out by the age of the car being scrapped. The data does not match up with the other data sources of `bilordning.dk`, indicating that they may have missing observations of car age. Where it is observed, we see that while the earliest subsidies were paid out to very old cars, the car age distribution after this looks somewhat stable. The biggest group is the 16–20 year old cars, but the number of 21–25 year old cars has grown from 3930 to 19526 from 2002 to 2009. Whether expected lifetime of cars has gone up or this was a transitory thing is hard to say from these descriptives alone.

Figure B.9 shows the number of ownership spells ending each year in our data, going back to 1992. The number increases from around 100,000 in 1992 up to over 400,000 in 1999, after which it appears to stabilize at this level. We note that there are fewer periods ending in the years prior to 1998, but not enough so to explain why we have so few scrappage incidents prior to 1998.

B.4 Driving

Driving in our data comes from the safety inspections administered by the Ministry of Transportation. They occur when the car is aged 4 first and then every 2nd year thereafter.

Table B.2: Car Scrappage by Year: Sample Data and Scrappage Subsidies

Year	Scrappage in Data	Subsidies
1996	3884	0
1997	4798	0
1998	47509	0
1999	136015	0
2000	120257	30439
2001	102258	68583
2002	105398	79836
2003	102452	86141
2004	110467	92700
2005	113246	98295
2006	127199	94268
2007	146709	91712
2008	141416	95747
2009	128017	93543




Figure B.8: Scrappage Subsidies Paid by Car Age (source: bilordning.dk)




Figure B.9: Number of Ownership Spells Ending

In practice, the test date varies by about 3 months around this. At these inspections, the odometer is measured and we find the vehicle kilometers traveled (VKT) as the first differences in the odometer readings.

Figure B.10 shows the average VKT conditional on the car age. We have split the data into 20 quantiles depending on the age of the car (for the observations where a car is present). Within each of these groups, we show the average VKT. Note that for the typical car, the VKT will be the same when the car is between zero and four years old. However, some cars may have an inspection before the planned one at four years, which explains why the average still changes before four years. The graph shows that households with older cars tend to drive less. The VKT increases up towards an age of four but recall that for the typical car, we only observe the average driving for the full period from zero to four years of age.

Figure B.11 shows the VKT by the household income. We have split the observations into 20 quantiles based on income (for the households where we observe VKT). Within each of these quantiles, we show the average VKT. Note that in the data used for estimation, we have split household income in two if the household owns two cars, and in three for three cars, etc. In Figure B.11, we show the relationship with the un-split income — the figure looks similar when we have split the income except for a small hump mid way through. Figure B.11 shows that high

Table B.3 shows regressions of vehicle kilometers traveled (VKT) on different sets of controls for the full sample. We find that the average elasticity of the price per kilometer (PPK, defined as the fuel price divided by the fuel efficiency in km/l) is at the lowest




Figure B.10: Vehicle Kilometers Traveled by Car Age




Figure B.11: Vehicle Kilometers Traveled by Real Income

−179% unless we control for a diesel dummy, in which case it drops to −41.9%. This big difference is intuitively clear; the difference in both fuel price and fuel efficiency is substantial between the gasoline and the fuel price segment. Without a dummy, we are attributing all differences in driving to the price variable and not allowing a level shift. On the other hand, from the point of view of the model, there should not be a level shift between the two segments unless it is due to endogenous selection based on the PPK variable in the sense that households needing to drive a lot choose a car that will allow them to do so cheaply.³² Nevertheless, we find these high price elasticities of driving puzzling, in particular in light of the findings of Munk-Nielsen (2015) and Gillingham and Munk-Nielsen (2015), who find much lower elasticities. We conjecture that adding more controls in line with those studies will lower the elasticity.

Table B.3: Regressions of vehicle kilometers traveled (VKT) on controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Price per km (PPK)	-107.3718*** (0.11)	-96.2963*** (0.19)	-99.9269*** (0.12)	-104.3433*** (0.11)	-86.6610*** (0.19)	-107.5702*** (0.24)	-9.4086*** (0.31)
PPK X income		-2.3936*** (0.03)			-2.1379*** (0.03)	-2.1749*** (0.03)	-2.1158*** (0.03)
PPK X income squared		0.0005*** (0.00)			0.0004*** (0.00)	0.0004*** (0.00)	0.0004*** (0.00)
PPK X dBoom=1						33.5003*** (0.23)	-5.8843*** (0.24)
Constant	117.8539*** (0.07)	109.1657*** (0.12)	116.1846*** (0.07)	104.0167*** (0.13)	97.7115*** (0.16)	111.3234*** (0.19)	47.4324*** (0.23)
Income (100,000)		1.8640*** (0.02)			1.5193*** (0.02)	1.5401*** (0.02)	1.4798*** (0.02)
Income squared		-0.0004*** (0.00)			-0.0003*** (0.00)	-0.0003*** (0.00)	-0.0002*** (0.00)
Car age			-0.2141*** (0.00)		-0.2593*** (0.00)	-0.2864*** (0.00)	-0.2889*** (0.00)
Car age squared			-0.0201*** (0.00)		-0.0184*** (0.00)	-0.0178*** (0.00)	-0.0213*** (0.00)
Age (head)				0.9297*** (0.01)	0.8270*** (0.01)	0.8259*** (0.01)	0.8155*** (0.01)
Age squared				-0.0138*** (0.00)	-0.0128*** (0.00)	-0.0128*** (0.00)	-0.0127*** (0.00)
$m_t = 1$						-21.4665*** (0.15)	4.1525*** (0.15)
Diesel dummy							17.3834*** (0.03)
N	15,018,013	15,018,013	15,018,013	15,018,013	15,018,013	15,018,013	15,018,013
R2	.0555	.0575	.0658	.0732	.0851	.0864	.102
Avg. PPK-elasticity ^a	-1.9860	-1.9859	-1.8483	-1.9300	-1.7859	-1.8131	-.4188

Selection: VKT in]0;1,000[and year in [1996;2009] and household age in [18;65]

^a: The avg. elasticity of driving wrt. the price per kilometer.

Income is measured in 100,000 real 2005 DKK.

³²The selection could also go the other way so that households needing to drive a lot would choose a car that would make the long drive as comfortable as possible and therefore go for a more luxurious car. Comfort and luxury tend to be correlated positively with vehicle weight and negatively with fuel efficiency.

C Appendix: Flexible Price Function Specification

This appendix describes a flexible specification for the price function for cars. We could estimate this price function as a first step along with the estimation of the structural parameters of the model using the Danish Register data, rather than using the price depreciation rates given to us from the Danish Automobile Dealers Association. Those depreciation rates may or may not be reasonable values to start the estimation with. The main drawback of using them is that they do not shift with changes in fuel prices or macro conditions. Below we describe a flexible price function that can allow fuel prices and macro shocks to enter and affect depreciation rates, and we have the ability to estimate the parameters using unconstrained optimization algorithms, yet the estimated price functions are constrained (via minimal functional form assumptions described below) to always be downward sloping.

We do assume that new car prices and scrap prices are determined exogenously. The exogenous new car price assumption is a consequence of the “small open economy” model for Denmark, where all cars are imported and we assume demand for new cars from Denmark is an insignificant share of worldwide demand for new cars. However it may be useful to allow new car prices to vary with macro shocks (which we initially assume to pertain to Denmark only, but which could be correlated with a worldwide macro shock, e.g. the 2008 Great Recession) and the specification below allows for this possibility.

Similarly we assume there is an infinitely elastic demand for vehicles as scrap, and this sets an exogenously determined scrappage price for cars, and this could also depend on fuel prices and macro shocks.

Recall the key state variables in the model: (a, m, p, τ) where τ is the type of car, a is age of car, m is the macro shock, and p is the fuel price. We conjectured that the equilibrium in the Danish car market could be found for prices of the form $P(a, m, p, \tau)$, i.e. we assumed that the price function is not a function of the age distribution of the vehicle stock but only of the current macro shock and fuel price. If $a = 0$ is a brand new car, then $\bar{P}(\tau) = P(0, m, p, \tau)$ is the “boundary condition” for the price of a new car under the small open economy assumption, where $\bar{P}(\tau)$ is the average suggested retail price of a new car of type τ . If we had enough time series data to detect any variation in new car prices with fuel prices or macro shocks, it may be possible to fit a function $\bar{P}(m, p, \tau)$ where new car prices shift with fuel prices and macro shocks (e.g. gas guzzlers sell at a discount when fuel prices are high, whereas high fuel efficiency cars sell at a premium when fuel prices are high, and luxury cars are discounted and economy cars sell at relatively higher prices during a recession, whereas luxury car prices are relatively higher and economy car prices are relatively lower during a recession, etc). But for now our data only allow us to identify $\bar{P}(\tau)$ which does not depend on (m, p) .

We may be able to estimate scrap prices $\underline{P}(\tau)$ from the model, but for now we will fix

this price at approximately 3,000 Danish Kroner, independent of τ or of (m, p) . It may be that export of old Mercedes, BMW to developing countries, or “collector value” implies a higher value than this floor scrap value for certain types of cars, but for now we go with this basic assumption of a constant scrap price for all types of vehicles, regardless of fuel prices or macro conditions.

To understand the basic flexible secondary price specification, first ignore the effect of (m, p) so that the prices are just a function of a , $P(a)$ (and for simplicity we suppress the car type indicator τ as well. If $a = 0$ is a new car and $a = 20$ is the oldest car allowed, we have the boundary conditions that $P(0) = \bar{P}$ and $P(20) = \underline{P}$. In the illustration below we set $\bar{P} = 180000$ and $\underline{P} = 3000$.

Let θ be an unconstrained 19×1 parameter vector. We will now write a specification for $P(a)$ that depends on these 19 unconstrained parameters θ in a way that guarantees that $P(a)$ is always decreasing in a and satisfies the boundary conditions $P(0) = \bar{P}$ and $P(20) = \underline{P}$. The specification that does this, $P(a, \theta)$ is given below

$$P(a, \theta) = \underline{P} + (\bar{P} - \underline{P}) \prod_{i=1}^a \rho(\theta_i), \quad a = 1, \dots, 19 \quad (\text{C.1})$$

where we define $\rho(\theta_0) \equiv 1$ and

$$\rho(\theta_i) = \frac{\exp\{\theta_i\}}{1 + \exp\{\theta_i\}} \quad (\text{C.2})$$


for $i = 1, \dots, 19$. Note that the θ_i can take any value in the interval $(-\infty, \infty)$ and for any vector $\theta \in R^{19}$ the implied price function $P(a, \theta)$ will be decreasing in a . Further we can impose restrictions to reduce the dimensionality of the vector θ . For example we could restrict θ to take the form

$$\theta = (\theta_1, \theta_1, \dots, \theta_1) \quad (\text{C.3})$$

so that $\theta \in R^{19}$ depends only on a single unknown parameter $\theta_1 \in R^1$. Or we could partition θ to depend on just two parameters (θ_1, θ_2) as follows

$$\theta = (\theta_1, \theta_1, \dots, \theta_1, \theta_2, \dots, \theta_2) \quad (\text{C.4})$$

so the first J_1 components of θ take the value θ_1 and the remaining $19 - J_1$ components of θ take the value θ_2 , and so forth. This gives us quite a bit of flexibility in how flexible we want to allow the price function $P(a, \theta)$ to be as a function of a . Even when the price function is restricted to depend on only a single parameter θ_1 , the implied price function $P(a, \theta_1)$ can assume many different shapes as θ_1 ranges over the interval $(-\infty, \infty)$ as illustrated in figure C below.



Now, taking this basic flexible specification for the price of cars as a function of age, we can allow these functions to shift with macro shocks and fuel prices in a flexible way also by a small modification of the basic functional form in equation (C.1) above. In addition to the 19×1 vector θ , let α be a $K \times 1$ vector that can flexibly parameterize the dependence of the price function on (m, p) . Let $f(m, p, \alpha)$ be some function of (m, p, α) such as linear-in-parameters $f(m, p, \alpha) = \alpha_1 + \alpha_2 m + \alpha_3 p$.

$$P(a, m, p, \theta, \alpha) = \underline{P} + (\overline{P} - \underline{P}) \prod_{i=1}^a \rho(m, p, \theta_i, \alpha), \quad a = 1, \dots, 19 \quad (\text{C.5})$$

where we define $\rho(\theta_0) \equiv 1$ and

$$\rho(m, p, \theta_i, \alpha) = \frac{\exp\{\theta_i + f(m, p, \alpha)\}}{1 + \exp\{\theta_i + f(m, p, \alpha)\}} \quad (\text{C.6})$$

for $i = 1, \dots, 19$. Note by construction we have $P(0, m, p, \tau) = \overline{P}(m, p, \tau)$.

C.1 Derivatives of the price function with respect to (θ, α)

Let θ_j be one of the independent subparameters (or components) of the 19×1 vector $\theta = (\theta_1, \dots, \theta_{19})$. In the case of parameter restrictions, such as the most restrictive specification $\theta = (\theta_1, \dots, \theta_1)$, then the θ vector would depend on only one independent subparameter θ_1 , whereas if θ depends on two free parameters (independent components) θ_1 and θ_2 then $\theta = (\theta_1, \dots, \theta_1, \theta_2, \dots, \theta_2)$. Suppose we have a specification where the overall 19×1 θ vector depends on J free parameters $(\theta_1, \dots, \theta_J)$, with the most flexible case being $J = 19$. Partition the set of indices $\{1, 2, \dots, 19\}$ into J subintervals

$\{1, 2, \dots, 19\} = (I_1, I_2, \dots, I_J)$ where $I_1 = \{1, \dots, \bar{I}_1\}$, and $I_2 = \{\bar{I}_1 + 1, \dots, \bar{I}_2\}$, and so on until $I_J = \{\bar{I}_{J-1} + 1, \dots, 19\}$. Then we have

$$\frac{\partial}{\partial \theta_j} P(a, m, p, \tau) = [P(a, m, p, \tau) - \underline{P}] \left[\sum_{i=1}^a [1 - \rho(m, p, \theta_i, \alpha)] I\{i \in I_j\} \right] \quad (\text{C.7})$$

$$\frac{\partial}{\partial \alpha} P(a, m, p, \tau) = [P(a, m, p, \tau) - \underline{P}] \left[\sum_{i=1}^a [1 - \rho(m, p, \theta_i, \alpha)] \frac{\partial}{\partial \alpha} f(p, m, \alpha) \right]. \quad (\text{C.8})$$

Of course we also have $\frac{\partial}{\partial \theta_j} P(0, m, p, \tau) = 0$ and $\frac{\partial}{\partial \alpha} P(0, m, p, \tau) = 0$ since $P(0, m, p, \tau) = \bar{P}(m, p, \tau)$ by construction, and the latter does not depend on (θ, α) .

C.2 Non-monotonic specification

We have found that it is difficult to estimate all parameters of the least restrictive monotonic specification above (i.e. where we have separate depreciation rates for all 19 age groups from age 1 to age 19 with a separate θ_a parameter for each value of a). The reason is that when there is rapid initial depreciation (i.e. large negative “early values” for θ_a , $a = 1, 2, 3, \dots$), there is less room for maneuvering for the values of the later depreciation parameters θ_a , $a = 15, 16, \dots, 19$. If the car’s secondhand price is already close to scrap by age 12, then the depreciation rate parameters for $a = 13, 14, \dots, 19$ hardly matter, and this shows up as parameters that have gradients close to zero and this tends to make the likelihood hessian matrix poorly conditioned (i.e. close to singular). We are able to estimate the first few depreciation parameters, such as restricted version of the a specification above where we estimate only $(\theta_1, \theta_2, \theta_3)$ where θ_1 governs depreciation for ages $1, \dots, I_1$, θ_2 governs depreciation over ages $I_1 + 1, \dots, I_2$, and θ_3 governs depreciation for the remaining ages $a = I_2 + 1, \dots, 19$.

But it might be useful to try a less restrictive specification of secondary market prices where we drop the monotonicity restriction. In this specification we do restrict prices to lie in the interval $[\underline{P}(\tau, m, p), \bar{P}(\tau, m, p)]$ but we do not require the price function to be monotonically decreasing. It will have unrestricted choices of depreciation parameters θ_a , $a = 1, \dots, 19$ but these parameters will be more “orthogonal” than in the case where we impose a monotonicity restriction as above, since a choice for θ_a does not restrict in any way the choices of possible prices in other age categories a' for $a' \neq a$.

This specification is rather simple: θ_a is just the parameter of a logit function that specifies the fraction of the distance between $\underline{P}(\tau, m, p)$ and $\bar{P}(\tau, m, p)$ the secondary market price $P(\tau, a, m, p)$ lies:

$$\rho_a(\theta_a) = \frac{\exp\{\theta_a\}}{1 + \exp\{\theta_a\}}, \quad (\text{C.9})$$

and

$$P(\tau, a, m, p, \theta) = \sum_{i=1}^{19} I\{i = a\} [\bar{P}(\tau, m, p)\rho_a(\theta_a) + \underline{P}(\tau, m, p)(1 - \rho_a(\theta_a))]. \quad (\text{C.10})$$

This specification will ensure that $\underline{P}(\tau, m, p) \leq P(\tau, a, m, p, \theta) \leq \bar{P}(\tau, m, p)$ for any choice of $\theta = (\theta_1, \dots, \theta_{19}) \in R^{19}$ but it does not enforce any monotonicity in $P(\tau, a, m, p, \theta)$ as a function of a .

The gradient of $P(\tau, a, m, p, \theta)$ with respect to θ_a is easy to compute using the fact that $\frac{\partial}{\partial \theta_a} \rho_a(\theta_a) = \rho_a(\theta_a)(1 - \rho_a(\theta_a))$.

D Appendix: Test equilibria

This appendix describes a simple infinite horizon model for constructing equilibria in a stationary (no macro or fuel price shocks) case with no transactions cost to provide a test bed to check that the equilibrium solver we develop finds the correct equilibrium. In the process of doing this, we discovered the possibility of multiple Pareto-ranked equilibria in the secondary market for cars.

Consider a simplified model where there is only one type of car (though of different ages) and consumers live forever. We assume any utility from driving is subsumed into the quasi-linear specification where the disutility of owning a car can be expressed in monetary equivalent units as akin to a “maintenance cost” $m(a)$ which is increasing in the age of the car a . Thus we can convert the utility maximization problem into a “dual” cost-minimization problem where a consumer chooses a trading strategy to minimize the discounted costs of holding a sequence of cars over an infinite horizon.

When there are no transactions costs, it will be optimal for the consumer to trade every period for a preferred vehicle age a^* . The per period cost of the strategy of holding a car of age a^* for one period and then selling it and buying another car of age a^* is

$$c(a^*) = m(a^*) + P(a^*) - \beta P(a^* + 1), \quad (\text{D.1})$$

that is, the one period holding cost $c(a^*)$ is the sum of the “maintenance cost” $m(a^*)$ plus the expected depreciation $P(a^*) - \beta P(a^* + 1)$, where $\beta \in [0, 1)$ is the consumer’s discount factor. The present discounted value of holding costs over an infinite horizon is then simply $c(a^*)/(1 - \beta)$.

If all consumers have the same discount factor β and have homogeneous preferences, then in equilibrium all consumers must be indifferent between holding any of the available ages of vehicles. Assume that cars that are older than some threshold age γ are scrapped, and we let $a = 0$ denote a brand new vehicle. Then cars of ages $0, 1, \dots, \gamma - 1$ will be held by consumers, and once a car reaches age γ it will be scrapped for the scrappage

price $\underline{P}$. As we have done in the paper, we assume there is an infinitely elastic demand for vehicles for their scrap value at price $\underline{P}$ and there is also an infinitely elastic supply of new vehicles at price $\bar{P}$. If we also assume that $m(a)$ is strictly monotonic, then this implies that for each age $a \in \{0, 1, \dots, \gamma - 1\}$ we have $P(a) \in [\underline{P}, \bar{P}]$. Clearly we must have $P(0) = \bar{P}$ and $P(\gamma) = \underline{P}$ and these prices are thus exogenously fixed. The remaining prices $P(1), P(2), \dots, P(\gamma - 1)$ are determined endogenously in equilibrium.

The equilibrium condition is that these prices must adjust to make consumers indifferent about holding any of the ages of vehicles, $a \in \{0, 1, \dots, \gamma - 1\}$, or

$$c(0) = c(1) = \dots = c(\gamma - 1). \quad (\text{D.2})$$

These indifference restrictions imply a system of $\gamma - 1$ linear equations in the $\gamma - 1$ unknowns $P(1), \dots, P(\gamma - 1)$. This system can be written in matrix form as

$$X \times P = Y \quad (\text{D.3})$$

where $P' = (P(1), \dots, P(\gamma - 1))$ and X is the $\gamma - 1 \times \gamma - 1$ matrix given by

$$X = \begin{bmatrix} (1 + \beta) & -\beta & 0 & \dots & 0 & 0 \\ -1 & (1 + \beta) & -\beta & 0 & \dots & 0 \\ 0 & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & \dots & -1 & (1 + \beta) & -\beta \\ 0 & 0 & \dots & 0 & -1 & (1 + \beta) \end{bmatrix}, \quad (\text{D.4})$$

and Y is a $\gamma - 1 \times 1$ vector given by

$$Y = \begin{bmatrix} m(0) - m(1) + \bar{P} \\ m(1) - m(2) \\ \dots \\ m(\gamma - 2) - m(\gamma - 1) \\ m(\gamma - 1) - m(\gamma) + \beta \underline{P} \end{bmatrix}. \quad (\text{D.5})$$

Notice that we do not impose a) monotonicity or b) the restriction that $P(a) \in [\underline{P}, \bar{P}]$ on the solution P to the linear system (D.3). Thus, we need to check if the solution has these properties. If it does, it is an equilibrium since the price vector results in all consumers being indifferent between holding any one of the available vehicles that are traded in the new or secondary markets, $a \in \{0, 1, \dots, \gamma - 1\}$. The equilibrium “quantities” are the holdings of vehicles of these different ages. It is easy to see that without any accidents or “endogenous scrappage” of cars prior to the scrappage threshold age γ , then the equilibrium or steady state age distribution of cars will be uniform on the interval $\{0, \dots, \gamma - 1\}$, so that a fraction $1/(\gamma - 1)$ of the total vehicle stock will be

of age a at the *beginning of each period*. This implies in particular that (assuming all consumers hold just one car) that the fraction $1/(\gamma - 1)$ of the population will buy a new car each period, and the corresponding fraction will scrap their cars, so the market will be in “flow equilibrium”. It will also be in “stock equilibrium” since the fact that consumers are indifferent about which age vehicle they own and hold, they can be arranged so that their demand for the different ages is also uniform, matching the supply. Thus there will be zero excess demand for any vehicle age $a \in \{0, 1, \dots, \gamma - 1\}$ at the price function given above.

We have demonstrated that multiple equilibria are possible in this model. That is, we can find different values of γ and different corresponding price vectors P_γ for which P_γ satisfies the linear system (D.3) and is also monotonically decreasing from $\bar{P}$ to $\underline{P}$. In the examples we have computed these equilibria can be Pareto-ranked, with the equilibria corresponding to larger values of γ being Pareto-preferred by consumers to equilibria with smaller values of γ . That is, per period holding costs are higher in equilibria where cars are scrapped “prematurely”.

However γ cannot be increased to arbitrarily large values for a fixed $m(a)$ function. Eventually for large enough γ , the solution P_γ to the linear system (D.3) is no longer monotonically decreasing from $\bar{P}$ to $\underline{P}$ and thus no longer constitutes an equilibrium.

We have found there is a largest possible γ for any $m(a)$ function, and this γ turns out to be the optimal scrappage threshold to a “social planning problem” where there is no secondary market and a single representative consumer simply chooses an age threshold at which to replace their current car with a brand new one. The value function $V(a)$ to this problem is given by

$$V(a) = \min [\bar{P} - \underline{P} + m(0) + \beta V(1), m(a) + \beta V(a + 1)]. \quad (\text{D.6})$$

It is easy to see from the Bellman equation above that $V(0) = m(0) + \beta V(1)$ and thus, any consumer who is “endowed” with a brand new car would never immediately replace it with another new one since this would involve the additional replacement cost $\bar{P} - \underline{P}$. However if $m(a)$ is increasing sufficiently rapidly there will be a finite age, γ , for which we have

$$V(\gamma) = \bar{P} - \underline{P} + m(0) + \beta V(1) = \bar{P} - \underline{P} + V(0). \quad (\text{D.7})$$

Thus, the optimal scrappage threshold γ is the smallest value of a at which it is optimal for the representative consumer to scrap their car and buy a new one.

Using this value function, we can define a *shadow price function* $P(a)$ by

$$P(a) = \bar{P} - [V(a) - V(0)]. \quad (\text{D.8})$$

Notice that this shadow price function statisfies $P(0) = \bar{P}$, $P(\gamma) = \underline{P}$, and $P(a)$ is

monotonically declining in a for the values of a for which $V(a)$ is monotonically increasing in a , which is the set of $a \in \{0, 1, \dots, \gamma - 1\}$. However it is not hard to see from the Bellman equation (D.6) that for $a < \gamma$ we have

$$V(a) = m(a) + \beta V(a + 1), \quad (\text{D.9})$$

which simply says that it is optimal for the consumer to keep their car if its age is younger than the optimal scrappage age γ . However using this condition, it is then easy to verify that the shadow price function (D.8) makes consumers indifferent between all car ages $a \in \{0, 1, \dots, \gamma - 1\}$,

$$m(a) + P(a) - \beta P(a + 1) = m(a') + P(a') - \beta P(a') \quad \forall a, a' \in \{0, 1, \dots, \gamma - 1\}. \quad (\text{D.10})$$

Thus it follows that the shadow price function (D.8) is an equilibrium in the secondary market, and we can show it is also the *Pareto dominant* equilibrium, i.e. the one in which consumers have the lowest holding cost and thus the highest discounted welfare.

E Appendix 5: Additional Results

This appendix contains additional results that have been omitted from the main results sections.

E.1 Estimates with Fixed Transaction Costs

Table E.1 shows estimation results where we have kept transaction costs fixed at 10,000 DKK plus 20% of the traded car's values. Comparing the parameter estimates to the preferred specification in the main text, where transaction costs are estimated, we in particular note the utility of money (θ_0), which is considerably higher here.

Figure E.1 shows the fit in terms of conditional choice probabilities. Compared to the preferred specification where the transaction cost is estimated, we see a considerable under-prediction of the keep decision. Moreover, the model produces much more probability mass for all car ages over 4 with the highest mis-match at car age 10. Turning to Figure E.2, we see that the keep probability predicted by the model changes much more with the car age than does the observed probability; at car age 4, the observed and predicted keep probabilities are about equal but while the data ends up at a probability of about 60% at the oldest car age, the predicted probability tends to zero.

Figure E.3 shows a simulation forward in time, keeping the choice set and price schedule fixed at the 2002 data values but drawing state variables from the conditional transition densities according to the model. The simulated car age distribution has no clear waves but rather shows synchronized, parallel shifts up in transactions in particular years. This

Table E.1: Structural Estimates — Fixed Transaction Costs

	Variable	Estimate	Std.err.
Model setup			
	Min. Hh. age	20	
	Max. Hh. age	85	
	# of car ages	25	
	# of car types	2	
	Clunkers in choicest	1	
β	Discount factor	0.95	
ρ	Inc. AR(1) term	1	
σ_y	Inc. s.d.	0	
ρ_p	Fuel price AR(1) term	1	
σ_y	Fuel price s.d.	0.0699	
Pr(0 0)	Macro transition	0.75	
Pr(1 1)	Macro transition	0.8	
	Accident prob.	0.0004	
λ	Logit error var.	1	
λ^{scrap}	Scrappage error var.	0.9	
Monetary Utility			
θ_0	Intercept	0.13984***	2.921e-06
θ_1	Inc.	-8.0175e-05***	2.597e-07
θ_2	Inc. sq.	5.996e-08***	2.915e-10
θ_3	Macro	-0.00055181***	3.881e-05
Driving Utility			
γ_0	Intercept	0.26509***	0.0005566
γ_1	Car age	0	
γ_2	Car age sq.	0	
γ_3	Hh. age	0	
γ_4	Hh. age squared	0	
γ_5	Macro	0	
γ_6	Macro	0	
γ_7	Macro	0	
ϕ	Squared VKT	0	
Car Utility			
$q(a)$	Car age, linear	0.34035***	0.0008795
$q(a)$	Car age, squared	-0.0013817***	3.947e-05
δ_1	Car type dummy	1.6588***	0.01175
δ_2	Car type dummy	0.00087379	0.01281
Transaction costs			
	Fixed cost	10	
	Proportional cost	0.2	
	N	169,733	

Figure E.1: Model Fit: Conditional Choice Probabilities




Figure E.2: Model Fit by State Variables





Figure E.3: Forward Simulation



pattern can be explained by the macro dummy shifting down the utility of money, making it more likely for all households to buy a new car, causing the upwards shift in the age distribution. However, the under-predicted keep probability means that households need not hold on to their cars in the following year.

E.2 Equilibrium Prices

In this section, we show additional results concerning the equilibrium price simulations. To re-iterate, the parameter estimates here are based on a first-stage estimation of the driving parameters (κ s) that are fixed in the second stage, where the structural parameters are estimated, including the fixed transaction cost. Finally, we solve for equilibrium prices in each year by setting expected excess demand equal to zero, clearing the market in each year. Figure E.4 shows simulations complementing figure 6.4 but showing all car age categories; in particular, the first- and final-year depreciations were omitted in Figure 6.4 because they make it hard to see what else happens in the figure. The large final-year depreciation may be to avoid too high scrapping earlier on.

E.3 Counterfactual Simulations

Figures E.5 and E.6 accompany Figures 6.3 and 6.6 in Section 6.5. Figure E.5 shows the macro state over the simulation and the fuel prices (which are held constant) and

Figure E.4: Simulations Under Equilibrium Prices: All Car Ages




Figure E.5: Forward Simulation with Equilibrium Prices





Figure E.6: Forward Simulation with Equilibrium Prices



In this section, we present results that are supplementary to the ones shown in section 6.5. Figure E.7 accompanies Figures 6.7 and 6.8 in showing the realized paths of the macro and fuel prices processes. Firstly, note that the fuel price is constant throughout the period except in the year 2012 where we counterfactually increase it by 50%.

To compare against the counterfactual simulation results in Figures 6.7 and 6.8, we show the corresponding graph for the actual data in Figure E.8. Note that the prices shown there are computed using the DAF suggested depreciation rates. The most important features to note are regarding purchases and scrappage; purchases clearly follow the car age distribution. In other words, we see more purchases (and thus sales) of cars age categories that are more abundant. Moreover, the scrappage distribution is distinctly different from the non-equilibrium model; in particular,

F Notation

This section provides an overview of the notation used in the paper.

Below are some of the core equations from the paper.

Figure E.7: Counterfactual Simulations: Macro and Fuel Price Processes




Figure E.8: Age Distribution, Scrappage and Purchases in the Data




Table F.1: Notation

s	Household age, $s = 20, \dots, 85$.
y	Household income.
x	Household characteristics, (s, y) .
m	Binary macro state, $m = 1$ for boom.
p	Fuel price.
k	Vehicle kilometers traveled (abbreviated VKT).
β	Annual time discount factor, $\beta = 0.95$.
τ	Car type. In the application: gasoline or diesel.
a	Discrete car age. In the application, $a \in \{0, \dots, 24\}$.
$P(\tau, 0)$	New car price (from the DAF data).
ζ_τ	Depreciation factor (from the DAF data).
$P(\tau, a, p, m)$	Car prices as they enter into the consumers' expectations.
$\underline{P}(\tau, p, m)$	Fixed scrap value for a type τ car. In the application, we assume that $\underline{P}(\tau, p, m) = \zeta_\tau^a P(\tau, 0)$, i.e. the used car price indicated by the DAF depreciation for a car of the oldest age.
$P(\tau, a, t)$	Car prices when we solve for equilibrium and allow them to vary freely over time, t , to equate supply and demand.
$\alpha(\tau, a, x)$	Accident probability.
d_s	Discrete decision about selling, $d_s \in \{-1, 0, 1\}$, where $d_s = 0$ means keeping, $d_s = 1$ means selling at the used-car market and d_s stands for scrapping the car.
d	Discrete car state, $d = (\tau, a)$. The no-car state is denoted $d = (\emptyset, \emptyset)$.
d'	Discrete decision, $d = (\tau, a, d_s)$ (depending on the context, we some times omit the scrappage decision, d_s , from d').
$D(d)$	Choiceset available to a household with car $d = (\tau, a)$.
$T(d', d, p, m)$	Trading cost function.
$c_T(\tau', a', p, m)$	Transactions cost.
$b_1(\tau', a, p, m)$	Proportional term in the transactions cost.
$b_2(\tau', a, p, m)$	Fixed term in the transactions cost.
λ	Scaling parameter in the scrappage probability.
ψ	Co-insurance rate for accidents where the car is totaled.
Utility parameters	
$\theta(y, m)$	Utility of money.
$\gamma(y, s, a, m)$	Driving utility, linear term.
ϕ	Driving utility, quadratic term.
θ_j	Parameters entering into $\theta(y, m)$, $j = 0, 1, 2, 3$.
γ_j	Parameters entering into $\gamma(y, s, a, m)$, $j = 0, 1, \dots, 7$.
κ_j	Parameters in the reduced-form driving equation. Interpreted as scaled versions of θ_j and γ_j by $-.5/\phi$ and $.5/\phi$ respectively.
ρ_j	Coefficients in the AR(1) equation for log income.
σ_y	Dispersion on the AR(1) error term for log income.
σ_p	Dispersion on the AR(1) error term for log fuel price.
$\varphi(f(d'), m', p', x')$	Value function integrated over the nested scrappage sub-problem (see equation (3.15)).
$\Gamma(S_{t+1} S_t, P_t)$	Transition density for the state variables of all households in the economy jointly.

Flow utility:

$$\begin{aligned}
u(k, \tau, a, s, p, m) &= \theta(y, m)[y - p^k(\tau, a, p, c^o)k - T] \\
&\quad + \gamma(y, s, a, m)k + \phi k^2 - q(a) + \delta_n \mathbb{1}(a = 0) + \delta_\tau, \\
\gamma(y, s, a, m) &= \gamma_0 + \gamma_1 a + \gamma_2 a^2 + \gamma_3 s + \gamma_4 s^2 + \gamma_5 m + \gamma_6 y + \gamma_7 y^2, \\
\theta(y, m) &= \theta_0 + \theta_1 y + \theta_2 y^2 + \theta_3 m.
\end{aligned}$$

Optimal driving:

$$\begin{aligned}
k^* &= \frac{1}{2\phi}(\theta_0 + \theta_1 y + \theta_2 y^2 + \theta_3 m)p^k(a, \tau) - \frac{1}{2\phi}(\gamma_0 + \gamma_1 a + \gamma_2 a^2 + \gamma_3 s + \gamma_4 s^2 + \gamma_5 m + \gamma_6 y + \gamma_7 y^2) \\
&= \kappa_0 + \kappa_1 a + \kappa_2 a^2 + \kappa_3 s + \kappa_4 s^2 + \kappa_5 m + \kappa_6 y + \kappa_7 y^2 + (\kappa_8 + \kappa_9 y + \kappa_{10} y^2 + \kappa_{11} m)p^{km}(a, \tau),
\end{aligned}$$

where the κ -parameters are the “reduced form” or “first stage” parameters. In the main estimation, they are kept fixed and used to predict the optimal driving, $k^*(x_{i,t})$, which is then used in the model. **Trading costs:**

$$\begin{aligned}
&T(d', d, p, m) = \\
&\left\{ \begin{array}{ll} 0 & \text{if } d' = (\tau, a, 0) \text{ or } d, d' = (\emptyset, \emptyset) \\ P(\tau', a', p, m) - P(\tau, a, p, m) + c_T(\tau', a', p, m) & \text{if } d' = (\tau', a', 1) \text{ and } d = (\tau, a) \\ P(\tau', a', p, m) - \underline{P}(\tau, p, m) + c_T(\tau', a', p, m) & \text{if } d' = (\tau', a', -1) \text{ and } d = (\tau, a) \\ -P(\tau, a, p, m) & \text{if } d' = (\emptyset, \emptyset, 1) \text{ and } d = (\tau, a) \\ -\underline{P}(\tau, p, m) & \text{if } d' = (\emptyset, \emptyset, -1) \text{ and } d = (\tau, a) \\ P(\tau', a', p, m) + c_T(\tau', a', p, m) & \text{if } d' = (\tau', a') \neq (\emptyset, \emptyset) \text{ and } d = (\emptyset, \emptyset) \end{array} \right.
\end{aligned}$$

Transactions costs:

$$c_T(\tau', a', p, m) = P(\tau', a', p, m)b_1(\tau', a', p, m) + b_2(\tau', a', p, m) \quad (\text{F.1})$$

References

- ADDA, J. AND R. COOPER (2000a): “Balladurette and Juppette: A Discrete Analysis of Scrapping Subsidies,” *Journal of Political Economy*, 108(4), 778–806.
- (2000b): “The Dynamics of Car Sales: A Discrete Choice Approach,” *NBER working paper 7785*.
- AKERLOF, G. A. (1970): “The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism,” *Quarterly Journal of Economics*, 84, 488–500.
- ALLCOTT, H. AND N. WOZNY (2012): “Gasoline Prices, Fuel Economy, and the Energy Paradox,” *NBER Working Paper 18583*, Cambridge, MA.
- ANDERSON, S. AND V. GINSBURGH (1994): “Price Discrimination via Second-hand Markets,” *European Economic Review*, 38(1), 23–44.
- ANDERSON, S. T., R. KELLOGG, J. M. SALLEE AND R. T. CURTIN (2011): “Forecasting gasoline prices using consumer surveys,” *American Economic Review Papers and Proceedings*, 101(3), 110.
- BENTO, A., L. GOULDER, M. JACOBSEN AND R. VON HAEFEN (2009): “Distributional and Efficiency Impacts of Increased US Gasoline Taxes,” *American Economic Review*, 99(3), 667–699.
- BERKOVEC, J. (1985): “New Car Sales and Used Car Stocks: A Model of the Automobile Market,” *RAND Journal of Economics*, 16(2), 195–214.
- BERRY, S., J. LEVINSOHN AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63(4), 841–890.
- BOND, E. (1982): “A Direct Test of the ‘Lemons’ Model: The Market for Used Pickup Trucks,” *American Economic Review*, 72(4), 836–840.
- BRESNAHAN, T. (1981): “Departures from Marginal Cost Pricing in the American Automobile Industry: Estimates for 1977–1978,” *Journal of Econometrics*, 17(2), 201–227.
- BUSSE, M. R., C. R. KNITTEL AND F. ZETTELMEYER (2013): “Are Consumers Myopic? Evidence from New and Used Car Purchases,” *The American Economic Review*, 103(1), 220–256.
- CHEN, J., S. ESTEBAN AND M. SHUM (2013): “When Do Secondary Markets Harm Firms?,” *American Economic Review*, 103(7), 2911–2934.
- CHO, S. AND J. RUST (2010): “The Flat Rental Puzzle,” *Review of Economic Studies*, 77(2), 560–594.

- DUBIN, J. AND D. MCFADDEN (1984): “An Econometric Analysis of Residential Electric Appliance Holdings and Consumption,” *Econometrica*, 52(2), 345–362.
- ENGERS, M., M. HARTMANN AND S. STERN (2008): “Are Lemons Really Hot Potatoes?,” *manuscript, University of Virginia*.
- (2009): “Annual Miles Drive Used Car Prices,” *Journal of Applied Econometrics*, 24(1), 1–33.
- ESTEBAN, S. AND M. SHUM (2007): “Durable-goods Oligopoly with Secondary Markets: The Case of Automobiles,” *RAND Journal of Economics*, 38(2), 332–354.
- FENG, Y., D. FULLERTON AND L. GAN (2005): “Vehicle Choices, Miles Driven, and Pollution Policies,” *NBER Working Paper 11553*, Cambridge, MA.
- GAVAZZA, A., A. LIZZERI AND N. ROKETSKIY (2014): “A Quantitative Analysis of the Used-Car Market,” *American Economic Review*, 104(11), 3668–3700.
- GILLINGHAM, K. (2012): “Selection on Anticipated Driving and the Consumer Response to Changing Gasoline Prices,” *Yale University, Working Paper*.
- GILLINGHAM, K. AND A. MUNK-NIELSEN (2015): “The Tail Wagging the Dog: Commuting and the Fuel Price Response in Driving,” *Working Paper*.
- GOLDBERG, P. (1995): “Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry,” *Econometrica*, 63, 891–951.
- (1998): “The Effects of the Corporate Average Fuel Efficiency Standards in the US,” *The Journal of Industrial Economics*, 46(1), 1–33.
- HENDEL, I. AND A. LIZZERI (1999a): “Adverse Selection in Durable Goods Markets,” *American Economic Review*, 89(5), 1097–1115.
- (1999b): “Interfering with Secondary Markets,” *RAND Journal of Economics*, 30(1), 1–21.
- HYMEL, K. M. AND K. A. SMALL (2015): “The Rebound Effect for Automobile Travel: Asymmetric Response to Price Changes and Novel Features of the 2000s,” *Energy Economics*, forthcoming.
- ISKHAKOV, F., T. JORGENSEN, J. RUST AND B. SCHJERNING (2015): “Estimating Discrete-Continuous Choice Models: The Endogenous Grid Method with Taste Shocks,” *manuscript*.

- JACOBSEN, M. (2013): “Evaluating U.S. Fuel Economy Standards in a Model with Producer and Household Heterogeneity,” *American Economic Journal: Economic Policy*, 5(2), 148–187.
- KONISHI, H. AND M. SANDFORT (2002): “Existence of Stationary Equilibrium in the Markets for New and Used Durable Goods,” *Journal of Economic Dynamics and Control*, 26(6), 1029–1052.
- KRUSELL, P. AND T. SMITH (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106(5), 867–896.
- MANNERING, F. AND C. WINSTON (1985): “A Dynamic Empirical Analysis of Household Vehicle Ownership and Utilization,” *RAND Journal of Economics*, 16(2), 215–236.
- MANSKI, C. (1980): “Short Run Equilibrium in the Automobile Market,” *Falk Institute Discussion Paper 8018*, Hebrew University of Jerusalem.
- (1983): “Analysis of Equilibrium Automobile Holdings in Israel with Aggregate Discrete Choice Models,” *Transportation Research*, 17B (5), 373–389.
- MANSKI, C. AND E. GOLDIN (1983): “An Econometric Analysis of Vehicle Scrappage,” *Transportation Science*, 17 (4), 365–375.
- MANSKI, C. AND L. SHERMAN (1980): “Forecasting Equilibrium Motor Vehicle Holdings by Means of Disaggregate Models,” *Transportation Research Record*, 764, 96–103.
- MUNK-NIELSEN, A. (2015): “Diesel Cars and Environmental Policy,” *University of Copenhagen Working Paper*.
- PETRIN, A. (2002): “Quantifying the Benefits of New Products: The Case of the Minivan,” *Journal of Political Economy*, 110(4), 705–729.
- REICH, G. (2013): “Divide and Conquer: A New Approach to Dynamic Discrete Choice with Serial Correlation,” *SSRN Working Paper 23715922*, Cambridge, MA.
- RUST, J. (1985a): “Equilibrium holdings distributions in durable asset markets,” *Transportation Research B*, 19(4), 331–345.
- (1985b): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55(5), 999–1033.
- (1985c): “Stationary Equilibrium in a Market for Durable Assets,” *Econometrica*, 53(4), 783–806.
- (1985d): “When is it Optimal to Kill Off the Market for Used Durable Goods?,” *Econometrica*, 54(1), 65–86.

Table F.2: Summary Statistics

Variable	N	mean	sd	p1	p50	p99
Age of H.	22041601.00	38.93	11.66	19.00	38.00	60.00
Real income (2005 kr)	22041601.00	403820.70	403550.81	29303.24	325821.59	1410501.00
Urban resident	22041601.00	0.32	0.47	0.00	0.00	1.00
Work distance of H.	22041601.00	20.81	86.96	0.00	24.00	104.89
Unemployment for H.	16242835.00	0.08	0.28	0.00	0.00	1.00
Dummy for couple	22041601.00	0.45	0.50	0.00	0.00	1.00
H. Work place shut down	9221767.00	0.03	0.18	0.00	0.00	1.00
Num of kids	22041601.00	0.61	0.97	0.00	0.00	4.00
Car age in years	7085310.00	7.27	4.88	0.00	7.00	20.00
Fuel price (period)	6362373.00	8.76	0.63	7.04	8.87	9.63
Fuel price (annual)	22041601.00	8.37	1.32	6.42	8.21	10.50
Dummy for diesel car	22041601.00	0.02	0.14	0.00	0.00	1.00
Total weight of car	7085310.00	1576.66	217.83	1125.00	1575.00	2100.00
Fuel efficiency (km/l)	3547818.00	14.21	2.49	9.45	13.82	22.70
VKT (km traveled/day)	7085310.00	46.52	24.22	4.13	43.53	124.39
Years to test	7085310.00	4.03	3.65	1.36	2.28	16.66

Notes: “H.” refers to the head of the household. All Danish kroner (kr) in 2005 kroner.

SCHIRALDI, P. (2011): “Automobile Replacement: A Dynamic Structural Approach,” *RAND Journal of Economics*, 42(2), 266–291.

SMALL, K. AND K. VAN DENDER (2007): “Fuel Efficiency and Motor Vehicle Travel: The Declining Rebound Effect,” *Energy Journal*, 28(1), 25–51.

STOLYAROV, D. (2002): “Turnover of Used Durables in a Stationary Equilibrium: Are Older Goods Traded More?,” *Journal of Political Economy*, 110(6), 1390–1413.

WEST, S. (2004): “Distributional Effects of Alternative Vehicle Pollution Control Technologies,” *Journal of Public Economics*, 88, 735.