

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

Kenneth Gillingham, *Yale University*[†]
Fedor Iskhakov, *Australian National University*[‡]
Anders Munk-Nielsen, *University of Copenhagen*[§]
John Rust, *Georgetown University*[¶]
Bertel Schjerning, *University of Copenhagen*^{||}

Preliminary Draft

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 scrapping 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 scrapping 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

[‡]Australian National University College of Business & Economics, 1018 HW Arndt Building, Canberra, ACT 0200, Australia, fedor.iskhakov@anu.edu.au

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

Contents

1	Introduction	1
2	Background and Data	3
2.1	Institutional Setting	4
2.2	Data	5
2.3	Descriptives	6
3	The Model	13
3.1	Household Dynamic Vehicle Choice Problem	19
3.2	Utility Specification	28
3.3	Specification of the Transition Densities	29
4	Estimation of the Model	30
5	Solving for Equilibrium Prices	36
5.1	Solving for Equilibrium Prices	36
5.2	Simulating Forward in Time	38
5.3	Non-Stationary Expectations	38
6	Results	39
6.1	Implementation	39
6.2	First-Stage Results	40
6.3	Structural Estimates	42
6.4	Equilibrium Prices: In-Sample	47
6.5	Counterfactual Simulations	51
7	Conclusion	53
A	Appendix: Income Transitions	55
B	Appendix: Background and Data	57
B.1	Institutional Background	57
B.2	Additional Descriptives	57
B.3	Scrappage	61
B.4	Driving	66
C	Appendix: Flexible Price Function Specification	67
C.1	Derivatives of the price function with respect to (θ, α)	71
C.2	Non-monotonic specification	72
D	Appendix 5: Additional Results	73
D.1	Estimates with Fixed Transaction Costs	73
D.2	Equilibrium Prices	75
D.3	Counterfactual Simulations	75

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, 1985b; 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 (1985b) 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. 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.

¹ Examples 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 issue.

One shortcoming of the vehicle data is that we do not observe the model year of the vehicle.

² Exceptions 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.

³ Note 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.

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 22). Therefore, importing of cars is not a major issue in our empirical setting.

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 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 1 shows the car age distribution, this is not necessarily informative about how

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

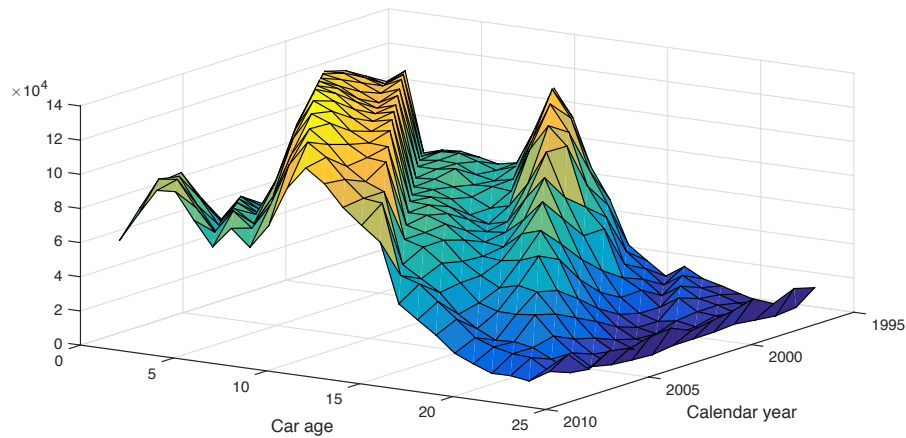


Figure 1: Car Age Distribution Over Time: “Waves”

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 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 clearer if we remove the new car sales from Figure 2 (see Appendix Figure 23).

Our Appendix Table 10 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 3 shows the new 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.⁵

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

⁵Appendix Figure 19 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 20 and 21 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).




Figure 2: Purchases by Car Age Over Time


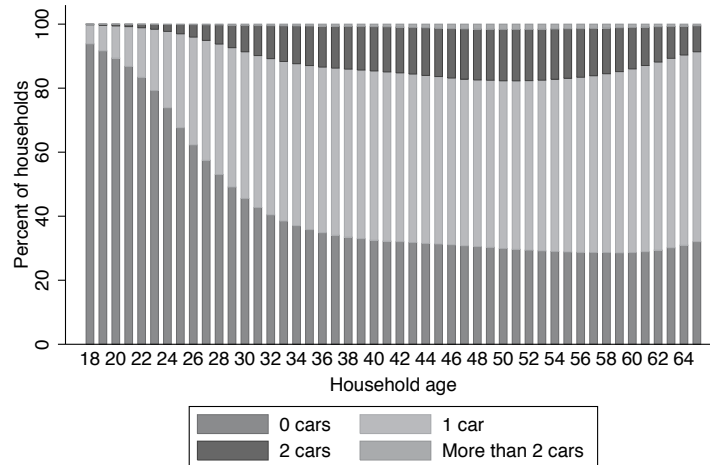


Figure 3: Car Characteristics over Time



Note: each bar shows the percentage of households within the given age category that own the particular number of cars in a particular year.

Figure 4: Number of Cars Owned by Household Age

age and income. First, note that only 12.1% of the households in our sample owns more than 1 car (Table 6). This is very low compared to the US but makes sense in light of the very high car prices (see Appendix Figure 22). 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 well-known 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 state. The two observations will split the household income equally to ensure that the total amount of resources in the economy remains stable.⁶

Figure 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 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,

⁶ If 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.




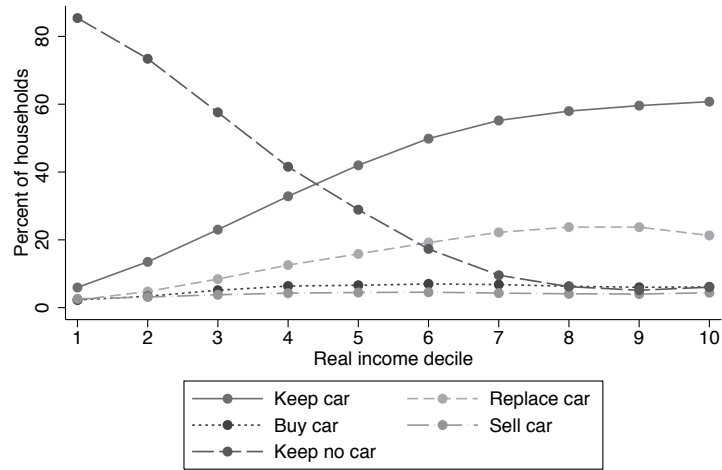
Figure 5: Number of Cars Owned by Income Decile

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 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 4 and 5 alone to determine whether the most important drivers are related to household age (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 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 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



Note: each curve shows the percentage of households within the given income decile that is in the particular car state in a given year.

Figure 6: Discrete Choice by Income Decile




Figure 7: Years of Ownership by Car Age at Purchase




Figure 8: Scrappage by Vehicle Age

figure shows that the mode of car scrappage occurs at car 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 9 shows the number of cars being scrapped in each year by the car age. When compared to the waves in Figure 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 9 is that the age distribution changes; so while Table 7 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 7). 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. In the Appendix we also include Figures showing correlations between VKT and several of our state variables.




Figure 9: Scrappage by Year and Car Age

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

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

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 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

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 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.

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 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 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

⁹In the estimation, we will set $\bar{a}(\tau) = 24$ for $\tau = 1, 2$ corresponding to gasoline and diesel.

scrapped as shown in Figure 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 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

$$D(d) = \left\{ (\tau, a, 0), \{(\emptyset, \emptyset, d_s), d_s \in \{-1, 1\}\}, \{(\tau, a, d_s), \tau \in \{1, \dots, \bar{\tau}\}, a \in \{0, \dots, \bar{a} - 1\}, d_s \in \{-1, 1\}\} \right\} \quad (1)$$

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 $d' = (\tau', a')$. Notice the decision $d' = (\tau', a', 1)$ corresponds to the household's decision to trade-in their current car $d = (\tau, a)$ for another car (τ', a') , and receive the secondary market price $P(\tau, a, p, m)$ for their existing car as a trade-in value. On the other hand, the household might also choose to scrap their current car if its physical condition is sufficiently poor and receive only the scrap value \underline{P} for it. This decision is denoted by $d' = (\tau', a', -1)$ where $d_s = -1$ corresponds to the decision to scrap the current car $d = (\tau, a)$.

Our model also accounts for accidents that are sufficiently severe to result in a total loss of the vehicle. Some accidents may be more minor and the damage is repairable, but here we are referring to accidents that are so severe the car is declared to be a “total loss” (typically by the household’s insurance company) and the household is unable to repair the damage and is simply given an insurance settlement for the expected market price (often less some deductible or coinsurance amount) of the vehicle involved in the accident. When such accidents happen, the household does not have the option of keeping the existing vehicle so the choice set is the one given in equation (2) but without the alternative $(\tau, a, 0)$ which corresponds to choosing to keep the current vehicle $d = (\tau, a)$. We denote this restricted choice set where keeping the currently owned car is not a feasible option by $D_R(d)$.

The choice set for a household that does not have a car $d = (\emptyset, \emptyset)$ is

$$D(d) = \left\{ (\emptyset, \emptyset), \{(\tau, a), \tau = 1, \dots, \bar{\tau}, a = 0, \dots, \bar{a}\} \right\} \quad (2)$$

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 cost* and *taxes* associated with the purchase of a new car. Moreover, 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

$$T(d', d, p, m) = \begin{cases} 0 & \text{if } d' = (\tau, a, 0) \text{ or } d, d' = (\emptyset, \emptyset) \\ P(\tau', a', p, m) - P(\tau, a, p, m) + c_T(\tau', a', p, m) & \text{if } d' = (\tau', a', 1) \text{ and } d = (\tau, a) \\ P(\tau', a', p, m) - \underline{P}(\tau, p, m) + c_T(\tau', a', p, m) & \text{if } d' = (\tau', a', -1) \text{ and } d = (\tau, a) \\ -P(\tau, a, p, m) & \text{if } d' = (\emptyset, \emptyset, 1) \text{ and } d = (\tau, a) \\ -\underline{P}(\tau, p, m) & \text{if } d' = (\emptyset, \emptyset, -1) \text{ and } d = (\tau, a) \\ P(\tau', a', p, m) + c_T(\tau', a', p, m) & \text{if } d' = (\tau', a') \neq (\emptyset, \emptyset) \text{ and } d = (\emptyset, \emptyset) \end{cases} \quad (3)$$

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

$$c_T(\tau', a', p, m) = P(\tau', a', p, m)b_1(\tau', a', p, m) + b_2(\tau', a', p, m) \quad (4)$$

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 $\varepsilon(d')$ representing the impact of idiosyncratic unobserved factors that affect the consumer’s choice, so the total current period utility becomes $v_s(d', d, p, m, x) + \varepsilon(d')$. Let $\varepsilon = \{\varepsilon(d') | d' \in D(d)\}$ 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 $V_s(d, p, m, x, \varepsilon)$ 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 as 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 $\varepsilon(d)$ and $\varepsilon(d')$ for $d \neq d'$ as well as serial independence in the overall vector stochastic process $\{\varepsilon_t\}$. 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 $\varepsilon_t(d)$ and $\varepsilon_t(d')$ 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 relevant for future specifications 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

$$V_s(d, p, m, x, \varepsilon) = \max_{d' \in D(d)} [v_s(d', d, p, m, x) + \varepsilon(d') + \beta EV_s(d', d, p, m, x, \varepsilon)] \quad (5)$$

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 $(\tilde{d}, \tilde{p}, \tilde{m}, \tilde{x}, \tilde{\varepsilon})$ 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 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

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.

$EV_s(d', p, m, x)$. This implies that we can write the Bellman equation as

$$V_s(d, p, m, x, \varepsilon) = \max_{d' \in D(d)} [v_s(d', d, p, m, x) + \varepsilon(d') + \beta EV_s(d', p, m, x)]. \quad (6)$$

Let $V_s(d', d, p, m, x)$ denote the *choice-specific value function*

$$V_s(d', d, p, m, x) = v_s(d', d, p, m, x) + \beta EV_s(d', p, m, x). \quad (7)$$

Then following Rust (1985a) we can rewrite the Bellman equation (5) in terms of the choice-specific value functions (7) as

$$V_s(d, p, m, x, \varepsilon) = \max_{d' \in D(d)} [V_s(d', d, p, m, x) + \varepsilon(d')]. \quad (8)$$

Equation (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

$$F(\varepsilon(\tau', a', -1), \varepsilon(\tau', a', 1)) = \exp \left\{ - \left[\exp\{-\varepsilon(\tau', a', -1)/\lambda\} + \exp\{-\varepsilon(\tau', a', 1)/\lambda\} \right]^\lambda \right\} \quad (9)$$

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

$$P(\tau, a, p, m) + \varepsilon(\tau', a', 1) \geq \underline{P}(\tau, p, m) + \varepsilon(\tau', a', -1). \quad (10)$$

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

$$\Pi_s \{d_s = 1 | d, d', p, m, x\} = \frac{\exp \{ \underline{P}(\tau, a, p, m) / \lambda \}}{\exp \{ P(\tau, a, p, m) / \lambda \} + \exp \{ \underline{P}(\tau, a, p, m) / \lambda \}}. \quad (11)$$

Note that the scrappage probability can be calculated independently of the overall solution of the dynamic programming problem given in equation (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

$$\begin{aligned} & \max [v_s((d', -1), d, p, m, x) + \varepsilon(d', -1), v_s((d', 1), d, p, m, x) + \varepsilon(d', 1)] = \\ & \lambda \log (\exp \{v_s((d', -1), d, p, m, x) / \lambda\} + \exp \{v_s((d', 1), d, p, m, x) / \lambda\}) + \varepsilon(d'). \end{aligned} \quad (12)$$

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 (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 (12)) and a single error term $\varepsilon(d')$ that has as Type 3 Extreme value distribution with scale parameter

$\lambda = 1$.

Using equation (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

$$v_s(d', d, p, m, x) = \lambda \log \left(\exp\{v_s((d', -1), d, p, m, x)/\lambda\} + \exp\{v_s((d', 1), d, p, m, x)/\lambda\} \right) + \varepsilon(d'). \quad (13)$$

Then with this redefinition/reduction, the Bellman equation (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

$$\Pi(d'|d, p, m, x) = \frac{\exp\{V_s(d', d, p, m, x)\}}{\sum_{d'' \in D(d)} \exp\{V_s(d'', d, p, m, x)\}}. \quad (14)$$

where $V_s(d', d, p, m, x)$ is the choice-specific value function (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 (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 (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:

$$\begin{aligned} \int_{\varepsilon'} V_{s+1}(f(d'), p', m', x', \varepsilon') q(d\varepsilon') &= \int_{\varepsilon'} \max_{d'' \in D(f(d'))} [V_{s+1}(d'', f(d'), p', m', x') + \varepsilon'(d'')] q(d\varepsilon') \\ &= \log \left(\sum_{d'' \in D(f(d'))} \exp \{V_{s+1}(d'', f(d'), p', m', x')\} \right) \\ &\equiv \varphi(f(d'), m', p', x'). \end{aligned} \quad (15)$$

Following Rust (1987) we can write the following recursion equation for the choice-specific value functions

$$\begin{aligned} V_s(d', d, m, p, x) &= v_s(d', d, m, p, x) + \\ &\quad \beta \sum_{m'} \int_{p'} \int_{x'} \varphi(f(d'), m', p, x') g(x'|x, m', p', m, p) h(p', m'|m, p) dx' dp' \end{aligned} \quad (16)$$

where $f(d')$ is given by

$$f(d') = \begin{cases} (\mathbf{0}, \mathbf{0}) & \text{if } d' = (\mathbf{0}, \mathbf{0}) \text{ or } d' = (\mathbf{0}, \mathbf{0}, d_s), d_s \in \{-1, 1\} \\ (\tau', \min[\bar{a}, a' + 1]) & \text{if } d' = (\tau', a') \text{ or } d' = (\tau', a', d_s), d_s \in \{-1, 0, 1\}. \end{cases} \quad (17)$$

As mentioned earlier, the continuation value (i.e., the expected discounted value of future utility, given by the expression multiplied by β in equation (17) above) depends only on d' and not on d . This is what equation (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 (8) and (17) we see that

$$EV_{s+1}(d', p, m, x) = \sum_{m'} \int_{p'} \int_{x'} \varphi(f(d'), m', p, x') g(x'|x, m', p', m, p) h(p', m'|m, p) dx' dp' \quad (18)$$

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 (18) above needs to be modified as follows

$$\begin{aligned}
EV_{s+1}(d', p, m, x) = & \quad (19) \\
& (1 - \alpha_s(d', x)) \sum_{m'} \int_{p'} \int_{x'} \varphi(f(d'), m', p, x') g(x'|x, m', p', m, p) h(p', m'|m, p) dx' dp' \\
& + \alpha_s(d', x) \sum_{m'} \int_{p'} \int_{x'} \varphi_R(f(d'), m', p, x') g(x'|x, m', p', m, p) h(p', m'|m, p) dx' dp'.
\end{aligned}$$

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:

$$D_R(d) = \{ \{(\emptyset, \emptyset, -1)\}, \{(\tau, a, -1)\}, \tau \in \{1, \dots, \bar{\tau}\}, a \in \{0, \dots, \bar{a} - 1\} \} \quad (20)$$

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 (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) 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

$$T_R(d', d, p, m) = \begin{cases} -P(\tau, a, m, p)(1 - \psi) & \text{if } d' = (\emptyset, \emptyset) \\ [P(\tau', a', p, m) - P(\tau, a, p, m)(1 - \psi) + c_T(\tau', a', p, m)] & \text{if } d' = (\tau', a', -1) \text{ and } d = (\tau, a) \end{cases} \quad (21)$$

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 utility* a household expects from owning a vehicle of type τ and driving a planned k kilometers, given by

$$u(k, \tau, a, s, p, m) = \theta(y, m)[y - p^k(\tau, a, p, c^o)k - T] + \gamma(y, s, a, m)k + \phi k^2 - q(a) + \delta_n \mathbb{1}(a = 0) + \delta_\tau \quad (22)$$

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 scrapping. 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 its age and type,

¹¹In the estimation, the function is monotone everywhere and predicts only strictly positive driving.

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 ,

$$\theta(y, m) = \theta_0 + \theta_1 y + \theta_2 y^2 + \theta_3 m.$$

Inserting $\gamma(y, s, a, m)$ and $\theta(y, m)$, in the equation for the optimal k , we obtain the following linear equation

$$k^* = \frac{1}{2\phi} (\theta_0 + \theta_1 y + \theta_2 y^2 + \theta_3 m) p^k(a, \tau) - \frac{1}{2\phi} (\gamma_0 + \gamma_1 a + \gamma_2 a^2 + \gamma_3 s + \gamma_4 s^2 + \gamma_5 m + \gamma_6 y + \gamma_7 y^2) \quad (23)$$

$$= \kappa_0 + \kappa_1 a + \kappa_2 a^2 + \kappa_3 s + \kappa_4 s^2 + \kappa_5 m + \kappa_6 y + \kappa_7 y^2 + (\kappa_8 + \kappa_9 y + \kappa_{10} y^2 + \kappa_{11} m) p^{km}(a, \tau), \quad (24)$$

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,

$$\log y_{it} \sim N(\mu_y, \sigma_y^2). \quad (25)$$

where μ_y is given by

$$\mu_y = \rho_1 \log y_{it-1} + \rho_2 s_{it} + \rho_3 s_{it}^2 + \rho_4 m_t + \rho_5 m_{t-1} + \rho_6 \mathbb{1}_{\{m_t=1 \wedge m_{t-1}=0\}} + \rho_7 \mathbb{1}_{\{m_t=0 \wedge m_{t-1}=1\}} \quad (26)$$

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

$$\log p_t \sim N(\log p_{t-1}, \sigma_p^2). \quad (27)$$

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

¹² To 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.

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

$$\underline{P}(\tau, p, m) = \zeta_{\tau}^{\bar{a}} P_0(\tau),$$

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 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

¹³ The scrappage subsidy paid out by the Danish Ministry of the Environment equals 1,500 DKK.

¹⁴ On 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.

¹⁵ Alternatively, one could use income data for the previous year to make sure that car decisions are made conditional on income already earned.

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 (1985a) and estimate the transition densities separately in a first stage.

We also include data on driving in our estimation. Recall that we observe odometer readings from biennial vehicle inspections in Denmark, so we need to derive the probability distribution for vehicle kilometers travelled between successive vehicle inspections. We assume that actual vehicle kilometers travelled in a year equals the desired or planned kilometers travelled from our model, which we denote by $k^*(x_{i,t}, \vartheta)$ plus an *i.i.d.* $N(0, \sigma_k^2)$ “optimization error” that captures the difference between intended or planned kilometers of travel at the start of the year and the actual kilometers driven at the end of the year. Let $k_{i,t}$ denote the *change* in the odometer reading on a car owned by household i at an inspection at date t :

$$k_{i,t} = o_{i,t} - o_{i,t-1} \tag{28}$$

where $o_{i,t}$ and $o_{i,t-1}$ are the odometer readings on the car at the inspections at dates t and $t-1$, respectively. Then by our assumption above, $k_{i,t}$ is a sum of two normally distributed random variables and its probability density, conditional on the household characteristics $(x_{i,t}, x_{i,t-1})$ that determine planned driving of the car in each of the intervening years $t-1$ and t between the biennial vehicle inspections, is $f(k_{i,t}|x_{i,t}, x_{i,t-1}, \vartheta)$ given by the convolution formula

$$\begin{aligned} f(k_{i,t}|x_{i,t}, x_{i,t-1}, \vartheta) &= \int_{-\infty}^{+\infty} \phi\left(\frac{k - k^*(x_{i,t}, \vartheta)}{\sigma_k}\right) \phi\left(\frac{k - k^*(x_{i,t-1}, \vartheta)}{\sigma_k}\right) dk \\ &= \phi\left(\frac{k_{i,t} - k^*(x_{i,t}, \vartheta) - k^*(x_{i,t-1}, \vartheta)}{\sqrt{2}\sigma_k}\right). \end{aligned} \tag{29}$$

This formula can be adjusted in a straightforward way if the car changes owner between successive inspection dates: if it is owned by household j at inspection date $t-1$ and by household i at inspection date t , then we use the formula above, but with $x_{j,t-1}$ substituted in place of $x_{i,t-1}$. If we adopt this convention, then we can, for the purposes of this part of the likelihood, let i denote a particular *vehicle* that we can potentially follow from household to household through a succession of inspection dates over the entire period of our sample. The ability to track individual cars across successive owners is possible due to the fact that Danish Register Data covers the

entire Danish population. However even for a subsample of the entire population (which creates the possibility that a particular car might “drop out” of our sample for a period where it is owned by a household that does not happen to appear in our household sample), we let I_i be the set of *successive* inspection dates that we observe a particular vehicle i that it is owned by one or more households in our sample. Then we let $L_k(\vartheta)$ be the part of the log-likelihood that captured this vehicle kilometer travelled information and it is given by

$$L_k(\vartheta) = \sum_{i=1}^{N_k} \sum_{t \in I_i} \log \left(f(k_{i,t} | x_{h_{i,t},t}, x_{h_{i,t-1},t-1}, \vartheta) \right), \quad (30)$$

where $h_{i,t}$ indexes the household in the sample that owns the car at inspection date t . Notice the outer sum in L_k is over the N_k cars that are owned by one or more households in our sample, and it is further restricted to the subsample of cars that we observe at a sequence of *successive* inspection dates, reflecting the fact that we cannot predict desired kilometers of travel $k^*(x_{h_{i,t},t})$ if the household $h_{i,t}$ that owns the car at inspection date t is not in our sample.

Now consider the component of the log-likelihood for the information we have on households’ ownership and trading decisions. For this component of the likelihood we will assume that we have a total of N households in our sample that we observe for a succession of years denoted by \mathcal{T}_i that may differ from household to household due to household dissolution and formation that differs across households over the period of our sample. The log-likelihood for this component depends on the sequence of household choices which are denoted by $\{d_{i,t}\}$ which includes the decision not to own a car. However, for those households that do own a car, we observe those who decide not to keep their current car, leading the household to sell or scrap. In the Danish Register Data we have unique identifiers for cars which enables us to track when a car is sold by one household and subsequently purchased by another household. When a car is scrapped, it is not purchased by any other household and so the car’s identification number disappears from the registry, and event that is referred to as a *deregistration*.

However cars can be deregistered for a number of other reasons besides a scrappage. Cars can be temporarily deregistered if they are traded in to a used car dealership for a short period until the car is sold to another household. Cars can also be deregistered if they are exported abroad, or if the owner of the car decides to simply “park” their vehicle and not use it (since it is illegal to drive a deregistered car in Denmark). For these reasons, not all vehicle deregistrations correspond to voluntary scrappage decisions. In attempt to isolate the “temporary deregistrations” we track all cars owned by the households in our sample as far forward in time as possible. For the cars that we determine to be temporarily deregistered, we do not classify the sales of these vehicles by the household who last owned them prior to the onset of the interval of temporary deregistration as “scrappage” events. However for cars that are deregistered “permanently” (i.e., where we cannot

observe them subsequently becoming registered again over the period of our sample, which ends in 2009, and we see no subsequent re-registration of the vehicle by 2011), we assume that these cars have either been scrapped or exported.

Unfortunately the Danish Register Data do not enable us to determine the reason why a particular vehicle is deregistered, and this is why it is somewhat difficult for us to determine if a particular car is only temporarily deregistered or permanently deregistered, such as would occur if the car were scrapped. In our analysis we do not distinguish between scrapping and exports for the purposes of calculating *equilibrium prices in Denmark*. In either case the car is removed from the stock of cars *in Denmark* and thus cannot contribute to the supply of vehicles any longer. We also ignore the import of used vehicles into Denmark, since the number of used car imports is negligible. Thus we treat all permanent deregistrations as equivalent to scrapping, i.e. the permanent removal of the car from Denmark.

Now consider a car that is sold by a household and we determine it to have been permanently deregistered after this sale so there are no other intervening owners between the sale date and the date of deregistration. Whether this car was scrapped or exported, it has been removed from the stock of vehicles in Denmark and thus can no longer contribute to the supply of vehicles. Unfortunately, we still have another problem in that we cannot observe whether this car was exported, scrapped “voluntarily” or scrapped “involuntarily.” In our study, we will treat exports as equivalent to scrapping, whether voluntary or involuntary. However we can econometrically handle the distinction between voluntary and involuntary scrapping by using a mixture probability for the households that sell and deregister their vehicles. We do this by estimating a flexibly parameterized accident probability $\alpha(x_{i,t})$ that can depend on the characteristics of the car and the household, $x_{i,t}$. Thus, if we observe a household who permanently deregisters their car, this occurs with probability

$$Pr\{d_{i,t}|x_{i,t}, \vartheta\} = \alpha(x_{i,t}) + [1 - \alpha(x_{i,t})]\Pi_s(P, \lambda) \quad (31)$$

where $\Pi_s(P, \lambda)$ is the scrapping probability given in equation (11) above.

If the household has a car and chooses to keep it, the probability for this observation is

$$Pr\{d_{i,t}|x_{i,t}, \vartheta\} = [1 - \alpha(x_{i,t})]\Pi(d_{i,t}|x_{i,t}, \vartheta), \quad (32)$$

where $\Pi(d_{i,t}|x_{i,t}, \vartheta)$ is a short-hand notation for the conditional choice probability given in equation (14) above. We multiply the choice probability by $1 - \alpha(x_{i,t})$ because the household could not have kept their car if it had been involved in a serious accident that forced the household into an involuntary scrapping decision. Finally, if the household does not have a car, or purchases another car, or chooses not to have a car and we do not observe a permanent deregistration of the

car following the sale of the existing vehicle, the probability for this observation is

$$\Pr\{d_{i,t}|x_{i,t}, \vartheta\} = \Pi(d_{i,t}|x_{i,t}, \vartheta), \quad (33)$$

since the accident probability is no longer relevant for these events.

Let ϑ contain all parameters jointly. The component of the log-likelihood for the information on household vehicle holding and trading decisions is given by

$$L(\vartheta) = \sum_{i=1}^N \sum_{t \in \mathcal{T}_i} \log \{\Pr(d_{i,t}|x_{i,t}; \vartheta)\}, \quad (34)$$

where the probability entering the likelihood is given by one of the appropriate formulas above, (31), (32), or (33), depending on whether we observe a household selling a car and permanently deregistering it or not, and \mathcal{T}_i denotes the subset of years where we observe household i . We 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 (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:

$$\begin{aligned} u[k^*(\kappa), \tau, a, s, p, m] &= \theta(y, m)[y - p^k(\tau, a, p, c^o)k^*(\kappa) - T] \\ &+ \gamma(y, s, a, m)k^*(\kappa) + \phi[k^*(\kappa)]^2 - q(a) + \delta_n \mathbb{1}(a = 0) + \delta_\tau. \end{aligned} \quad (35)$$

Then we use the following 2nd stage likelihood function,

$$L^{2\text{step}}(\vartheta) = \sum_{i=1}^N \sum_{t \in \mathcal{T}_i} \log \left\{ \Pr(d_{i,t}|x_{i,t}, ; \vartheta, k = k^*(\kappa)) [\Pr(d_{i,t,s}|x_{i,t}, k = k^*(\kappa))]^{\mathbb{1}\{d_{i,t,s} \neq 0\}} \right\}, \quad (36)$$

where the conditioning on $k = k^*(\kappa)$ is to indicate that the model should be solved using the flow utility given in (35). Note that we are still searching over the same parameters as when we use the full likelihood from equation (34); 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 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 (1985b) 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

$$ED(\tau, a, S, P) = \sum_{i=1}^N \Pr [d'_i = (\tau, a) | d_{i,t}, x_{i,t}; P] - \sum_{i=1}^N \{1 - \Pr [d'_i = 0 | d_{i,t}, x_{i,t}; P]\} \mathbb{1}(d_{i,t} = (\tau, a)), \quad (37)$$

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.¹⁶ The first term in equation (37) 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,

$$ED(S_t, P_t) = 0. \quad (38)$$

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 (38) using real data, so existence is proven constructively for our situation. The conditions under which equilibrium prices exist are left for future work.

When working with the non-parametric specification in (38), 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} .

¹⁶ As 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 $ED(\tau, \bar{a}, S, P) = 0$.

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 (37) 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 (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.¹⁷

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.

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

¹⁷ The 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 .

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) .¹⁸

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 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.

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.¹⁹ 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 (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)

¹⁸ Krusell 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.

¹⁹ The 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.

income process and the estimated coefficients are shown in Table 3. 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) had 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 4). 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 4). 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.²⁰ 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 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. $\underline{P}(\tau, p, m) = \zeta_\tau^{\bar{a}} P(\tau, 0, p, m)$.

We choose to use the “two stage” estimation procedure outlined in Section 4: we start by estimating the κ -parameters in the driving equation (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 (35), and find the structural parameters by maximizing (36).

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

²⁰ Recall that the most common ownership length is 5 years (Figure 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 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.

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 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 1. In Table 8, we show regressions corresponding to the first stage specification in Table 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).²¹ 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 9. Most notably, the fixed transaction 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.²² 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

²¹ Including 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.

²² We 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 .

decisions.²³ 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 10 shows the model fit in terms of the choice probabilities (observed and predicted), here shown for the 2002 cross-section.²⁴ 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 4 and 5). We 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 11 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


²³ Specifically, 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.

²⁴ We 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 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 choicset	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	

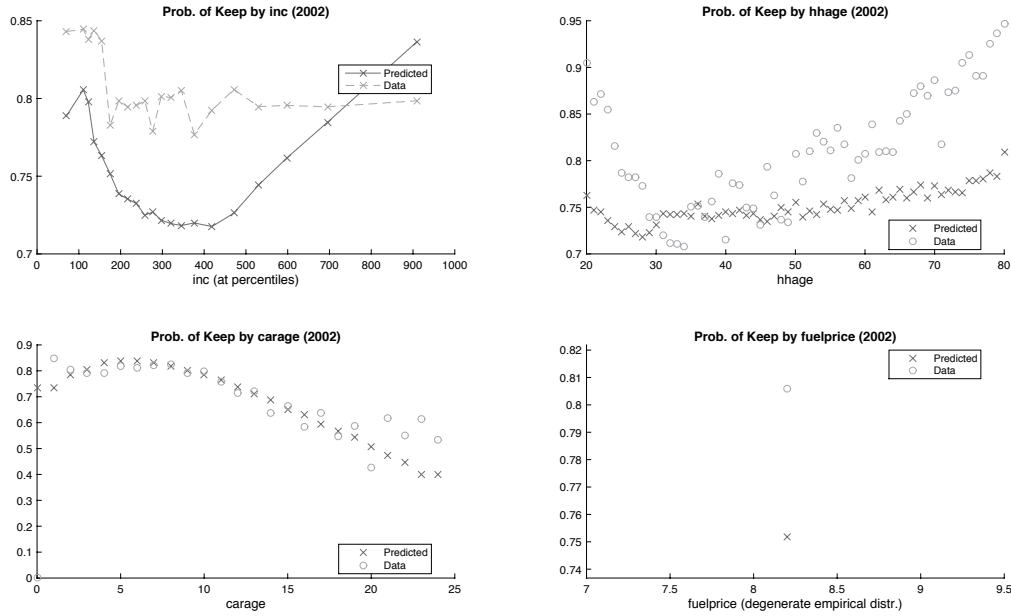
Figure 10: Model Fit: Conditional Choice Probabilities (CCPs)



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 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.²⁵

²⁵ For 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

Figure 11: Model Fit by State Variables: The Keep-decision




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 12 and the simulated macro process is shown in Figure 34.

First off, we do not see the clear macro waves in the car age distribution in 12 that we observe in the actual data (Figure 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

not to own a car.

Figure 12: Forward Simulation from the Non-equilibrium Model




probabilities by car age that was indicated in Figure 10. Finally, note that the only thing creating the small “waves” present in Figure 12 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 D.1 presents results from the model where transaction costs are kept fixed at lower values. Table 9 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 32).

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 13. 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 un-affected 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.

Figure 13: Equilibrium Prices: Gasoline Cars

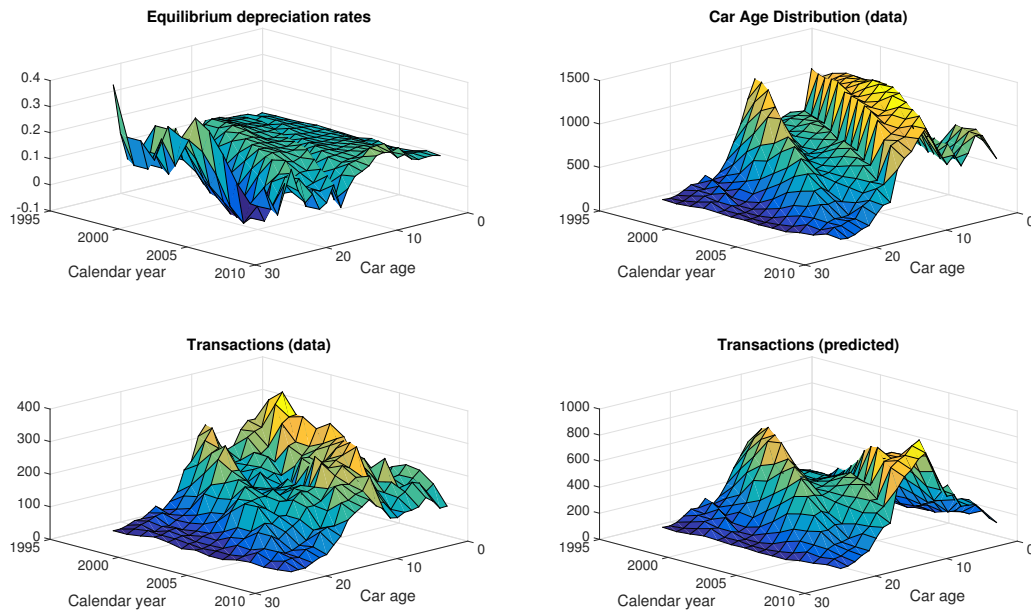


From closer inspections, we found that the model fits quite poorly in 2008 and predicts that too many car-owning households should sell their cars. The explanation may well be the spike in real fuel prices for both gasoline and diesel in 2008 (cf. Figure 19); if this causes all households to want to 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 14 shows 4 panels; the first panel shows the negative log differences of figure 13, 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.²⁶ The upper right panel shows the car

²⁶ When 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

Figure 14: Simulations Under Equilibrium Prices: Gasoline Cars



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.²⁷ 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. 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.²⁸


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 house-

higher car ages without ever converging. We consider this a problem related to the 1% subsample.

²⁷ To 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.

²⁸ In 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 15: Forward Simulation with Equilibrium Prices



hold behavior moving forward (similarly to 12). Figure 15 shows the car age distribution in this simulation. When compared with Figure 12, 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 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 1, new car sales are persistently higher throughout the booms and persistently low throughout the busts. Figure 34 shows the macro state process and the (constant) fuel prices for this simulation. The macro process is the same as was used for Figure 12. Figure 35 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 to use cars in general, and in particular fuel-inefficient ones.²⁹ 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%.³⁰ 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 16 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 12, 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 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 36 shows the simulated paths of the macro state and the exogenous fuel price process to aid the interpretation of Figure 16.

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 36 (and fuel prices are constant except for the exogenous increase of 50% in period 2012).

²⁹ Figure 28 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.

³⁰ These 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 16: Counterfactual Simulation: Non-equilibrium Prices

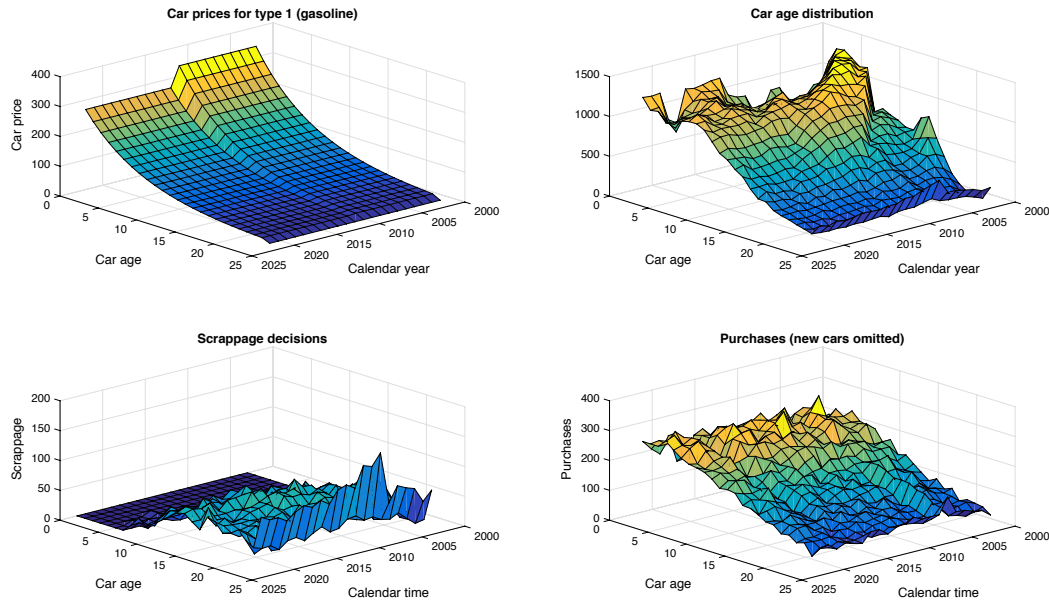
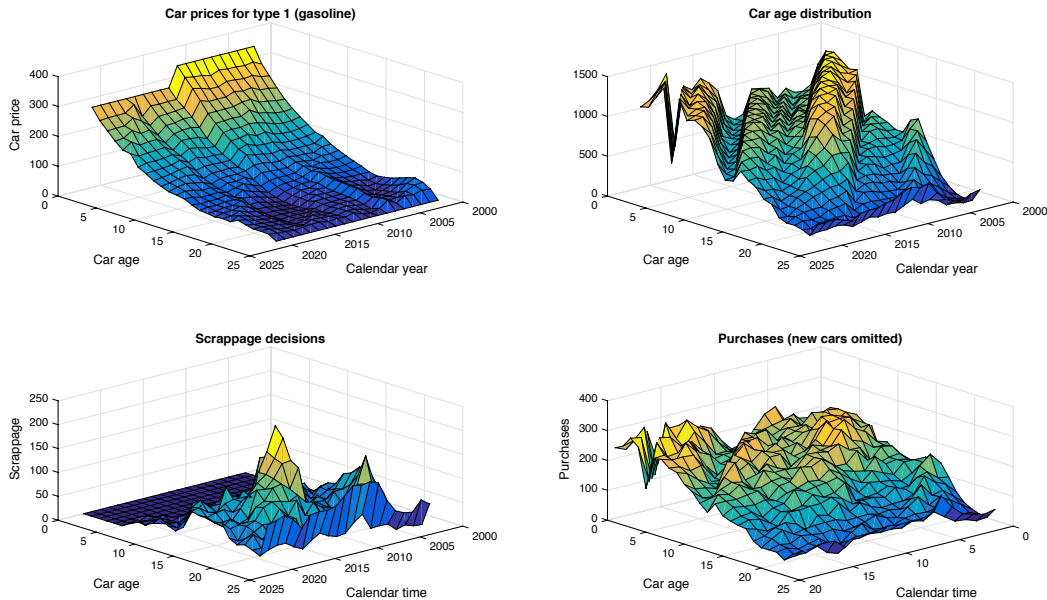


Figure 17 shows the equilibrium simulation. These simulations differ markedly from the non-equilibrium counterparts. The car age distribution displays clear waves traveling through the distribution that look very much like the waves we see in the real data. The equilibrium price schedule is shown over 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, so do 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

Figure 17: Counterfactual Simulation: Equilibrium Prices




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 18 shows a rotated view of the car age distribution in Figure 17, which makes it easier to see the new car sales replacing the old wave of cars being scrapped.

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

Figure 18: Counterfactual Simulation of the Car Stock under Equilibrium Prices



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 3 shows the results from the estimation of the equation

$$\log y_{it} = \rho_0 + \rho_1 \log y_{it-1} + \rho_2 s_{it} + \rho_3 s_{it}^2 + \rho_4 m_{it} + \text{error}_{it}. \quad (39)$$

We find that controlling for the age profile of income, the AR coefficient is $\hat{\rho}_1 = 0.853$ (Table 3). 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 (recession) 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 4 and 5. 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 3: 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 4: 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
r2	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 19 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 20 and for diesel in Figure 21. 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 22 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 23 shows the number of transactions by car age and over time. When compared to the car age distribution in Figure 1, we clearly see that the “waves” appear in both graphs. This indicates

Table 5: 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
r2	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 19: Real Fuel Prices Over Time

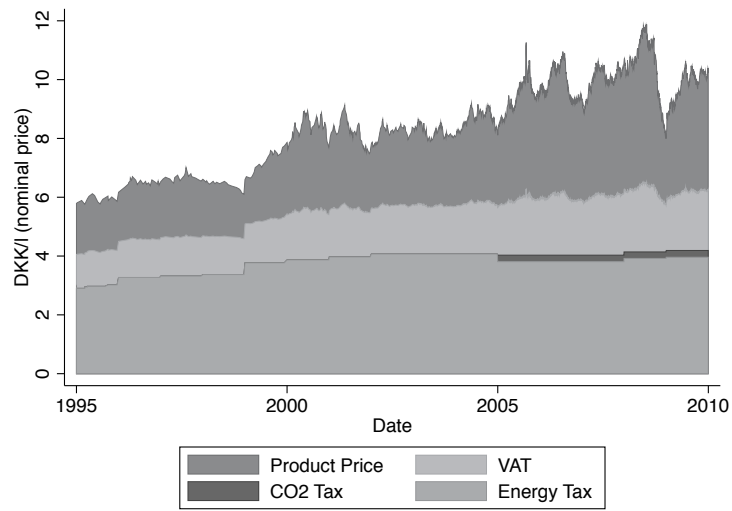


Figure 20: Composition of the Gasoline Price (Octane 95)




Figure 21: Composition of the Gasoline Price (Octane 95)




Figure 22: MRSP For a Toyota Avensis: Differences in Europe




Figure 23: Purchases by Car Age Over Time (new cars omitted)

that transactions tend to follow the age distribution.

Table 6 shows the shares of households owning zero, one, two or more than two cars for each year in our sample.

Figure 24 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 25 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 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.


Table 6: 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



Note: each bar shows the percentage of vehicles that have had the particular number of owners until they turned 15 years old in 2009. Only vehicles whose first owner is observed are used (61,919/92,021 of the vehicles). Truncated at 15 owners.

Figure 24: Number of owners for 15 year old cars in 2009



Note: each bar shows the percentage of cars and vans that were scrapped or sold to another household at each age.

Figure 25: Number of Scrappages and Number of Transactions by Car Age

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 8. We see the same spikes in scrappage in even years that correspond with the inspection years.

Table 7 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

Table 7: 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

may explain the higher number of scrappages later in our sample and perhaps particularly some of the younger scrapped cars, we see in Figure 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 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 26 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 27 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.

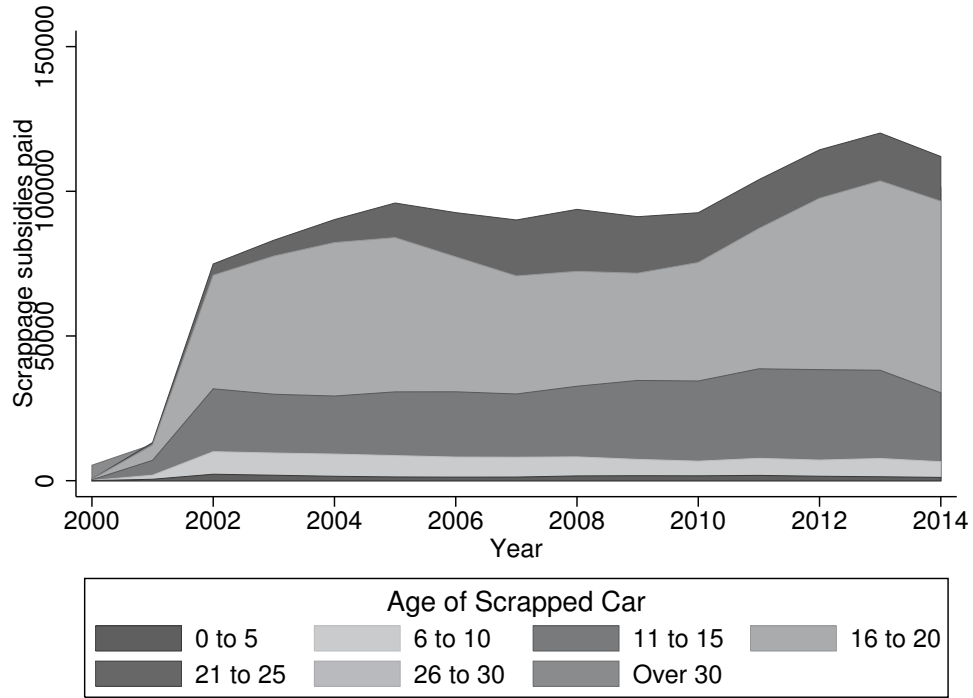


Figure 26: Scrapage Subsidies Paid by Car Age (source: bilordning.dk)




Figure 27: Number of Ownership Spells Ending




Figure 28: Vehicle Kilometers Traveled by Car Age

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. 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 28 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 29 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 29, 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 29 shows that high

Table 8 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 -179% unless we control




Figure 29: Vehicle Kilometers Traveled by Real Income

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.³¹ 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.

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

³¹ The 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.

Table 8: 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	-4.188

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.

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

$$P(a, \theta) = \underline{P} + (\bar{P} - \underline{P}) \prod_{i=1}^a \rho(\theta_i), \quad a = 1, \dots, 19 \quad (40)$$

where we define $\rho(\theta_0) \equiv 1$ and

$$\rho(\theta_i) = \frac{\exp\{\theta_i\}}{1 + \exp\{\theta_i\}} \quad (41)$$

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

$$\theta = (\theta_1, \theta_1, \dots, \theta_1) \quad (42)$$

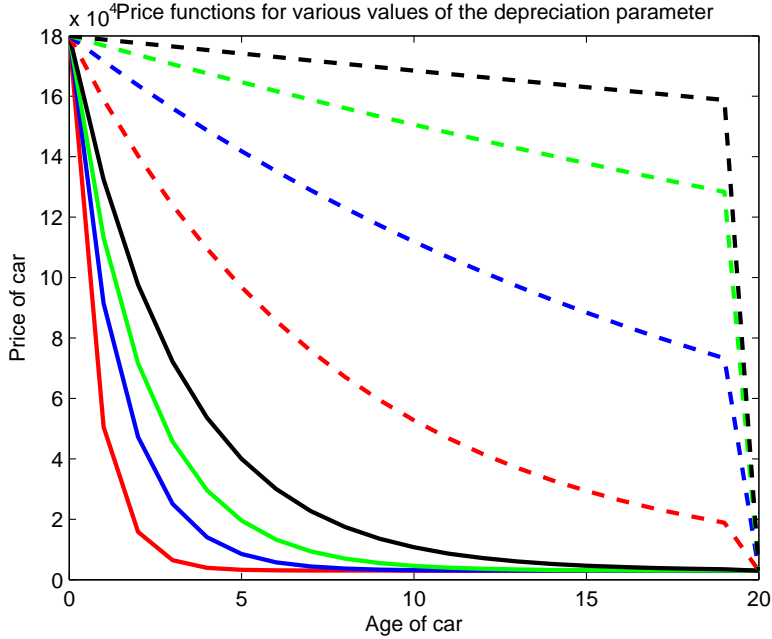
so that $\theta \in R^{19}$ depends only on a single unknown parameter $\theta_1 \in R^1$. Or we could partition θ to depend on just two parameters (θ_1, θ_2) as follows

$$\theta = (\theta_1, \theta_1, \dots, \theta_1, \theta_2, \dots, \theta_2) \quad (43)$$

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 (40) 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 $f(m, p, \alpha) = \alpha_1 + \alpha_2 m + \alpha_3 p$.

$$P(a, m, p, \theta, \alpha) = \underline{P} + (\bar{P} - \underline{P}) \prod_{i=1}^a \rho(m, p, \theta_i, \alpha), \quad a = 1, \dots, 19 \quad (44)$$



where we define $\rho(\theta_0) \equiv 1$ and

$$\rho(m, p, \theta_i) = \frac{\exp\{\theta_i + f(m, p, \alpha)\}}{1 + \exp\{\theta_i + f(m, p, \alpha)\}} \quad (45)$$

for $i = 1, \dots, 19$. Note by construction we have $P(0, m, p, \tau) = \bar{P}(m, p, \tau)$.

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

$$\frac{\partial}{\partial \theta_j} P(a, m, p, \tau) = [P(a, m, p, \tau) - \underline{P}] \left[\sum_{i=1}^a [1 - \rho(m, p, \theta_i, \alpha)] I\{i \in I_j\} \right] \quad (46)$$

$$\frac{\partial}{\partial \alpha} P(a, m, p, \tau) = [P(a, m, p, \tau) - \underline{P}] \left[\sum_{i=1}^a [1 - \rho(m, p, \theta_i, \alpha)] \frac{\partial}{\partial \alpha} f(p, m, \alpha) \right]. \quad (47)$$

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:


$$\rho_a(\theta_a) = \frac{\exp\{\theta_a\}}{1 + \exp\{\theta_a\}}, \quad (48)$$

and

$$P(\tau, a, m, p, \theta) = \sum_{i=1}^{19} I\{i = a\} [\bar{P}(\tau, m, p)\rho_a(\theta_a) + \underline{P}(\tau, m, p)(1 - \rho_a(\theta_a))]. \quad (49)$$

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 .

Figure 30: Model Fit: Conditional Choice Probabilities



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 5: Additional Results

This appendix contains additional results that have been omitted from the main results sections.

D.1 Estimates with Fixed Transaction Costs

Table 9 show 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.

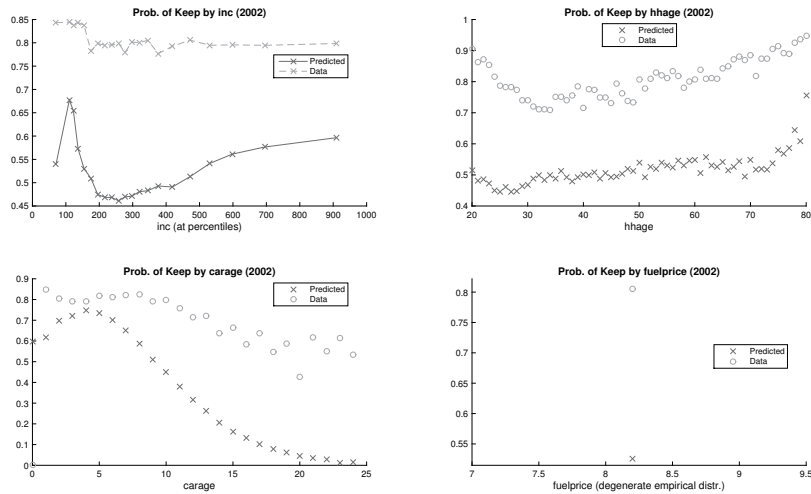
Figure 30 shows the fit in terms of conditional choice probabilities. Compared to the preferred specification where the transaction cost is estimated, we see a considerable under-prediction of the keep decision. Moreover, the model produces much more probability mass for all car ages over 4 with the highest mis-match at car age 10. Turning to Figure 31, 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 32 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

Table 9: 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	
$\text{Pr}(0 0)$	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	0.13984***	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		74	169,733

Figure 31: Model Fit by State Variables



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 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.

D.2 Equilibrium Prices

In this section, we show additional results concerning the equilibrium price simulations. To reiterate, 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 33 shows simulations complementing figure 13 but showing all car age categories; in particular, the first- and final-year depreciations were omitted in Figure 13 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.

D.3 Counterfactual Simulations

Figures 34 and 35 accompany Figures 12 and 15 in Section 6.5. Figure 34 shows the macro state over the simulation and the fuel prices (which are held constant) and

In this section, we present results that are supplementary to the ones shown in section 6.5.

Figure 32: Forward Simulation




Figure 33: Simulations Under Equilibrium Prices: All Car Ages

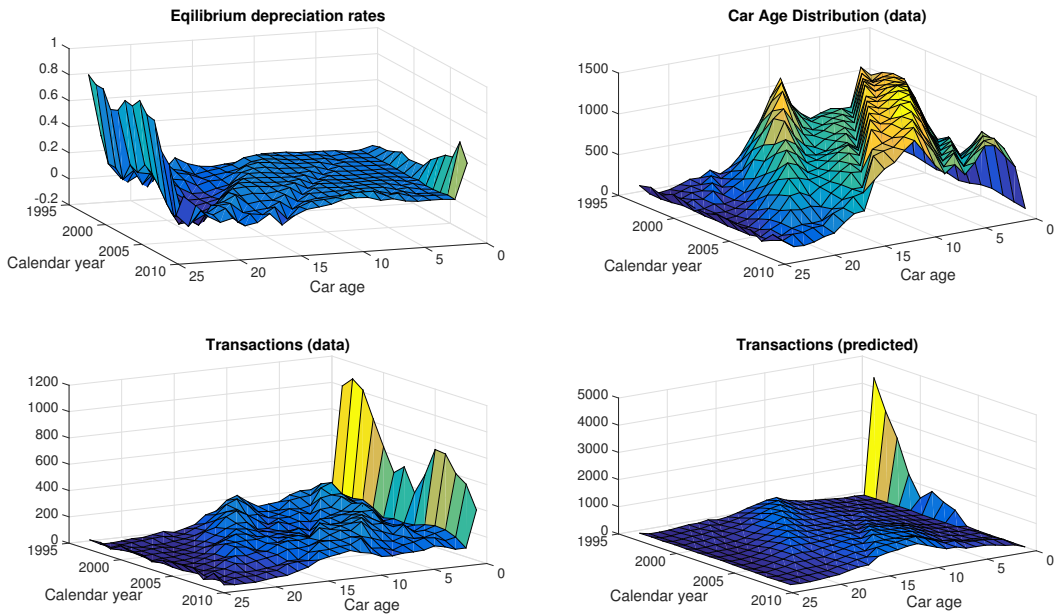


Figure 34: Forward Simulation with Equilibrium Prices

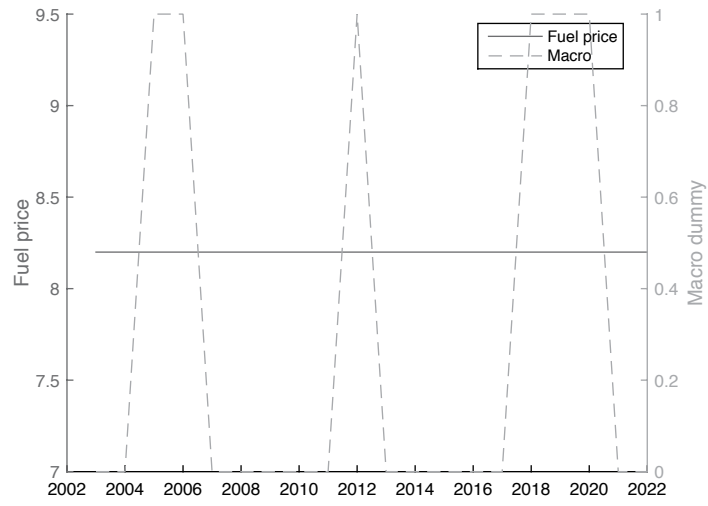


Figure 35: Forward Simulation with Equilibrium Prices

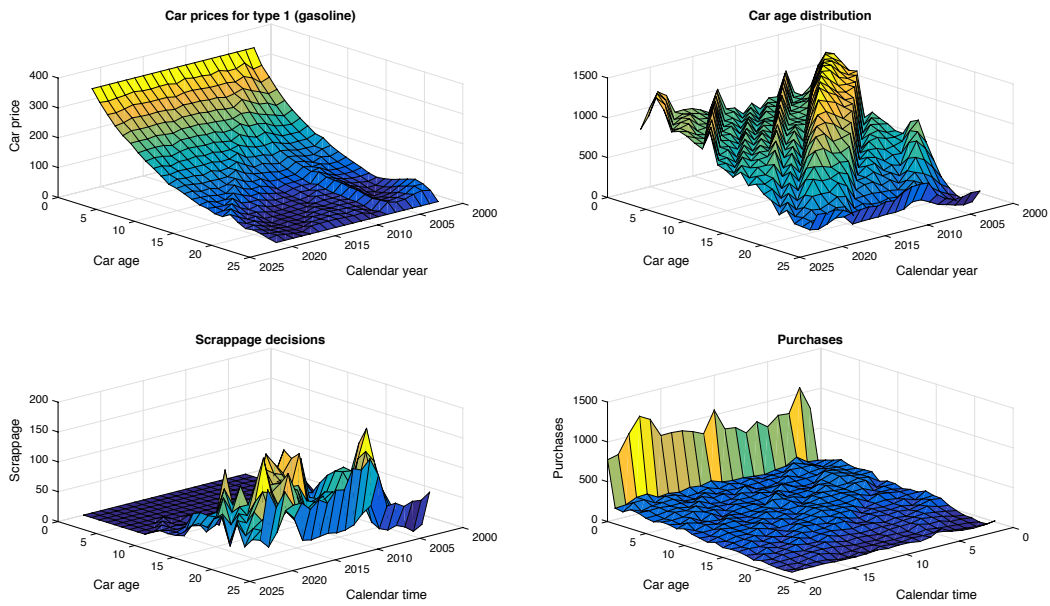


Figure 36: Counterfactual Simulations: Macro and Fuel Price Processes


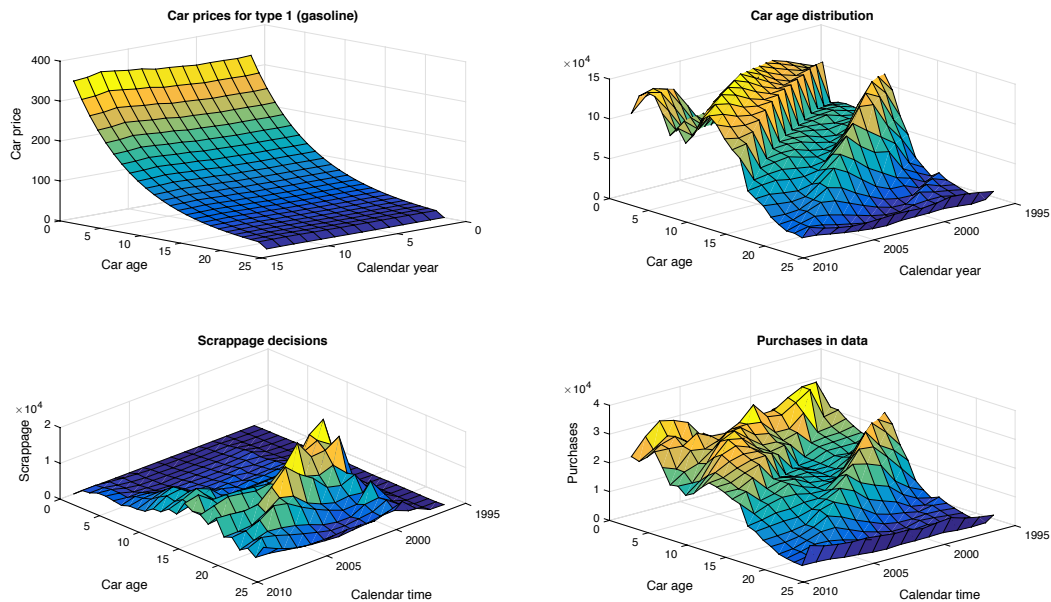


Figure 36 accompanies Figures 16 and 17 in showing the realized paths of the macro and fuel prices processes. Firstly, note that the fuel price is constant throughout the period except in the year 2012 where we counterfactually increase it by 50%.

To compare against the counterfactual simulation results in Figures 16 and 17, we show the corresponding graph for the actual data in Figure 37. 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,

Figure 37: Age Distribution, Scrappage and Purchases in the Data



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.
- 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.
- 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.
- ENGERS, M., M. HARTMANN AND S. STERN (2008): “Are Lemons Really Hot Potatoes?,” *manuscript, University of Virginia*.
- ESTEBAN, S. AND M. SHUM (2007): “Durable-goods Oligopoly with Secondary Markets: The Case of Automobiles,” *RAND Journal of Economics*, 38(2), 332–354.
- 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.
- 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.
- 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.

Table 10: 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.

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.

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): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55(5), 999–1033.

——— (1985b): "Stationary Equilibrium in a Market for Durable Assets," *Econometrica*, 53(4), 783–806.

SCHIRALDI, P. (2011): "Automobile Replacement: A Dynamic Structural Approach," *RAND Journal of Economics*, 42(2), 266–291.

STOLYAROV, D. (2002): "Turnover of Used Durables in a Stationary Equilibrium: Are Older Goods Traded More?," *Journal of Political Economy*, 110(6), 1390–1413.