Automobile Replacement: A Dynamic Structural Approach*

Pasquale Schiraldi
London School of Economics

The Toyota Centre
Suntory and Toyota International Centres for Economics and Related Disciplines
London School of Economics and Political Science
Houghton Street
London WC2A 2AE
Tel: (020) 7955 6674

EI/49
June 2010

* This paper has benefitted greatly from discussions with Marc Rysman. Special thanks are also due to Victor Aguirregabiria, Michael Manove, Martin Pesendorfer, Philipp Schimdt-Dengler and seminar participants at several institutions. All remaining errors are mine. Financial support from Bart Lipman, Michael Manove, the Institute for Economic Development - Boston University and the University of Bari (Italy) is greatly appreciated.

Correspondence: Pasquale Schiraldi, Department of Economics, London School of Economics, Houghton Street, London WC2A 2AE.
E-mail:p.schiraldi@lse.ac.uk. Phone: +44(0)2079557584.

Abstract

This paper specifies and estimates a structural dynamic model of consumer demand for new and used durable goods. Its primary contribution is to provide an explicit estimation procedure for transaction costs, which are crucial to capturing the dynamic nature of consumer decisions. In particular, transaction costs play a key role in determining consumer replacement behavior in both primary and secondary markets for durable goods. The unique data set used in this paper has been collected by the Italian Motor Registry and covers the period from 1994 to 2004. It includes information about sales dates for individual cars over time as well as the initial stock of cars in the sample period. Identification of transaction costs is achieved from the variation in the share of consumers choosing to hold a given car type each period, and from the share of consumers choosing to purchase the same car type that period. Specifically, I estimate a random coefficients discrete choice model that incorporates a dynamic optimal stopping problem in the spirit of Rust (1987). I apply this model to evaluate the impact of scrappage subsidies on the Italian automobile market in 1997 and 1998.

1 Introduction

In many durable goods industries, such as that of automobiles, used products are often traded in decentralized secondary markets. The U.S. Department of Transportation reports that in 2004 13.6 million new vehicles and 42.5 million used vehicles were sold in the U.S.A.; in the same year 2.5 million new vehicles and 4.7 million used vehicles were sold in Italy. Transactions in the secondary market may occur because the quality of a durable deteriorates over time and current owners sell their product in order to update to their preferred quality. Alternatively the level of required maintenance and/or the probability of failure may increase as the automobile ages, making replacement of the current unit desirable.

Durability and the presence of second-hand markets introduce dynamic considerations into both producers' output decisions and consumers' purchase decisions in the automobile market. Empirical models of demand for durable goods have focused mostly on the market for new products (See Berry, Levinshon and Pakes (1995) — henceforth BLP and Bresnahan (1981)). Using sophisticated simulation techniques embodied in the logit framework, these models are able to allow for general patterns of substitution across differentiated products. However, they do not usually account for the intertemporal dependence of consumers' decisions that characterize markets for durable goods. They either ignore the secondary market and its dynamics altogether or lump used goods into a composite outside option. In spite of their importance and although the auto market is one of the most studied in the literature (Bresnahan (1987), BLP (1995), Goldberg (1995), Petrin (2002)), there have been relatively few empirical models of secondary markets for used goods.

An important feature of the automobile market is that the stock of cars held by consumers is persistent over time. If a consumer owns a car in one year then it is likely that she will hold the same car the following year as well. The persistence of consumer holdings of automobiles, when durables depreciate over time, arises because of the presence of transaction costs such as search costs, taxes, asymmetric information, switching costs, etc, which vary over time. If there are no frictions a consumer would choose a quality that maximizes her utility in each period and have no incentive to hold it across multiple periods once the quality of the good depreciates. However, these frictions are present and they tend to make replacement infrequent because consumers try to economize on the costs associated with these frictions.

Any model that tries to explain the pattern of consumer holdings in a market for semi-durable goods must explicitly account for dynamic consumer considerations and the cost of the replacement decision. The model that I present incorporates both of these features. More specifically, I assume consumers incur different kinds of costs upon the replacement of their automobile (i.e. taxes, search costs, dealer compensations etc.). I do not separately identify the different components of these costs but they are lumped together in a unique parameter to be estimated[1]. The quality of used goods is assumed to be common knowledge among the agents, hence the model does not account explicitly for the presence of adverse selection. Information about resales and prices along with ownership data of used cars provides a potential source of identification for the transaction costs which has not been explored in the previous literature.

I use a data set containing information about the Italian car market to examine how unobserved heterogeneity and transaction costs affect replacement behavior. In particular, I observe the pattern of sales and ownerships for each individual car type in the sample over a period of 11 years. The possibility of following the history of each vehicle in the sample is due to the presence, in the data, of a unique identification number assigned to each unit. The data are from the Province of Isernia in Italy and are collected by the Motor Vehicle Department. Identification of transaction costs is achieved from the difference between the share of consumers choosing to hold a particular car type each period, and the share of consumers choosing to purchase the same car type that period. The presence of these two market shares for each car type represents the main strength of my unique data set. These market shares are obtained in the model by aggregating consumers' optimal decision that take the depreciation of automobiles over time into account. With full information this depreciation is captured by the decline in prices; then the pattern of sales and holdings, along with the pattern of prices in the data is used to identify the transaction costs. The structural model explicitly accounts for this information and provides an estimation of these costs for each product at each point in time.

The contribution of this paper to the durable goods literature is twofold. First, it is the first paper which studies replacement behavior in the presence of secondary markets, using aggregate data, while allowing for heterogeneity across consumers and endogeneity of price in a dynamic setting. Second, it shows how the combination of ownership and purchase data is useful to infer the size of transaction costs. Transaction costs play a central role in the analysis of market structure and industry conduct for a variety of industries. The proposed methodology can be used to measure transaction costs in the context of other industries as well.

Finally, I investigate the effect of scrappage subsidies offered by the Italian government to stimulate the early voluntary removal of used cars in 1997 and 1998. Such subsidies were temporary and offered in exchange for used cars of delineated vintages to reduce environmental pollution and stimulate car sales. Scrappage subsidies have been very popular in the European Union as well as in the United States and Canada. The possibility that such programs will be expanded has evoked a debate surrounding their effects on car markets

[1] While a unique parameter is estimated for each used car j in each time period t from the structural model, the nature of these costs is subsequently analyzed.

and consumers' welfare. The model is used to investigate the impact of such policies on consumers' demand for new and used vehicles. The point of doing the analysis is to quantify both the extent of the immediate incentive to replace and the subsequent effects of this policy on sales as the distribution of car ages evolves. The model allows me to illustrate and study the short and long run effects of the scrappage policies on new and used car sales and government revenues accounting for richer dynamics where new and used car markets interact.

I estimate a discrete choice logit model over a set of products with random coefficients on observable product characteristics that incorporates a dynamic optimal stopping problem in the spirit of Rust (1987) using market-level data. The random coefficients allow me to relax the so-called independence of irrelevant alternative (IIA) property (see BLP (1995), Browstone and Train (1999)) and allows the preferences to be correlated across vehicles. Thus I construct a generalized method of moments (GMM) estimator to deal with potential price endogeneity. This is possible provided that one can recover the unobserved product characteristics.

Berry (1994) suggests the use of a contraction mapping to find the mean product utilities. I use a similar contraction mapping to invert both the market share of purchases, and the market share of consumer holdings for each product in each period. The first market share refers to the share of consumers who decide to acquire a car j conditional on buying/replacing a vehicle. The market shares of consumer holdings refer to the share of consumers who keep car j conditional on owning that car. Both market shares for each car type deliver information about the mean level of utility and the mean level of transaction costs. As suggested by the model, if transaction costs are paid by buyers, the market share of consumer holdings conveys information on the mean product utility, whereas the market share of purchases will, in addition, convey information on transaction costs. For each product, I solve for the vectors of mean product characteristics and transaction costs that make the predicted shares match the observable ones.

Because no individual level data is available, I need to compute the aggregate predicted share of each product at any time period. Doing so requires integrating over the individual heterogeneity and consumer holdings once the consumer decides to replace her current vehicle. Then, I allow consumers to solve a dynamic optimization problem based on expectations about the stochastic process that governs the transition across different states of the durables and the market evolution. As in Rust (1987), the consumer's decision problem is formulated as an optimal stopping problem. Therefore, the consumer decides the optimal time period in which to replace her current vehicle with a different one. In my analysis, the consumer's decision to replace a car depends on her expectation about the future value of the product she currently owns and on the perceived distribution about the future set of products available.

The emphasis on the presence of a second-hand market with transaction costs and good depreciation distinguish the present model from BLP and Gowrisankaran and Rysman (2009) — henceforth GR. GR (2009) extended Melnikov's (2001) model to include consumer heterogeneity and examine the pattern of sales after the introduction of new digital cameras and DVD players in a dynamic setting. As in those models, the major simplifying assumption here is that consumers perceive the evolution of product characteristics to be a simple first order Markov process, where the distribution of the next period's product characteristics is a polynomial function of a simple statistic: the logit inclusive value (Melnikov, 2001 and Carranza, 2007). Gordon (2009) allows consumers to have the possibility of replacing the good and he does not allow for price endogeneity and heterogeneity across consumers.

There are recent studies that deal with the implications of durability and secondary markets on the dynamics of car demand. Esteban and Shum (2006) estimate a model with forward-looking consumers and firms. They reduce consumer heterogeneity to a single dimension, and do not consider the presence of transaction costs. Having a single dimension and considering a vertically differentiated market places strong restrictions on the substitutability among cars in consumers' choice sets.

Durables sold in second-hand markets are typically highly differentiated in quality and this captures some of the motivations for consumer holdings. Stolyarov (2002) uses a dynamic model with transaction costs to replicate the pattern of resales in the used car market. His model restricts consumer heterogeneity to a single dimension, but does allow for the possibility of infrequent replacement. He looks at a stationary environment in which all the goods are homogenous in all aspects but the age. Transaction costs increase deterministically over time. The model is calibrated to match the cross sectional pattern of resales. It does not allow transaction costs to be different across different cars and time. Adda and Cooper (2000) study the optimal decision rules from a dynamic discrete-choice model to explore the effects of scrappage subsidies on new car demand in France. In their model, consumers are homogenous so that in equilibrium, agents will choose either to keep the car or to replace it with a new one by scrapping their old car. Hence, in their model, in equilibrium there is no active secondary market. Consequently, a richer analysis of the scrappage design could be carried out that accounts for the sale dynamics in the primary and secondary market with different implications in terms of the impact of this policy. Finally, Hendel and Lizzeri (1999), Porter and Sattler (1999) and Schiraldi (2009) study vertical differentiated models in which durable goods live for just two periods, so that used goods of all ages are lumped together and derive some testable implications.

Complementary to these works, I contemporaneously allow for the presence of heterogenous consumers under multiple dimensions, for the possibility of the price to be correlated with the unobservable characteristics and for the presence of frictions on the secondary market given that durables depreciate over time. I use aggregate data to estimate the demand parameters and the distribution of transaction costs across models and over time.

The remainder of this paper is organized as follows. Section 2 analyses data. Section 3 discusses the model and the method of inference. Section 4 presents the results. Section 5 investigates the effect of scrappage subsidies on the Italian automobile and section 6 concludes.

2 The automobile market

2.1 Data

The Italian automobile market is the fourth largest market in the world (after the US, Japan and Germany) with about 2 million cars sold every year. Most cars sold are manufactured by the FIAT Group that controls the following brands: FIAT, Lancia, Alfa Romeo, Innocenti, Autobianchi, Ferrari and Maserati. The FIAT Group's share was more then 50% in 1990 and has gradually decreased since then. Volkswagen, the second largest manufacturer had a 14% market share; Ford between 7% and 11%; Citroen/Peugeot and Renault about 7% each; Opel between 5% and 8% and BMW/Mercedes between 3% and 4%.

The data set covers the period from January 1994 to December 2004 for the Province of Isernia in Italy. I have information on prices and characteristics of all new and most popular used cars sold in Italy. This information comes from *Quattroruote,* the main monthly automobile publication in Italy. Quantity data are provided by *ACI*, an association that runs the registration records for the Department of Motor Vehicles in Italy. Information about household income, population and price indexes for inflation are available at the Bank of Italy website and at the National Institute of Statistics website.[2] I report some demographics of the population in Table 1.

For all units in the sample, I observe the initial stock in 1994 and all subsequent individual transactions (sales, scrappage decisions, etc.), for each transaction I observe whether or not a car dealer was involved. I observe the manufacturer, the model, the engine displacement (cc), the car size, the first registration year, the plate for each car and the data track sales dates for individual cars over time. For the cars scrapped in 1997 and 1998, I have information on whether the owner opted to buy a new car and availed of the government subsidy. If the owner of a car moves to a location outside Isernia or sells it to a buyer living outside the Province, then that particular unit is excluded from the sample in the subsequent periods. It is similarly excluded if the owner decides to scrap the car. Analogously cars coming from

[2] www.bancaditalia.it, www.istat.it

outside Isernia are included in the sample in the years following the purchase of these cars.

In 1994, the first period of the sample, I observe an initial stock of 37,980 vehicles. Over the sample period I observe 82,254 transactions net of the transactions made by car dealers. To achieve a manageable dimensionality, I group them into 2,178 categories based on the year, on the vehicle's age (0,..,10) where 0 stands for a new car and 10 groups together all the cars 10 years or older[3], engine displacement (*small* if cc<=1300, *medium* if 1300<cc<=1800, *large* if cc>1800) and type of fuel: gasoline or diesel and origin of manufacturers.[4] In particular, I consider three possible macro-groups of manufacturers:

- the Italian FIAT-Group that controls the following brands (all located in Italy): FIAT, Lancia, Alfa Romeo, Innocenti, Autobianchi, Ferrari and Maserati
- manufacturers located in Germany: BMW, Mercedes, Opel, Volkswagen, Audi and Porsche
- a residual group that is mostly accounted for by Ford, Peugeot, Renault and Seat (the Korean and Japanese manufacturers have a very tiny market share due to the presence of quotas)

Up until 2000 *Quattroruote* provided price information only for cars that were up to 8, or in some cases 9, years old. I fill in the missing prices by assuming for each car model a subsequent depreciation rate (i.e. beyond the 8th or 9th year) equal to the depreciation rate the car experienced in the previous period.

In the empirical analysis, I focus on the market for passenger cars, excluding trucks, vans, minivans, SUVs and luxury cars (like Ferrari and Lamborghini), in part because I do not have price information for them. The total proportion of these cars is less than 2% of the initial stock and about 2% of all the transactions over the 11 years. Furthermore, I assume that the owners of a 10-year old car receive the market price of that car type irrespective of whether they decide to sell or scrap the car.

Table 3 reports some descriptive statistics about prices and quantities of new and used cars in the data. Figure 1 shows the pattern of sales of new and used cars in the data. The total amount of new units purchased suddenly jumped in 1997 when the government introduced the scrappage policy. The scrappage policy, which involved subsidizing car replacement, was aimed at increasing road safety, reducing environmental pollution and stimulating car sales.

[3] I assume that a 10-year old car no longer depreciates and provides the same utility to the consumer. Therefore, I also assume that the price is the same across cars older than 10 years except for the stochastic component ξ_{jt}.

[4] The choice of engine displacement as a key characteristic to identify the different products seems natural in this context for two reasons. First, the scrappage-policies was designed according to this characteristic (as explained later) and second, until 1999 property taxes were paid based on the size of the engine displacement.

From January 1997 until September 1997 the government awarded a bonus, the amount of which depended on the size (engine displacement) of the new replacement bought. The cash subsidy accruing to consumers was conditional on buying a new car and the burden was jointly borne by the government and the car manufacturer. The program was scheduled to expire in September 1997 but was extended until the end of the year. In 1998, a similar scheme, lasting from February to September, was introduced. Observe that the purchases of used cars slowed down in 1997 and 1998 and there was a steep increase in the following years. The increase in the number of used cars traded indicates a more active second-hand market over time. The increase in the volume of used cars transactions is explained by the reduction of taxes to pay upon registration over the considered time horizon.

2.2 A closer look at transaction costs

The distinguishing feature of durable goods, and in particular of automobiles, is its potential for resale. In the absence of some sort of market frictions or information asymmetries, consumers have no incentive to hold their durables across multiple quality levels. Each heterogeneous consumer will choose a durable from the product spectrum so as to maximize her net surplus. Hence, if there are no frictions in the market it should observe a high turnover rate. Specifically, if the good depreciates every period (as in my model) consumers will never hold their durable more than one period but they will always update to their preferred car quality (see for example Hendel and Lizzeri 1999a, Hideo and Sandfort 2002, Rust 1985 among others). While in my model consumers incur different kinds of costs upon the replacement of their automobile (i.e. taxes, search costs, dealer compensations etc.), the quality of used goods is assumed to be common knowledge among the agents. Therefore no adverse selection is present in the used car market. The transaction costs associated play a crucial role in explaining the consumers' decisions to replace their car. Each period, consumers assess the quality of the durable they own. If the gain in utility from updating their holdings, net of prices, exceeds transaction costs, consumers sell their used goods in the secondhand market and replace them with durables of the preferred quality. It follows that the high level of transaction costs reduces the frequency of replacement. The two driving assumption are: (1) cars depreciate after every year; (2) consumer preferences (for car characteristics) are perfectly persistent.

A first look at the data, and in particular at the average resale ratio (i.e., the percentage of the stock of a given type, or brand, of a given age of car resold in a period; durable good trading volume) across all cars in each year gives an idea about the presence and size of transaction costs. We can observe, by looking at Figure 2, that the average resale ratio per year varies between 0.15 and 0.25.[5] This is substantially lower than 100% trade as would be predicted by vertical differentiation models with consumers' heterogeneity but without frictions, and suggest the presence of a relevant level of transaction costs (for a comparison between markets with and without frictions see Anderson and Ginsburgh 1994, Hendel and Lizzeri 1999b, Stolyarov 2002 among others). Figure 3 reports the same ratios at a more disaggregate level: I report the resale ratios for Fiat and Volkswagen compact cars. Figure 3 suggests the presence of different levels of transaction costs for different models/types. A strength of my model and my estimation procedure is that it can recover the whole distribution of transaction costs for different car types in each point in time.

Consumers with a preference for quality wish to replace durables that deteriorate. When quality deterioration is small, some buyers of the new durable retain ownership of their unit rather than incur the costs associated with transaction of the used durable goods. When quality deterioration of a durable good is large relative to the transaction costs of the used durable good market, however, more owners of the durable will wish to sell their units and purchase a new unit in the current period. Thus the durable good's volume of trade is also directly related to the quality depreciation. With full information, prices reflect this deterioration.[6] Exploiting this insight, I regress the resale ratios for different models/ages of cars on the price depreciation and car characteristics (the age of the car, the engine size (cc)). To control for unobservable quality, I also add model dummies and time dummies. Results are reported in Table 4. Notice that the volume of trade is positive correlated with price depreciation as we should expect in a model with transaction costs.[7] The presence of transaction costs in the market could cause a potential endogeneity problem in the previous regression, the transaction costs which are omitted could potentially be correlated with the price depreciation. Starting from 2001, Quattroruote publishes a rating (that runs from 1 to 5) which reflects how easy is to trade each given car on the secondary market. To further analyze the presence of transaction costs, I collected this data from 2002 to 2004 and then ran the same regression as above, also including this additional variable. More specifically, I included these ratings (with negative sign) which should capture the presence of transaction costs related to specific car models. I report the results in Table 5. As expected, the volume of trade is negative correlated with such a variable, whereas the sign of price depreciation does not change.

[5] The unit of observation is model/age for different car segments: Sub-compact cars, Compact cars, Medium size cars, Med/Full-size cars, Full-size cars.

[6] See Hendel and Lizzeri 1999a and Gilligan 2004 for further discussion about the effect of asymmetric information.

[7] The model does not explicitly address the potential adverse selection problem. As noted in Gilligan (2004), the presence of transaction costs and consequently a positive correlation between volume of trade and price depreciation does not necessarly exclude also the presence of asymmetric information in the market.

Having presented evidence for the presence of transaction costs, the next sections are devoted to structural estimate a model for consumers' replacement to measure the size of transaction costs and derive better demand estimates and consequently more suitable price elasticity. Finally I use the model to investigate the scrapping policy.

3 Model and Inference

Consider an infinite horizon model and finite types of durable goods (BMW, Mercedes, FIAT, and so on). The good is durable, but it depreciates over time. A physical stochastic process describes the transformation of the condition of the vehicle in period t to its condition in period $t+1$.

Each consumer is assumed to consume, at most, one unit of the good. Since products degrade over time, a given consumer will desire to replace her durable over time, either with a brand new durable or with a secondhand one. In the model, consumers have perfect information about durables so that there is no lemon problem. In addition there is a perfectly divisible good (money), which is treated as numeraire. Consumers maximize the expected lifetime utility using a discount factor $\beta \in (0, 1)$.

Let $\mathbf{j}_t$ denote the set of new cars available in period t and $J_t = \{j : j \in \{\cup_{\tau=1}^{t} \mathbf{j}_\tau\}\}$ this denotes the set of all possible products attainable in period t in the primary or secondary market. In every period there is always the possibility to opt for the outside option, i.e. $j = 0$, which corresponds to not owning a car.

At the beginning of each period, each consumer i may or may not have a car endowment from previous purchases. If she does not have any vehicle, she simply decides whether or not to purchase one. If she has a car endowment, immediately upon entering period t the durable depreciates according to the exogenous depreciation process. Then the consumer decides whether to hold, sell or scrap that car. If she gets rid of the car (via scrap or sale), she also decides whether or not to purchase a different car among the $J_t \cup \{0\}$ products present in the primary and secondary market in period t (including the outside option). In either case, she faces a similar (though not identical) decision problem in time $t+1$. Since consumers can delay purchase, they face a dynamic optimization problem of when, if ever, to purchase any given (new or used) car available. The consumer's choice maximizes her expected discounted utility conditional on her information and endowment in each given period.

Each product $j \in J_t$ is characterized by observed physical characteristics x_{jt} (for example engine displacement, fuel, age, size, etc.), the unobserved (by the econometrician) product characteristic ξ_{jt}, the price p_{jt} and the unobserved (by the econometrician) transaction cost τ_{jt}. I assume that the transaction cost is paid by the consumer (along with the price) every time that she purchases a car and it captures the presence of searching costs, financial costs, switching costs, asymmetric information and so on. I also assume that no transaction costs are paid if the consumer opts for the outside option and $p_{0t} = 0$.

Consumers are heterogenous in the evaluation of how intensively they prefer car characteristics and in their price sensitivity. Consumers also have an idiosyncratic shock to their preferences for each good and in each period. Let $\boldsymbol{\epsilon}_{it} = (\epsilon_{i0t}, \epsilon_{i1t}, ..., \epsilon_{iJ_t t})$ be the vector of idiosyncratic shocks of consumer i for period t and capture horizontal differentiation, which are $i.i.d.$ across (i, j, t).

A consumer i derives the following one-period utilities for each of the possible choices at time t. If the consumer i keeps the car she already owns (the car $k \in J_{t-1}$) she gets utility

$$\tilde{u}_{it}^{k} = x_{kt}\,\alpha_i^x + \xi_{kt} + \epsilon_{ikt} \tag{1}$$

If the consumer sells her car and purchases a different car $j \in J_t$, she gets utility

$$u_{it}^{kj} = x_{jt}\,\alpha_i^x + \xi_{jt} - \alpha_i^p p_{jt} - \tau_{jt} + \alpha_i^p p_{kt} + \epsilon_{ijt} \tag{2}$$

If she replaces the car she pays the price and the transaction costs for the new car, p_{jt} and τ_{jt}, and cashes the value of her endowment, p_{kt}. If, instead, the consumer sells the car she owned and does not purchase a replacement, she gets utility

$$u_{it}^{k0} = \alpha_i^p p_{kt} + \epsilon_{i0t} \tag{3}$$

A consumer who does not hold any product in period t obtains a mean flow utility which I set equal to zero and in particular $\tilde{u}_{it}^{0} = \epsilon_{i0t}$. Assume that the error term, ϵ_{ijt}, is independent across consumers, products and time and is Type I extreme value distributed. Finally α_i^p represent consumer i's price sensitivity for cars and α_i^x is an individual-specific preference for car characteristics.

The same product in subsequent years differs by the age and by its unobservable characteristics, hence both these elements capture the depreciation of durables over time. The depreciation is not deterministic because the unobserved product characteristic evolves stochastically over time.

Formally, consumer i who initially owns a durable k seeks an infinite sequence of decision-rules μ_t to maximize the expected, present discounted sum of future utility, or

$$\max_{\{\mu_t\}_{t=0}^{\infty}} E\left\{\sum_{t=0}^{\infty} \beta^t g_i\left(k, \mu_t\right) | \Omega_{it}, \boldsymbol{\epsilon}_{it}\right\}$$

where $\mu_t = \{d, j\}$ where d, denotes a consumer's replacement decision at time t, $d = 0$ (keep), $d = 1$ (replace) and $j \in J_t \cup \{0\}$ is the optimal replacement at time t if $d = 1$. Ω_{it}

includes current product attributes and prices, product availability, the year, and any other market characteristics which may affect the firms product pricing, entry, exit, or change in attributes.[8] In general, it includes all variables at time t in consumer i's information set that affect her utility or value for waiting. I assume that Ω_{it+1} evolves according to some Markov process $P(\Omega_{it+1}|\Omega_{it})$ that will account for firm optimizing behavior. Finally

$$g_i(k,\mu_t) = \begin{cases} \tilde{u}_{it}^{k} \text{ if } d=0 \\ u_{it}^{kj} \text{ if } d=1 \text{ and } j \in J_t \cup \{0\} \end{cases}$$

In each period, a consumer chooses her optimal action given her initial endowment k, preferences, current product qualities, prices, product availability, and expectations over future values of these characteristics and in particular over the stochastic value of her endowment. To solve the consumer's problem, I must solve for the value function $\hat{V}_i$ which is the unique solution to Bellman's equation. A consumer's value function from being on the market for a car, conditional on following her optimal policy and her initial endowment k, is given by:

$$\begin{aligned} \hat{V}_i(k,\boldsymbol{\epsilon}_{it},\Omega_{it}) = \max\{ & \underbrace{\tilde{u}_{it}^{k} + \beta E[\hat{V}_i(\boldsymbol{\epsilon}_{it+1},\Omega_{it+1},k)|\Omega_{it},\boldsymbol{\epsilon}_{it})]}_{\text{Keep the same car}}, \\ & \underbrace{\max_{j\in J_t} u_{it}^{kj} + \beta E[\hat{V}_i(\boldsymbol{\epsilon}_{it+1},\Omega_{it+1},j)|\Omega_{it},\boldsymbol{\epsilon}_{it})]}_{\text{Buy optimal replacement today}}, \\ & \underbrace{u_{it}^{k0} + \beta E[\hat{V}_i(\boldsymbol{\epsilon}_{it+1},\Omega_{it+1},0)|\Omega_{it},\boldsymbol{\epsilon}_{it})]}_{\text{Sell the car and consume the outside option}}\} \end{aligned} \tag{4}$$

If $k=0$ the consumer value function is

$$\begin{aligned} \hat{V}_i(0,\boldsymbol{\epsilon}_{it},\Omega_{it}) = \max\{ & \underbrace{\max_{j\in J_t} u_{it}^{kj} + \beta E[\hat{V}_i(\boldsymbol{\epsilon}_{it+1},\Omega_{it+1},j)|\Omega_{it},\boldsymbol{\epsilon}_{it})]}_{\text{Buy optimal car today}}, \\ & \underbrace{\tilde{u}_{it}^{0} + \beta E[\hat{V}_i(\boldsymbol{\epsilon}_{it+1},\Omega_{it+1},0)|\Omega_{it},\boldsymbol{\epsilon}_{it})]}_{\text{Still consume the outside option}}\} \end{aligned} \tag{5}$$

The state space of the problem is too large for the consumer's full dynamic automobile replacement decision problem to be computationally solvable. Hence in the next subsection, I make some assumptions in order to reduce the dimensionality of the state space.

3.1 Simplifications and assumptions

The goal of the present subsection is to introduce and discuss different assumptions to simplify the optimal dynamic problem, to reduce the dimension of the state space and to make it computationally tractable.

[8] In paricular, I assume Ω_{it} includes also the attribute of the good of the good hold by the consumer.

The first step towards simplifying the problem is to write the value function in a more convenient way. It is opportune to subtract the price of the car owned from equation (4), then substituting for (1), (2) and (3) we can rewrite it as follows:

$$\begin{aligned}\hat{V}_i(k, \boldsymbol{\epsilon}_{it}, \Omega_{it}) - \alpha_i^p p_{kt} \quad = \quad & \max \Big\{ \epsilon_{i0t} + \beta E[\hat{V}_i(\boldsymbol{\epsilon}_{it+1}, \Omega_{it+1}, 0)|\Omega_{it}, \boldsymbol{\epsilon}_{it})] \ , \\ & \max_{j \in J_t} x_{jt}\ \alpha_i^x + \xi_{jt} - \alpha_i^p p_{jt} - \tau_{jt} + \epsilon_{ijt} + \beta E[\hat{V}_i(\boldsymbol{\epsilon}_{it+1}, \Omega_{it+1}, j)|\Omega_{it}, \boldsymbol{\epsilon}_{it}], \\ & x_{kt}\ \alpha_i^x + \xi_{kt} - \alpha_i^p p_{kt} + \epsilon_{ikt} + \beta E[\hat{V}_i(\boldsymbol{\epsilon}_{it+1}, \Omega_{it+1}, k)|\Omega_{it}, \boldsymbol{\epsilon}_{it}] \Big\}\end{aligned} \tag{6}$$

If I redefine $V_i(k, \boldsymbol{\epsilon}_{it}, \Omega_t) \equiv \hat{V}_i(k, \boldsymbol{\epsilon}_{it}, \Omega_t) - \alpha_i^p p_{kt}$ and substitute it back into the previous equation I have[9]:

$$\begin{aligned}V_i(k, \boldsymbol{\epsilon}_{it}, \Omega_{it}) = & \max \{ \epsilon_{i0t} + \beta E[V_i(\boldsymbol{\epsilon}_{it+1}, \Omega_{it+1}, 0)|\Omega_{it}, \boldsymbol{\epsilon}_{it})] \ , \\ & \max_{j \in J_t} x_{jt}\ \alpha_i^x + \xi_{jt} - \alpha_i^p p_{jt} - \tau_{jt} + \epsilon_{ijt} + \beta E[\alpha_i^p p_{jt+1} + V_i(\boldsymbol{\epsilon}_{it+1}, \Omega_{it+1}, j)|\Omega_{it}, \boldsymbol{\epsilon}_{it}], \\ & x_{kt}\ \alpha_i^x + \xi_{kt} - \alpha_i^p p_{kt} + \epsilon_{ikt} + \beta E[\alpha_i^p p_{kt+1} + V_i(\boldsymbol{\epsilon}_{it+1}, \Omega_{it+1}, k)|\Omega_{it}, \boldsymbol{\epsilon}_{it}] \}\end{aligned} \tag{7}$$

where the expected price in $t+1$ enters linearly because prices enter additively in the current utility function (see equations 1, 2 and 3). Notice that $p_{0t} = 0$ implies $V_i(0, \boldsymbol{\epsilon}_{it+1}, \Omega_{it+1}) = \hat{V}_i(0, \boldsymbol{\epsilon}_{it+1}, \Omega_{it+1})$.

Hence equation (5) is equal to:

$$\begin{aligned}V_i(0, \boldsymbol{\epsilon}_{it}, \Omega_{it}) \quad = \quad & \max\{ \epsilon_{i0t} + \beta E[V_i(\boldsymbol{\epsilon}_{it+1}, \Omega_{it+1}, 0)|\Omega_{it}, \boldsymbol{\epsilon}_{it}], \\ & \max_{j \in J_t} x_{jt}\ \alpha_i^x + \xi_{jt} - \alpha_i^p p_{jt} - \tau_{jt} + \epsilon_{ijt} + \beta E[\alpha_i^p p_{jt+1} + V_i(\boldsymbol{\epsilon}_{it+1}, \Omega_{it+1}, j)|\Omega_{it}, \boldsymbol{\epsilon}_{it}] \}\end{aligned} \tag{8}$$

Notice that Ω_{it} already includes the vector of prices, hence there is no need to explicitly include the resale price as a state variable, once the change of variable is performed. This transformation is not strictly necessary in order to estimate the model but it is convenient to reduce the computational costs of the estimation. I defer further discussion after the definition of the logit inclusive value.

In order to evaluate consumer i's choice at time t, I need to formalize consumer i's expectations about the utility from future products and from the product that she may potentially own. I assume that consumers have no information about the future values of the idiosyncratic unobservable shocks ϵ_{ijt} beyond their distribution. The set of products, their prices, their characteristics and transaction costs vary across time due to entry and exit, technological progress and changes in prices for existing products, according to optimal price decisions. Consumers are uncertain about the future product attributes, but rationally

[9] Notice Ω_{it} already includes also the vector of prices, hence I do not need to explicitly include p_{kt} as a state variable.

expect them to evolve, based on the current market structure. Consequently, the dynamic consumers' optimization problem potentially depends on the whole set of information, Ω_{it}, available in period t and the particular endowment k of each consumer i at time t.

The main issue in the estimation procedure is the "curse of dimensionality" usually associated with these kinds of problems. To simplify the problem I make some assumptions in line with the existing literature. As in Rust (1987), let $EV_i(j, \Omega_{it}) = \int_{\boldsymbol{\epsilon}_{it}} V_i(j, \boldsymbol{\epsilon}_{it}, \Omega_{it}) dP_{\boldsymbol{\epsilon}}$ $\forall j \in J_{t-1} \cup \{0\}$ denote the expectation of the value function, integrated over the realization of $\boldsymbol{\epsilon}_{it}$, which follows from Rust's conditional independence assumption. The next step to reduce the dimensionality of the state space is to identify a few variables that can summarize the information available in each moment and describe how consumers form their expectation based on these elements. This simplification will be done by introducing a state variable (the net augmented utility flow) which captures the depreciation of the good over time, along with the more common logit inclusive value[10], a scalar-valued sufficient statistic used in the literature to characterize the distribution of future payoffs. This defines the net augmented utility flow as:

$$\phi_{ijt} \equiv x_{jt}\, \alpha_i^x + \xi_{jt} - \alpha_i^p (p_{jt} - \beta E_t \left[p_{jt+1} \right]) \tag{9}$$

where $(p_{jt} - \beta E_t \left[p_{jt+1} \right])$ is the rental price of car j in period t. The rental price accounts for the cost of keeping a particular good j for a single period of time. The net augmented utility flow, ϕ_{ijt}, captures the mean flow utility derived by the consumer i from keeping the durable net of the rental price; it includes both elements of consumer characteristics and elements of product characteristics.[11] I also define the mean net augmented utility flow as $\hat{\phi}_{jt} = x_{jt}\, \alpha^x + \xi_{jt}$ which is a product specific term common to all consumers and will be used in the subsection (3.2.1).[12]

In a durable-goods setting, where the quality of the goods changes over time and there is the possibility of reselling, consumers maximize the utility derived from the good in any particular period net of the implicit rental price paid in that period to keep the good. Hence the net augmented utility flow seems a natural index, that captures the per period quality adjusted by the price that consumers account to make their decisions.

Finally, I use the aggregation proprieties of the type 1 extreme value distribution of $\boldsymbol{\epsilon}_{ijt}$ to express the expectation of the Bellman equation in a relatively simple form. In particular, the expected value of the best choice from several options in a logit model, can be expressed

[10] The logit inclusive value was first introduced by Melnikov (2001) and subsequently used in the literature of dynamic demand models.

[11] Note that for the outside option $\phi_{i0t} = 0$.

[12] $\hat{\phi}_{jt}$ is a product specific term common to all consumers whereas ϕ_{ijt} includes both the product specific terms and consumers specific terms. In particular given $\hat{\phi}_{jt}$, the simulated draws v_i and the vector of nonlinear parameters $\{\alpha^p, \sigma\}$, ϕ_{ijt} is easily derived: $\phi_{ijt} = \hat{\phi}_{jt} + \sum_j x_{jt}\, \sigma v_i - \alpha_i^p (p_{jt} - \beta E_t \left[p_{jt+1} \right])$.

as the logarithm of the sum of the mean expected discounted utility of each option and for each consumer i:

$$\delta_{it} = \ln\left(\sum_{j\in J_t} \exp\left(\phi_{ijt} - \tau_{jt} + \beta E[EV_i(\Omega_{it+1}, j)|\Omega_{it}]\right)\right) \tag{10}$$

The logit inclusive value is the maximum expected utility from buying one of the J_t products present in the primary and secondary market in period t. Notice the importance of the transformation made at the beginning of this subsection, this allows me to uniquely define the net augmented utility flow independently from the fact that the car is purchased or owned by consumers. Moreover, without the above transformation, the logit inclusive function would also have been function of the price of the car owned. Consequently prices should have been accounted explicitly among the state variables, with a substantial increment of the computational time needed to estimate the model.[13]

To reduce the dimensionality of the state space and describe how consumers form their expectation, I assume that each consumer perceives the evolution of the net augmented utility flow and the logit inclusive value to evolve, according to a first-order process that depends on the previous value of the variables themselves:

Assumption 1. *Each consumer i perceives that $(\phi_{ikt}, \delta_{it})$ and can be summarized by a first-order Markov process:*

$$G_i\left(\phi_{ikt+1}, \delta_{it+1}|\Omega_{it}\right) = G_i\left(\phi_{ikt+1}, \delta_{it+1}|\phi_{ikt}, \delta_{it}\right)$$

where G_i is consumer specific.

Similarly to the Inclusive Value Sufficiency assumption introduced by GR, the previous assumption implies that all states characterized by the same pair $\{\phi_{ikt}, \delta_{it}\}$ have the same expected value. This assumption can be interpreted as an assumption that consumers are boundedly rational and use only a subset of the data potentially available to them, in forming their expectations. Although reducing the state space dramatically, this assumption may not be consistent with an underlying supply model. Many different quality or market characteristics could potentially lead to the same value of $\{\phi_{ikt}, \delta_{it}\}$, which may have different implications in the evolution of the industry, and nevertheless will imply the same expectation about their evolutions in the present model. For the estimation of the model, I assume that the Markov processes take the following linear functional form:

$$\delta_{it+1} = \rho_{1i} + \rho_{2i}\delta_{it} + \eta_{it} \tag{11}$$

[13] In particular, the logit inclusive value would have been $\delta_{it}^k = \ln\left(\sum_{j\in J_t} \exp\left(\phi_{ijt} - \tau_{jt} + \alpha_i^p p_{kt} + \beta E[EV_i(j, \boldsymbol{\epsilon}_{it+1}, \Omega_{it+1})|\Omega_{it}]\right)\right)$

where $\phi_{ijt} \equiv x_{jt}\ \alpha_i^x + \xi_{jt} - \alpha_i^p p_{jt}$ is for the cars available to purchase on the market and $\phi_{ikt}^k \equiv x_{jt}\ \alpha_i^x + \xi_{jt}$ is for the car owned.

$$\phi_{ijt+1} = \gamma_{1i} + \gamma_{2i}\phi_{ijt} + \mu_{it} \tag{12}$$

where η_{it} and μ_{it} are jointly normally distributed and ρ_{1i}, ρ_{2i}, γ_{1i} and γ_{2i} are incidental parameters specific to each consumer i. Focusing on the evolution of the logit inclusive value, similar functional forms as (11) have been used in the existing dynamic literature (see Melnikov, 2001; Hendel & Nevo, 2006). However, the implication of the assumption is more closely related to GR (2009) because the specification of δ_{it} includes not only prices and characteristics of the products available, but also future optimal decision-making.[14] In Melnikov, the simplification results from the assumption that there is no repeat purchase, implying that the choice of a product is the final choice made by the consumer. Hendel & Nevo achieve this result by specifying δ_{it} only over (exogenous) characteristics that affect the consumer in the current period and do not affect dynamic decision making. In the present model, I allow all quality characteristics of a purchased product to affect future upgrading decisions. The drawback of this definition is that it makes assumptions on the evolution of δ_{it} potentially more restrictive: the assumption reflects consumer decision-making (flow but continuation values), which is endogenous to the model.[15] However in the present framework, consumers' expectations are also based on the evolution of the net augmented utility flow which, in the spirit of the previous papers, includes only the price and characteristics of the product owned by consumers and loosens the more restrictive implications of (11) alone used in GR (2009).

Using the previous assumptions, I can write $EV_i(\Omega_{it})$ as $EV_i(\phi_{ikt}, \delta_{it})$ and rewrite the Bellman equations (7) and (8) for consumer i as:

$$\begin{aligned} EV_i(\phi_{ikt}, \delta_{it}) \;\; = \;\; & \ln(\exp(\delta_{it}) + \exp(\phi_{ikt} + \beta E[EV_i(\phi_{ikt+1}, \delta_{it+1}|\phi_{ikt}, \delta_{it})) \\ & + \exp(\beta E[EV_i(0, \delta_{it+1}|\phi_{ikt}, \delta_{it}))) \end{aligned} \tag{13}$$

$$EV_i(0, \delta_{it}) = \ln(\exp(\delta_{it}) + \exp(\beta E[EV_i(0, \delta_{it+1}|\delta_{it}))) \tag{14}$$

The aggregate demand for a product is determined by the solution to the consumer's optimization problem. Specifically the probability that a consumer of type i with good k $\in J_{t-1} \cup \{0\}$ purchases a good $j \in J_t \cup \{0\}$ is:

$$d_{it}^{kj} = \frac{\exp\left(\phi_{ijt} - \tau_{jt} + \beta E\left[EV_i\left(\phi_{ijt+1}, \delta_{it+1}\right)|\delta_{it}, \phi_{ijt}\right]\right)}{\exp(\delta_{it}) + \exp\left(\beta E\left[EV_i\left(0, \delta_{it+1}\right)|\delta_{it}\right]\right) + \exp\left(\phi_{ikt} + \beta E\left[EV_i\left(\phi_{ikt+1}, \delta_{it+1}\right)|\delta_{it}, \phi_{ikt}\right]\right)} \tag{15}$$

Let $\widetilde{d}_{it}^{k}$ denote the probability that a consumer of type i with good $k \in J_{t-1} \cup \{0\}$ chooses not to make a purchase and retain her existing product:

[14] Similar assumption has also been used by Shcherbakov (2008) to study switching costs between cable and satellite television.

[15] See Gowrisankaran and Rysman (2009) for further discussion.

$\widetilde{d}_{it}^{k} =$

$$\frac{\exp(\phi_{ikt} + \beta E\left[EV_i\left(\phi_{ikt+1}, \delta_{it+1}\right)|\delta_{it}, \phi_{ikt}\right])}{\exp(\delta_{it}) + \exp\left(\beta E\left[EV_i\left(0, \delta_{it+1}\right)|\delta_{it}\right]\right) + \exp\left(\phi_{ikt} + \beta E\left[EV_i\left(\phi_{ikt+1}, \delta_{it+1}\right)|\delta_{it}, \phi_{ikt}\right]\right)} \tag{16}$$

where $\sum_{j \in J_t \cup \{0\}} d_{it}^{kj} + \widetilde{d}_{it}^{k} = 1$.

Let s_{jt}^{D} denote the unconditional market share of consumers that purchase j in period t and let $\widetilde{s}_{kt}^{H}$ be the proportion of consumers holding a durable k in period t. Hence, the total proportion of consumers having good j at the end of period t is $s_{jt} = \widetilde{s}_{jt}^{H} + s_{jt}^{D}$. Respectively s_{jt}^{D} and $\widetilde{s}_{kt}^{H}$ are obtained by integrating d_{it}^{kj} and $\widetilde{d}_{it}^{k}$ over consumer preferences and summing uo d_{it}^{kj} over all existing products:

$$s_{jt}^{D} = \int_{\upsilon_i} \sum_{k \in J_{t-1} \cup \{0\}} d_{it}^{kj} s_{ikt-1} dP_{\upsilon} \tag{17}$$

$$\widetilde{s}_{kt}^{H} = \int_{\upsilon_i} \widetilde{d}_{it}^{k} s_{ikt-1} dP_{\upsilon} \tag{18}$$

where s_{ikt} is the fraction of consumers of type i that own product k at the end of period t. In particular:

$$s_{ijt} = \widetilde{s}_{ijt}^{H} + s_{ijt}^{D} \tag{19}$$

where $\widetilde{s}_{ijt}^{H}$ and s_{ijt}^{D} are obtained as in equations (17) and (18) without integrating over the consumer heterogeneity. The market size M_t is observed and evolves deterministically over time.

3.2 Estimation

I set the discount factor $\beta = 0.9$ and the total market size M equal to the adult population in the area. I have also used the observed prices as proxy for the expected prices in computing the rental value when I perform the estimation. As in Berry, Levinson and Pakes (1999), I assume that the price sensitivity varies with income. Accordingly, I assume that α_i^p has a time-varying distribution that is a lognormal approximation to the distribution of income in this region of Italy in each year. If y_i is a draw from this lognormal income distribution, then $\alpha_i^p = \frac{\alpha^p}{y_i}$, where α^p is a parameter to be estimated. In this way, price sensitivity is modeled as inversely proportional to income. This allows me to use the exogenously available information on the income distribution to increase the efficiency of our estimation procedure. Moreover, I assume that consumers differ in the preference for the age of the car. More specifically, I assume that consumer preference for the age of the car α_i^{age} are independently distributed normally with mean α^{age} and standard deviation $\sigma_{\alpha^{age}}$, i.e. $\alpha_i^{age} = \alpha^{age} + \upsilon_i \sigma_{\alpha^{age}}$ where $\upsilon_i \sim N(0,1)$.

Following Berry's (1994) strategy, I recover the set of unobservable product characteristics and transaction costs (ξ_{jt}, τ_{jt}) for any parameter vector θ that perfectly rationalize the model's predicted market shares[16], and then employ a generalized method of moments (GMM) estimator, via forming conditional moments. I leverage the dynamic nature of my data and assume that the unobservable characteristics for each automobile evolve according to an exogenous Markov process, and these innovations in product unobservables, ξ_{jt}, (and not the product unobservables themselves) are uncorrelated with a vector of instruments.[17] More specifically, I assume:

Assumption 2. *Unobservable product characteristics for each automobile evolve according to a first-order autoregressive process, where the error terms*

$$\varsigma_{jt} = \xi_{jt} - \lambda \cdot \xi_{jt-1}$$

are independent of each other and

$$E\left[Z_{jt}\varsigma_{jt}\right] = 0$$

where Z_{jt} *are instruments.*

The drift of this process is set to 0 since it is not separately identified from the constant in the mean utility and λ is a parameter to be estimated. The instruments used are described in subsection 3.2.2. Moreover, I add a set of moments, chosen to improve the identification of consumers' price sensitivity: the fraction of people who used the scrappage scheme to replace their old automobile with a brand new one in 1997 and 1998.

Let $\theta_1 = \{\alpha^p,\ \lambda, \sigma_{\alpha^{age}}\}$ be the nonlinear and $\theta_2 = \{\alpha^x\}$ be the linear parameters, then $\theta = \{\theta_1, \theta_2\}$ are all the parameters to estimate. The GMM estimator is given by:

$$\hat{\theta} = \arg\min_{\theta}\ G\left(\theta\right)' W^{-1} G\left(\theta\right) \tag{20}$$

where $G\left(\theta\right)$ is a vector of stacked moments and W is the weighting matrix.[18] The computa-

[16] Having two market shares to match for the same product, I can pin down two vectors of error term: ξ_{jt} which allows me to match the market share of purchasing and τ_{jt} which allows me to match the market share of holding. See the discussion in the subsection 3.2.1.

[17] This assumption is also discussed in BLP (1994) and used by Sweeting (2007) and Lee (2009). An alternative approach would have been to assume that the demand shock ξ_{jt} is orthogonal to the observable product characteristics and the standard excluded instruments suggested by BLP and based on rival characteristics. However, there are no firms producing second-hand goods. Hence the intuition behind instruments that are excluded from demand but included in supply to identify the demand curve cannot be found in used-goods markets. I thank the referee for this discussion.

[18] The diagonal elements of the weighting matrix should be the inverse of the variance of the moment. For variance of micromoments, I use $(p)(1-p)/N^o$, where p is the value of the moment in the data, and N^o is the number of consumers elegible for the subsidy. As this variance is very small, our weighting matrix puts a high weight on the micromoments so our estimation algorithm attempts to match these very closely. See BLP (2004) and Petrin (2002) for more details on calculating weighting matrices when combining micro moments with aggregate moments.

tion of the objective function requires knowledge of the weight matrix, W, which in general requires knowledge of either the true value of the parameters or consistent estimates of these. There are several solutions to this problem. I follow Nevo's (2000) two-step approach: I first assume homoscedastic errors and therefore the optimal weight matrix is proportional to $Z'Z$. I can then compute an estimate of the vector θ and use this estimate to compute a new weight matrix to perform the second and final estimation of the parameters. As discussed in Nevo (2000), θ_2 can be expressed as function of θ_1, therefore the nonlinear search is performed only over θ_1 using a non-derivative based Nelder and Mead (1965) simplex algorithm.

3.2.1 Computation

This section outlines the algorithm used to jointly estimate the parameters of the model and the distribution of transaction costs. Using an approach similar to GR (2009), I combine BLP's (1995) procedure to recover the unobserved product characteristics ξ_{jt} and transaction costs τ_{jt} with Rust's (1987) fixed point algorithm to solve consumers' dynamic optimization problems. Once ξ_{jt} is recovered, the objective function in (20) can be computed. Hence, the nonlinear search is performed to recover the parameters of the model. Figure 4 shows an overview of the computation algorithm described in detail below.

For a given vector of θ_1 and a set of random draws, the mean net augmented utilities $\hat{\phi}_{jt}$ and the transaction costs τ_{jt} which rationalize predicted market shares to observed market shares of consumers' holdings and of consumers' purchases, are found via the contraction mappings similar to those in BLP. For each iteration of the mappings, consumer beliefs over the evolution of the logit inclusive value δ_{it+1} and the net augmented utility flow ϕ_{ijt} are updated. More specifically for each iteration of the BLP mappings (hence for each value of $\hat{\phi}_{jt}$, τ_{jt}) and given the nonlinear parameters θ_1, the set of simulated draws is used to calculate the logit inclusive values as in equation (10), and net augmented utility[19] as in equation (9). Both these variables are used to estimate the coefficients of the Markov process regressions in (11) and (12). These coefficients are then used to construct the transition matrix and to calculate the expected value function from (13) and (14) by iteration (Rust, 1987). Hence individual probabilities are computed as in (15) and (16). The number and identity of consumers for each product available on the market evolves according to individual probability of buying or keeping product j predicted by the model. The individual probabilities are aggregated as in (17), (18) and (19) to form predicted market-level purchase and holding probabilities. Finally, the aggregate shares are used to update $\hat{\phi}_{jt}$ and τ_{jt}. The procedure iterates until $\hat{\phi}_{jt}$ and τ_{jt} converge at which ξ_{jt} is recovered from the final value of $\hat{\phi}_{jt}$ via linear regression.[20] Then the objective function in (20) is computed and the nonlinear search

[19] Specifically $\phi_{ijt} = \hat{\phi}_{jt} + \sum_j x_{jt}\, \sigma v_i - \alpha_i^p (p_{jt} - \beta E_t [p_{jt+1}])$.

[20] A potential alternative approach to utilizing multiple nested-fixed point routines includes Mathematical

of θ_1 is performed using the Nelder and Mead simplex algorithm to recover the parameters of the model.

The main innovation with respect to the algorithm in GR is the use of two BLP-type contraction mappings to recover $\hat{\phi}_{jt}$ and τ_{jt}.[21] Berry (1994) suggests inverting the product market shares to recover the implied mean utilities for each good. Hence, given K different market shares it is possible to recover K different mean utilities. In my setting, I observe $2K$ market shares for K used products: I observe for each used car the market share of consumers' purchase and market share of consumers' holding. Hence I invert these shares and I can recover the mean (net augmented) utilities and the transaction costs for each good. The mean net augmented utilities $\hat{\phi}_{jt}$ which rationalize predicted market shares to observed market shares of consumers' holdings, are found via iteration of the following equation,

$$\hat{\phi}_{jt}^{'} = \hat{\phi}_{jt} + \psi_1 \left(\ln \left(\check{s}_{jt}^{H} \right) - \ln \left(\widetilde{s}_{jt}^{H} \left(\hat{\phi}_{jt}, \tau_{jt}, \theta \right) \right) \right) \tag{21}$$

The car mean utilities net of the transaction costs $(\phi_{jt} - \tau_{jt})$ rationalize the predicted market shares to the observed market share of consumers' purchases,

$$(\hat{\phi}_{jt} - \tau_{jt})' = (\hat{\phi}_{jt} - \tau_{jt}) + \psi_2 \left(\ln \left(\check{s}_{jt}^{D} \right) - \ln \left(s_{jt}^{D} \left(\hat{\phi}_{jt}, \tau_{jt}, \theta \right) \right) \right) \tag{22}$$

Having computed $\hat{\phi}_{jt}$ from equation (21), I can recover τ_{jt} from the previous equation by difference. $s_{jt}^{D}\left(\hat{\phi}_{jt}, \theta\right)$ and $\widetilde{s}_{jt}^{H}\left(\hat{\phi}_{jt}, \theta\right)$ are computed from equations (17) and (18) and ψ_1 and ψ_2 are *tuning* parameters while $\check{s}_{jt}^{H}$ and $\check{s}_{jt}^{D}$ are the corresponding shares observed in the data. I have found that the speed of convergence of equation (21) is higher than (22), and to avoid instability in the convergence process I set $\psi_2 = (1 - \beta)^2$ and $\psi_1 = 1 - \beta$.

As in GR, there is no proof of the existence of a unique fixed point. However, no problems with convergence or multiple solutions were encountered.

More computational details. To perform the iterative calculation, I discretize the state space $(\phi_{ikt}, \delta_{it})$ and compute the transition matrix following Tauchen (1986). Specifically, I compute the value function by discretizing ϕ_{ikt} into 20 evenly-spaced grid points and δ_{it} into 20 evenly-spaced grid points and allowing 400 points for the transition matrix. I specify that δ_{it} and ϕ_{ikt} can take on values from 15% below the observed values to 15% above. I have examined the impact of easing each of these restrictions and found that they have very small effects on the results.

Program with Equilibrium Constraints (MPEC), which has been shown to yield computational advantages in related problems (Su and Judd, 2008; Dube, Fox, and Su, 2008)

[21] The other innovation is due to the presence of two state variables which characterize consumers' dynamic optimization problems. Consequently the transition matrix is computed based on the estimated coefficients of two Markov process regressions, (11) and (12).

I assume that η_{it} and μ_{it} are uncorrelated[22] to reduce the computational burden of the model in constructing the transition matrix.

Since the estimation algorithm is computationally intensive and computational time is roughly proportional to the number of simulation draws, I use importance sampling to reduce sampling variance, as in BLP and GR. Finally, instead of drawing i.i.d. pseudo-random normal, I use Halton sequences to further reduce the sampling variance. In practice, I use 80 draws. Results for the base specification do not change substantively when I use more draws.

3.2.2 Identification

Here I present a heuristic discussion of the intuition for identification. As discussed in the introduction, the persistence in demand is driven by the presence of transaction costs. The key assumptions are: cars depreciate after every year and consumer preferences are perfectly persistent, therefore a consumer who faces no frictions will always prefer after one period to resale the car to upgrade to his preferred quality.

The parameters in the utility function, a and σ, are identified analogous to BLP. Among people that choose to buy any car, I look at the share that choose each product. As the set of available cars on the market and their prices change, market shares change. The extent to which consumers are attracted to any particular characteristic in the x vector identifies α. The extent to which they substitute from products with similar x variables identifies σ. For example, if the price falls for a particular car and it attracts the market share from similar cars, we will find that σ is large. That is, I match this feature by saying that there is substantial consumer heterogeneity in preferences. If the product attracts market share from a diffuse set of products, that is similar to the standard logit model, and I find that σ is small. Note that as in GR but different from BLP, my model makes use of substitution across time periods. For instance, a price decline in this period leads to low sales for similar products in the next period, that also leads us to find that σ is large.

The challenge that is unique to my model is to separately identify the product unobservable characteristic, ξ_{jt}, from the unobservable transaction costs, τ_{jt}. To separate these two unobservable components, I use the information from the two shares: the share of consumers choosing to hold a given car type each period, and the share of consumers choosing to purchase the same car type in the same period.

As the car depreciates there will be consumers who would like to re-optimize and choose

[22] I have examined this assumption using the Monte Carlo technique. The correlation between these two error terms converges rapidly to zero as J increases. In particular, they are uncorrelated for $J > 150$ which is the minimum number of products available in the market and used for the estimation. The Monte Carlo results are provided upon request.

a different alternative that better matches their taste. However, the replacement decision depends on the size of the transaction costs; the maximum utility that a consumer can obtain by replacing the car is reduced by the size of the transaction cost he will pay. With full information, prices reflect the deterioration in quality. Conditional on replacing the car, $\xi_{jt} - \tau_{jt}$ is identified from demand among consumers that choose to buy good j. In my model, purchase behavior identifies the mean utilities of products, net of transaction costs. To identify ξ_{jt}, I use a separate vector of market shares: the share of consumers who choose not to sell their car but to keep it at the beginning of each period. The intuition is that comparing the mean utilities of what is available on the market, to those of the products that consumers hold, would predict a much larger set of sales than we see in the data. This discrepancy is explained by the transaction costs.[23] More generally, a model without transaction costs implies consumers never want to hold goods for more than one period, so the relative distribution of sales and holdings identifies the distribution of transaction costs.

This argument is heuristic. Since all of the elements of the model are solved for simultaneously, all of the variation in the data contributes to the identification of each parameter. For instance, transaction costs are, in part, determined by purchase decisions: when consumers make a purchase, they rationally predict the transaction costs they will realize when they eventually sell it.

Discussion on assumptions and simplifications. The tractability of the estimation is achieved on the assumption that the logit error term is i.i.d. across time, individuals and products. One would expect there to be some factors that are not observed by the researcher that affect each of the decision makers' choices. Random coefficients generate persistent unobserved heterogeneity over time, alleviating this problem. Moreover, the random coefficients reduce the undesirable features of the IIA of the logit model. The logit errors provide a possible source of the lack of resale. It is technically possible for a consumer to realize a sequence of logit errors such that she does not want to sell her car. However, this will not be sufficient to explain consumers holding behavior. If I do not account for the presence of transaction costs, the comparison among mean utilities of what is available on the market to those of the products that consumers hold would predict a much larger set of sales than what is observed in the data.

An issue to deal with is the initial distribution of consumer types across different car types. To account for the initial distribution, I estimate a static random coefficient model without transaction costs. Then, I use the resulting distribution of consumer types as the initial distribution of consumers across different car types for the full model estimation. I

[23] From the perspective of BLP, one might view τ_j as coming from a set of dummies in the outside option for the car that the consumer holds.

have tried different options[24], the results are quite similar.

Finally, over the 11 years which I consider, there has been the introduction of the scrappage policy in 1997 and 1998 which potentially requires the introduction of a third state variable in the model. I do not introduce it for two reasons: first the policy was introduced in 1997 for a few months and then renewed again in 1998 for a few months with stringent requirements in the way to benefit of the policy. The subsidies were awarded to consumers who had owned a car for at least one year. First, this requirement restricts the possibility that consumers could have modified their replacement behavior, in advance, to take advantage of a law that was not issued yet. Second, data availability and computational costs. An introduction of a third state variable would have required a richer specification than (4) with more parameters to estimate using only 11 points given the 11 years observed.

Instruments. Valid instruments must be correlated with the regressors but uncorrelated with the time t unobservable innovation. Although observed product characteristics may be endogenous with respect to unobserved characteristics ξ_{jt} and I assume that – with the exception of prices – these observed characteristics will be exogenous with respect to changes in these unobserved characteristics.

The innovation in the unobserved characteristics can be expressed as pseudo-differences in the mean utilities $\hat{\phi}_{jt}$, more specifically:

$$\varsigma_{jt} = (\hat{\phi}_{jt} - \lambda\hat{\phi}_{jt-1}) - \alpha^x(x_{jt} - \lambda x_{jt-1})$$

where $\hat{\phi}_{jt}$ are the mean net augmented utilities for each product j at time t. Hence, pseudo-differences in x_{jt} are used as instruments to identify α^x. These pseudo-differences are valid instruments if consumers cannot predict the future value of ς when making their decision at time t.

Since current prices may be correlated with these innovations in product unobservables, I will use lagged prices as an instrument, p_{jt-1}. Lagged prices are valid instruments as long as the price of new goods and price of used cars are respectively set by firms, or determined in the secondary market without accounting for future values of ς which cannot be forecasted either by firms or consumers. Following the same logic, I also used as instruments the initial stock for each model at the beginning of each period, s_{ijt-1}, and the market share of purchased product in the previous period, s^D_{jt-1}. Finally, the lagged value of $\hat{\phi}_{jt}$ is also used as an instrument to further help identify λ.

[24] I have used the initial distribution implied by the dynamic model without transaction costs.

4 Results and Implications

4.1 Parameter Estimates: utility specification

Table 6 reports the parameter estimates associated with the characteristics of the cars as in the utility specification. Multiple specifications are provided: columns (1) and (2) report the estimates of the full dynamic model respectively, with and without micro moments; column (3) reports the estimates of a dynamic model without the transaction costs; column (4) the estimates of the static model.

By looking at the first column, signs of coefficients are as expected, with utility decreasing from the price and the age of the car. The price coefficient is estimated non linearly and the magnitude is -53.52. A consumer obtains a positive flow utility from owning a car (relative to the outside option) with a mean constant term of 4.04. The age of the car reduces the utility. The heterogeneity in preference for age among consumers is captured by σ^{age}. The coefficient on engine size of 2.89 shows that consumers prefer cars with a higher cc engine. Dummies for location suggest consumers' preference for German cars.[25] The dummy on fuel shows that people prefer a gasoline rather than a diesel engine. The positive coefficient on the fuel dummy interacted with time trend is capturing the increasing utility over time to buy diesel cars. Over the considered time window, there is a substantial reduction in the taxes owed to the government, especially for diesel engine cars; the model is able to capture the increasing appeal for these vehicles due to this tax reduction.[26] Future resale prices are needed to obtain the rental price each year, I use observed future prices as a proxy for the expected ones. As in assumption 2, unobservable product characteristics for each automobile evolve according to a first-order autoregressive process, λ is estimated to be 0.39 which shows a significant persistence in the unobservable over time. The micro moments improves the idenfication of the consumer heterogeneity and the price coefficient: σ^{age} becomes significant in the specification with micro moments.

Column (3) provides estimates from the dynamic model where no transaction costs are paid to replace the automobile, hence there are no frictions in the market. The dynamic model without transaction costs has a simple analytical solution. For each consumer i the probability of choosing alternative j given $k \in J_{t-1} \cup \{0\}$ is

$$d_{it}^{kj} = \frac{\exp(x_{jt}\,\alpha_i^x + \xi_{jt} - \alpha_i^p p_{jt} + \alpha_i^p p_{kt} \cdot I(j \neq k) + \beta E_t\left[EV_{ij}\left(...\right)\right])}{\sum\limits_{j \epsilon J_t \cup \{0\}} \exp(x_{jt}\,\alpha_i^x + \xi_{jt} - \alpha_i^p p_{jt} + \alpha_i^p p_{kt} \cdot I(j \neq k) + \beta E_t\left[EV_{ij}\left(...\right)\right])}$$

If there are no transaction costs, the problem is no longer state dependent and $E_t\left[EV_i\left(k, .\right)\right] =$

[25] The higher quality of new and used cars produced in Germany is in line with the findings of Emons & Sheldon (2003) .

[26] In particular, the property tax fell progressively by more than 50%.

$E_t\left[EV_i\left(0,.\right)+\alpha_i^p p_{kt+1}\right]$.[27] Replacing the previous equality in the discrete choice probability and simplifying it, I can write:

$$d_{it}^{kj}=d_{it}^{j}=\frac{\exp(x_{jt}\,\alpha_i^x+\xi_{jt}-\alpha_i^p(p_{jt}-\beta E_t\left[p_{jt+1}\right]))}{\sum\limits_{j\epsilon J_t\cup\{0\}}\exp(x_{jt}\,\alpha_i^x+\xi_{jt}-\alpha_i^p(p_{jt}-\beta E_t\left[p_{jt+1}\right]))}\ \forall k\in J_{t-1}\cup\{0\}$$

The probability of purchasing any good j does not depend on the car owned. It is similar to the static model but for presence of the expected price in the flow utility function as in the dynamic model with transaction costs. The very imprecise price coefficient (which becomes insignificant) and the unexpected sign on the characteristics that enter the mean utility, suggest that the data cannot easily be explain by a dynamic model where consumers are allowed to frequently replace their goods.

Column (4) reports the estimates from the static model with random coefficients when consumers choose between different types of new and used cars and they do not face any dynamic decisions and they do not pay any transaction costs. The price coefficient drops consistently with respect to the full dynamic model with transaction costs, and the coefficient estimates attached to the fuel dummy, and the fuel dummy interacted with time lose their statistical significance. These parameters reflect some dynamic consideration that the static model is not able to capture. Other coefficients seem plausible and have the same signs as in the full specification. The random coefficient attached to the consumer preference for the age becomes insignificant.

4.1.1 Price elasticities

In Table 7 and 8 I present the average[28] own- and cross-price long-run elasticities simulated from the full dynamic model and I compare them with the static price elasticities and the dynamic model without transaction costs.[29] I calculate the long-run elasticities using permanent changes in the prices of a product. In particular, I allow for a permanent change in the price of a new model, as well as the future price of the same model in the used market in the following years[30], keeping the percentage of the depreciation in price of the same model across different ages unchanged. Permanent changes capture the long-term effects of the change on the consumer's expectations. The elasticities were simulated for the dynamic model as follows. First, I use the observed quantities to solve the consumer problem and

[27] Notice that $EV_i\left(0,.\right)$ no longer depends from the good held by consumers and it is constant across products

[28] Market shares are used to weight the price-elasticities.

[29] Price elasticities for the dynamic model without micro moments are similar in magnitude to the ones with micro moments.

[30] For example, I increase the price of new Fiat compact in 2000, the 1-year old Fiat compact in 2001, the 2-year old Fiat compact in 2002 and so on.

estimate a baseline level of demand. Second, I generate a permanent change in the price path of each product which will affect both the new and used market. I then re-solved the dynamic model for the optimal consumers' behavior, allowing the consumers to update their beliefs. Finally, I simulated new choice probabilities, using them to compute the change in choice probabilities relative to the initial values, and to compute the price elasticities. The reported estimates are the average price own-price elasticities for new products in 1995, distinguished by market segment and country of origin. The average cross-price elasticities for new products are reported in Table 8. The cross-price elasticities are reported within market segment and for different car ages.

Table 7 reports the average own-price elasticities. The average elasticity for the static model is -3.05. Using a price change of 5%, the estimates show that the myopic model underestimates price elasticities (in absolute value) on average by 140%. Myopic consumers do not consider the future utility of owning a product, as well as the possibility for consumers to timing their purchases, which leads to a downward biased price, hence myopic consumers underreact to a permanent price change. The own-price elasticities computed from the static model and the full dynamic model are all greater than one in absolute value. Price elasticities of the dynamic model without transaction costs are significantly smaller than the static model, due to the bias in the price coefficient which is about 4 times smaller in magnitude than the static model and not significatively different from 0 (see Table 6). As for the static model, the dynamic model without transaction costs does not rationalize the possibility of consumers to wait and timing their purchase.

Table 8 shows the average value of the cross-price elasticities after a permanent change in price. The first row of Table 8 reads as follows: after a 1% permanent increase in the price of new cars belonging to the small car segment, there is on average a 17% increase in demand of new cars within the same segment (first column), a 10% increase in demand of 3-year old cars within the same segment (second column), a 0.7% increase in demand of 6-year old cars (third column) in the same segment, and so on. It is interesting to observe that the older are the cars within each segment the lower are the cross-price elasticities: older cars are poorer substitutes for the new ones. Differently from a static model (and from the dynamic model without transaction costs), the dynamic model (with transaction costs) can generate negative cross-price elasticities. The dynamic model by explicitly solving the sequence of consumers' decisions, endogenizes the distribution of consumer holdings across different types of vehicles and time. A permanent increase in prices of new and used cars will reduce the appeal for consumers to wait and replace their car (or buy one) in the near future, which is reflected in a lower continuation value. In particular, the model generates negative cross price elasticities mostly for older cars as consumers switch from buying older cars towards buying newer cars which will be held for longer periods.

4.2 Transaction costs

Transaction costs in my model (in the absence of asymmetric information) measure all possible frictions (i.e. taxes, search costs, dealer compensations etc.) which consumers incur upon replacement. I estimate the whole distribution of transaction costs. Specifically, I estimate the average cost consumers pay to purchase a car j in period t. The monetary value of the transaction costs is obtained by dividing the estimated transaction costs by the average price coefficient, once integrated over the income distribution.

As can be seen in Figure 5, the magnitude of transaction costs declines over time.[31] The effect is the result, among other factors, of a progressive reduction of the taxes paid upon the transaction and a reduction of the interest rate due to the introduction of the European currency. The average transaction cost was about €3000 in 1994 decreasing to €2000 in 2004, the average standard deviation across time is about €580. The distribution of transaction costs is shown in Figure 6; it shows a peak in the level of transaction costs between €1900 and €2600. The minimum level of the cost is about €1000.

It is also instructive to look at the distribution of the transaction cost/price ratio across different models and different ages. Figure 7 shows that there is a peak between 20% and 40%, and most of the models show a level of transaction costs between 10% and 80% of the respective level of prices.

Finally, I use the panel aspect of the data to track how transactions costs vary for car models over time. I take the car models that I observe for more than 5 years and I compute the standard deviation around the mean transaction cost for each one of them, the average standard deviation across these models is about €390. This seems a moderate fluctuation associated with the changes in the car age, the evolution of the market and the presence of aggregate shocks.

It is important to notice that the estimates do not refer to the costs that are actually paid upon transaction, but rather to the costs of a hypothetical purchase of a particular car j. In the model, people choose to buy a car only when the payoff shocks are favorable. The unexplained part of the utility flow, ϵ_{ijt}, may be viewed as either a preference shock, or a shock to the cost (or both), with no way to distinguish between the two. The net cost paid upon a transaction is therefore less than the amounts reported above.

External validation. The level of transaction costs explains the high persistence in the stock of cars held by consumers . The results imply that on average, a consumer keeps her automobile for about 7 years. This result is obtained without accounting for the truncations

[31] The estimate of the transaction costs are relative to used cars only. For new cars, I cannot identify the size of the transaction costs. In the estimation procedure, I assume that whoever purchases a new car pays taxes and other costs of registration as specified by *Quattroruote*. These costs vary between €350 and €800 according to the type of the new vehicle purchased.

in the data. Are these figures reasonable? According to the information published in the magazine *Quattroruote* in 1998, the explicit costs to sustain upon a transaction of a used car varies between €1000 and €4000. The composition of these costs is the following: financial costs about €400; *Quattroruote* reports that on average, the money borrowed to buy a used car in 1998 was €5000 and the spread over a safe interest rate was about 8%. The taxes and expenses to pay upon the transaction varied between €340 and €1600 according to the size and the type of cars. The dealer compensation for trading a used car also varied between €300 and €2000 according to the model. In addition, one has to account for the hidden costs like search costs, asymmetric information and so on. The above analysis confirms that the estimations of the model seem to have the right magnitude and transaction costs as would be expected, to play a substantial role in a consumers' replacement decision. In Table 10, I compare the transaction cost estimates with the taxes and the dealer compensations as reported in *Quattroruote* relative to few models. The difference can be explained by the presence of financial costs, search costs and similar costs also sustained by the seller of a used car.

Decomposition of transaction costs. Next I try to investigate the composition of the transaction costs in more detail. Table 9 reports the parameter estimates of transaction costs regressed on a set of variables. We can observe that the coefficient associated with the stock of each car type in percentage terms is negative and highly significant. This result indicates that having more cars in the market reduces the costs associated with finding the right match. More specifically, an increase of 1% in the stock of cars available reduces the transaction costs by €310. This relation captures one of the essential characteristics of a decentralized market: traders must incur costs to search for trading opportunities. Thinner markets cause higher search costs. Instead, the matching between buyers and sellers becomes easier in a thicker market where larger stocks of cars are available. In this sense, cars with a thicker market are more liquid. The reason is that cars with a thin market are more difficult to sell, and they have higher option values: consumers choose to hold on to them for longer periods. Hence as expected, the transaction costs decrease in the stock of each type of car available. Moreover, the effect of trading frictions transmits to transaction prices by decreasing on average their level, i.e. cars with lower transaction costs have higher demand and higher levels of price.

The variable *Diesel*Time trend* captures the reduction of taxes over time, relative to the car with Diesel engines as discussed above. The costs are increasing in the engine displacement, as higher taxes and fees are usually associated with bigger cars. Notice the transaction costs display a decreasing trend over time, confirming that the used car market has become more active since 1994. This is consistent with the information displayed in Figures 1 and 2. The effect is the result of a progressive reduction in the taxes to pay upon registration, the enhancement of Internet transactions and the introduction of the Euro — and the consequent reduction of the interest rate and transaction costs across EU countries.

Finally, there is a negative coefficient associated with the FIAT dummy which may reflect the presence of a dense network of FIAT dealers as well as a lower maintenance costs associated with the national manufactured cars that reduces the risk and cost of buying a used vehicle.[32]

The range of transaction costs is in line with evidence found elsewhere and it varies with vehicles and market characteristics in an intuitively plausible way, lending support to validity of the estimates

5 The Scrappage Policy

Scrappage subsidies have been particularly popular in the European Union (EU), as well as in the United States and Canada. These policies were aimed to reduce pollution by forcing an early retirement of old and polluting cars but they were also aimed, in some cases, at stimulating the national car industries. Typically, these subsidies were between €500 and €1,500 and eligibility to participate in the program was a function of the vehicle's age (e.g. the automobile must be 10 years old or older). During the 1990s, most EU countries offered scrappage subsidies. France, Greece, Hungary, Ireland, Italy and Spain required that to be eligible for these subsidies, the replacement vehicle had to be new. These policies are called *cash-for-replacement* schemes. On the other hand, Denmark and Norway as well as the United States and Canada, did not impose any constraints on the type of replacement vehicle — they followed a *cash-for-scrappage* scheme.[33]

There has been a debate regarding the overall effects of these policies on car markets and consumers' welfare, especially considering that these programs could be expanded in scope and duration. The model can help understanding their implications and effects. In particular, in this section, I study the effect of the scrappage program implemented in Italy in 1997 and 1998 which is summarized by Table 11. When the policy was implemented there were about 34,000 potential cars elegible for the subsidy in the region and about 2,400 were actually scrapped to get the discount on new vehicles.

[32] In the transaction costs regression I allow for the transaction costs to be correlated with the unobservable characteristics of the car-type.

[33] See European Conference of Ministers of Transport Publications (1999), EPA (1998) and Hahn (1995) for a comprehensive description of the different scrappage subsidy programs in the United States and Europe.

5.1 Policy Evaluation

Using the framework developed and the estimates obtained from the previous sections, I proceed to examine the impact of the replacement scheme implemented in Italy. The model with micro-moments replicates closely the impact of the scrappage policy on the number of new vehicles purchased with the subsidy. In theory, policies that subsidize the replacement of old cars operate through the optimal scrapping age and the requirement in terms of replacement choice: *cash-for-replacement* schemes required that to be eligible for these subsidies, the replacement vehicle had to be new, whereas *cash-for-scrappage schemes* did not impose any constraint on the type of replacement vehicle. The initial effects of these policies depend on the fraction of cars older than the (new) optimal scrapping age. The subsequent effects then reflect the evolution of the cross-sectional distribution once the policy change has occurred. The idea that the policy will create an incentive to scrap older cars is obvious. The point of doing the analysis is to quantify both the extent of the immediate incentive to replace and the subsequent effects of this policy on new and used car sales as the distribution of car ages evolves. The contribution of my model is to illustrate and study the effect of the scrappage policies on sales and government revenues accounting for richer specifications, where new and used car markets interact, and consequently to evaluate richer design policies. I find that these policies will boost sales of new cars in the short run but, at the same time, set in motion variations in the cross-sectional distribution of car ages that create rich long-run effects. In particular, bursts of activity associated with temporary scrapping subsidies are short-lived: car production is reduced in future periods. In contrast, different policies may have different effects on used car sales. *Cash-for-scrappage* schemes produces qualitatively similar effects as for the new car sales whereas *cash-for-replacement* schemes has a more long lasting effect of reducing the sales of used cars.

To evaluate the cash for replacement scheme in terms of my model and study the dynamics and implications of the new and used automobile markets, I perform the following analysis. I first simulate a baseline situation with no subsidies offered and I use it as the benchmark. Hence I compare the new and used car sales and revenues (both in the short run and long run) obtained from the baseline model with the results obtained from 3 different policy scenarios. The 3 policy scenarios I consider are the following: the *cash-for-replacement* scheme with the same requirements and subsidies as those adopted by the Italian Government; the *cash-for-replacement* scheme with same level of subsidies as before but the elegible consumers must have a car older than 8 years (rather than 10); and the *cash-for-scrappage* scheme with the subsidies fixed to 25% of the one implemented by the Italian Government, for cars older than 10 years.

Figure 8 displays the aggregate sales of new cars for the different schemes compared with the baseline model. The simulation shows that these policies burst the aggregate sales of new cars followed by a contraction in sales. Sales remain lower for the next few years. The bigger is the short-run effect of the subsidy in expanding the demand, the bigger is the contraction in sales that follows in the future. In particular, the cash for replacement 8 and 10, increases the new car sales by 126% and 97% in 1997 and 69% and 51% in 1998 respectively. It follows a contraction in sales of 11% and 6% respectively in 1999 which remain lower than the sales in the baseline model until 2003, with a total contraction of 25% and 16%. The cash for scrappage policy has a much smaller impact on the sales on new cars with a total increase of 9%, which is followed by a total reduction in sales of 6%. The effect lasts until 2003. In 2004, there is a slight increase in the number of new cars bought as the generation of cars from the policy reform are starting to be replaced/scrapped. Across all policies, the contraction in sales is smaller in magnitude compared to the increase due to the implementation of the scrappage scheme. The model accounts for the fact that consumers using the subsidy to buy a new cars are the ones who would most likely have bought either a used car, kept the same one or chosen the outside option.

This finding is reinforced by looking at the used-car market sales. Figure 9 shows that the *cash-for-replacement* schemes reduce the sales in the used car market and this negative effect on sales is more persistent than the contraction of the new car sales. As I discussed above, these policies are aimed at consumers with low income and/or low sensitivity to the car age. Hence they reduce the demand for used cars in the short and long run. The *cash-for-replacement 8 years* has a more negative effect than the 10 years one. Both schemes determine a contraction in sales of respectively 23% and 27%. In contrast, the *cash-for-scrappage* bursts the sales of used cars as well as the sales of new cars with a bigger effect on the first group (63% vs 9% in two years). After the expansion in sales, there is a contraction on the used market with a total reduction in sales of 9%.

Compared with the study of Adda and Cooper (2000), my model would predict a smaller contraction in new car sales following a replacement subsidy. In their model, consumers are homogenous and there is no active second-hand market hence the subsidy leads consumers to anticipate their replacement decision causing a contraction in future sales of new cars. In my model, the subsidy will affect more consumers who would have purchased a used car otherwise, hence the contraction in sales is more evident in the secondary market rather than in the primary market. Moreover, my model provides a more complete framework to evaluate policy design according to their requirement decisions and the consequent impact on new or used car markets and revenues.

The second exercise is to study the effect of these policies on the Government revenues. The evaluation should account for both the short term effect and for the long term dynamics. On the one hand, the Government supports a cost in implementing the scrappage policy equal to the subsidies disbursed, on the other hand, the Government collects the V.A.T. on the new car sales which was 19% in 1997 and 20% from 2000 until 2004. Moreover, the Province (the local government) collects a tax for each purchase of a new and a used car equal to €77.47 in 1997 and to €150.81 from 1998 until 2004. Hence the evaluation of the policy should also take the direct redistribution of revenues from the central to the local government into account.

I consider the same policies implemented above. Figure 10 shows the net effect on government revenues compared with the baseline option for the three policies. Figure 11 displays instead the net change in revenue (compared with the baseline model) collected by the Province. Finally, Table 12 reports the present discounted value of the Government revenue, of the Province revenue calculated for the baseline model, as well as the changes due to the implementation of the various policies. On average, the cash for replacement policies increase the revenue of the central government, whereas the cash for scrappage scheme determined a fall due to a lower impact on new car sales and the consequent reduction on the V.A.T. collected, along with the larger number of subsidies awarded. Even if the *cash-for-replacement 8 years* determine a bigger increase in the revenue in the short run, compared with the *cash-for-replacement 10 years,* the net present discounted value of revenues under the two policies are similar, due to a bigger contraction in sales under the 8-years policy. At the local government level, the impact is exactly the opposite since the cash for scrappage policy causes an overall greater jump in total sales.

Notice that the *cash-for-scrappage* scheme has similar effect on the automobile market due to the reduction of transaction costs of the same magnitude. Hence the Government could achieve the same results by temporally reducing the taxes to pay upon transactions, or by making the secondary market more liquid by improving competition among dealers and/or reducing the search costs.

The previous simulations show that these policies have quite a different impact on the number of used cars scrapped and on the number of new cars purchased. A more complete evaluation of the costs and benefits of the scrappage policies should account for the emission reductions, and the change in unemployment rate which is a potential area for future research.

6 Concluding Remarks

This paper presents a structural model of dynamic demand for automobiles that explicitly accounts for the replacement decision of consumers in the presence of a second-hand market. The model incorporates the feature that consumer replacement is costly due to the presence of transaction costs. In addition, it allows for rational expectations about future product attributes, heterogeneous consumers with persistent heterogeneity over time and endogeneity of prices. The data set that I use for the estimation provides information about sales for individual cars over time as well as information about prices and characteristics of cars. The estimation of the transaction costs is achieved from the difference between the share of consumers that choose to hold a given car type each period and the share of consumers that choose to purchase the same car type in each period. The dynamic aspect of the model and the presence of transaction costs are essential to explain the sales pattern in the primary and in the secondary market. If these costs were ignored, it would not be possible to explain the high persistency in the stock of cars held by consumers.

The model is particularly useful in analyzing the effects of policies directed at modifying the replacement decisions that in turn have an impact on the overall distribution of vehicle holdings. My approach highlights the quantitative response to individual agents to policy variations in a model where new and used markets interact. Furthermore, the consequent evolution of the cross-sectional distribution creates persistent effects of the policies. Consequently, policy analysis is much more difficult in this setting since the evolution of the cross-sectional distribution must be taken into account. That said, the analysis illustrates that taking these dynamics into account is feasible and instructive for both the design and evaluation of policies.

References

[1] Adda, J., & Cooper, R. (2000). Balladurette and Juppette: A Discrete Analysis of Scrapping Subsidies. Journal of Political Economy, 108, 778-806.

[2] Anderson, S. and V. Ginsburgh. 1994. "Price discrimination via Second-Hand Markets" European Economics Review. January, 38:1 pp. 23-44.

[3] Berry, S. (1994). Estimating discrete choice models of product differentiation. RAND Journal of Economics, 25, 242–262.

[4] Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile prices in market equilibrium. Econometrica, 63, 841–890.

[5] Berry, Steven, James Levinsohn, and Ariel Pakes. (1999). Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy. American Economic Review 89(3): 400-30.

[6] Berry, S., Levinsohn, J., & Pakes, A. (2004). Estimating differentiated product demand systems from a combination of micro and macro data: The market for new vehicles. Journal of Political Economy, 112, 68–105.

[7] Bresnahan, T. F. (1981). Departures from marginal cost pricing in the American automobile industry: Estimates for 1997-1978. Journal of Industrial Economics, 17, 201–277.

[8] Bresnahan, F. T. (1987). Competition and Collusion in the American Automobile Industry: The 1955 Price War. Journal of Industrial Economics, June, 457-482.

[9] Brownstone, D. & Train, K. (1999). Forecasting New Product Penetration with Flexible Substitution Patterns. Journal and Econometrics, 89, pp109-129

[10] Carranza, J. (2007). Estimation of demand for differentiated durable goods. Unpublished manuscript, University of Wisconsin-Madison.

[11] Dube, J.-P., J. T. Fox, and C.-L. Su (2008). Improving the Numerical Performance of BLP Static and Dynamic Discrete Choice Random Coe cients Demand Estimation. Mimeo.

[12] Emons, W. & Sheldon, G. (2003). The Market for Used Cars: A New Test of the Lemons Model. Unpublished manuscript, University of Bern.

[13] Esteban, S. & Shum, M. (2006). Durable-goods oligopoly with a secondary markets. Forthcoming RAND Journal of Economics.

[14] EPA (1998). Transportation Control Measures: Accelerated Vehicle Retirement. Transportation Air Quality Center Technical Overview.

[15] European Conference of Ministers of Transportation (1999): Clearer Cars. Fleet Renewal and Scrappage Scheme. European conference of Ministers of Transport Publications, France.

[16] Gentle, J. E. (2003). Random number generation and Monte Carlo methods (2 ed.). New York: Springer. Unpublished manuscript, Boston University.

[17] Gilligan, T.W., (2004). Lemons and leases in the used business aircraft market. Journal of Political Economy 112, 1157–1180.

[18] Goldberg, P. K. (1995). Product Differentation and Oligopoly in Iternational Markets: the case of the U.S. Automobile Industry. Econometrica, 63, 891-951.

[19] Gordon, B. (2009). A Dynamic Model of Consumer Replacement Cycles in the PC Processor Industry. Marketing Science Vol. 28, No. 5, pp. 846-867

[20] Gowrisankaran, G. & Rysman, M. (2009). Dynamics of Consumer Demand for New Durable Goods. Unpublished manuscript, Boston University.

[21] Hahn, R. W. (1995). An Economic Analysis of Scrappage. RAND Journal of Economics, 26, 183-196.

[22] Hendel, I. & Lizzeri, A. (1999a). Interfering with Secondary Markets. Rand Journal of Economics. Spring, 30:1, pp. 1-21.

[23] Hendel, I. & Lizzeri, A. (1999b). Adverse Selection in Durable Goods Markets. American Economic Review, 89, 1097-1115.

[24] Hendel, I. & Nevo, A. (2006). Measuring the implications of sales and consumer stockpiling behavior. Econometrica, 74, 1637–1673.

[25] Hideo, K. & Sandfort, M. T. (2002). Existence of stationary equilibrium in the markets for new and used durable goods. Journal of Economics Dynamics & Control, 26, 1029-1052.

[26] Lee, R. S. (2009). Dynamic Demand Estimation in Platform and Two-Sided Markets, mimeo, NYU.

[27] Nelder, J. A., & Mead R. (1965). A Simplex Method for Function Minimization. Computer Journal, 7, 308-313.

[28] Melnikov, O. (2001). Demand for differentiated products: The case of the U.S. computer printer market. Unpublished manuscript, Cornell University.

[29] Nevo, A. (2000). A practitioner's guide to estimation of random coefficients logit models of demand. Journal of Economics & Management Strategy, 9, 513–548.

[30] Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. Journal of Political Economy, 110, 705–729.

[31] Porter, R. and Sattler, P. (1999). Pattern of Trade in the Market for Used Durables: Theory and Evidence. Working Paper no. 7149. Cambridge, Mass.: NBER.

[32] Rust, J. (1985): Stationary Equilibrium in a Market for Durable Assets. Econometrica, 53, 783-805.

[33] Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. Econometrica, 55, 999–1033.

[34] Schiraldi, P. (2009). Second-Hand Markets and Collusion by Manufacturers of Semi-durable Goods. Unpublished manuscript, London School of Economics

[35] Shcherbakov, O. (2008). Measuring switching costs in the television industry. Unpublished Manuscript, Yale University.

[36] Stolyarov, D. (2002). Turnover of Used Durables in a Stationary Equilibrium: Are Older Goods Traded More?, Journal of Political Economy, vol. 110(6), pages 1390-1413, December.

[37] Sweeting, A. (2007). Dynamic Product Repositioning in Differentiated Product Industries: The Case of Format Switching in the Commercial Radio Industry. Mimeo.

[38] Su, C.-L., and K. L. Judd (2008). Constrained Optimization Approaches to Estimation of Structural Models. Working paper, CMS-EMS, Kellogg School of Management.

[39] Tauchen, G. (1986). Finite state markov chain approximations to univariate and vector autoregressions. Economic Letters, 20, 177–181.

	Mean	St. Deviation
Population	74114	363.33
Income per Household	€21547	€3610
Family Size	2.70	1.23

Table 1: Province of Isernia from 1994 to 2004

	Mean	Std
Resale ratio	0.247	0.2862
N. Obs.		40148

Table 2 – Resale Ratio

Year	Vehicle type	Quantities sold	Average Price in Euros
1994	New	1297	13,000
	Used	2503	4,000
1995	New	1292	12,300
	Used	3014	4,100
1996	New	1259	13,200
	Used	2658	4,900
1997	New	2141	13,400
	Used	2713	5,150
1998	New	2195	13,350
	Used	2980	5,100
1999	New	3023	14,600
	Used	3272	5,050
2000	New	2086	14,200
	Used	3090	5,300
2001	New	2092	14,600
	Used	2336	5,400
2002	New	2244	16,700
	Used	3345	6,000
2003	New	1908	16,400
	Used	3462	6,750
2004	New	2111	17,050
	Used	3933	6,700

Table 3 – Descriptive statistics. Market shares are used to weight prices.

Estimation: Volume of trade (1994-2004)

Parameters		
Constant	-0.0091	(0.1665)
Depreciation	0.1617**	(0.0378)
Age	-0.0251**	(0.0015)
CC	0.1165**	(0.0121)
Diesel	-0.0093	(0.0096)
Fiat	0.0336**	(0.0098)
Model dummies	Yes	
Year dummies	Yes	

Standard errors in parentheses; statistical significance at 5% level indicated with **, and at 10% with *. R-squared 0.41. Obs. 1648

Table 4 – Resale Ratio Regression

Estimation: Volume of trade (2002-2004)

Parameters		
Transaction costs	-0.0434**	(0.0211)
Constant	-0.1175	(0.2659)
Depreciation	0.5527**	(0.1164)
Age	-0.0385**	(0.0053)
CC	0.0977	(0.0366)
Diesel	-0.0318	(0.2949)
Fiat	0.0727**	(0.0283)
Model dummies	Yes	
Year dummies	Yes	

Standard errors in parentheses; statistical significance at 5% level indicated with **, and at 10% with *. R-squared 0.27. Obs. 437

Table 5 – Resale Ratio Regression

Estimation Results: Utility

PARAMETERS	Dynamic Model with micro moments	Dynamic Model without micro moments	Dynamic Model with no Transaction costs	Static Model
Constant	4.04** (0.45)	3.92** (0.46)	-6.45** (0.57)	-4.25** (0.30)
Log(Age)	-2.29** (0.11)	-2.42** (0.11)	-0.76** (0.13)	-2.20** (0.09)
Engine size (CC)	2.89** (0.30)	3.05** (0.31)	-0.19 (0.24)	1.18** (0.15)
Fiat	0.10 (0.13)	0.12 (0.13)	0.48** (0.23)	0.62** (0.16)
German	0.51** (0.16)	0.53** (0.17)	-0.01 (0.23)	0.55** (0.15)
Diesel	-3.42** (029)	-3.47** (0.29)	0.41 (0.56)	0.11 (0.38)
Diesel*Time trend	0.33** (0.03)	0.34** (0.03)	0.01 (0.05)	-0.01 (0.04)
NON LINEAR PARAMETERS				
(Price- Expected Price), α_p	-53.52** (5.43)	-55.90** (12.91)	-2.04 (4.52)	-
Price, α_p				-9.32** (6.13)
Log(Age), σ_{age}	0.32** (0.08)	0.16 (0.45)	0.24 (0.88)	0.03 (4.64)
λ	0.39** (0.04)	0.38** (0.05)	0.80** (0.23)	0.75** (0.02)

Standard errors in parentheses; statistical significance at 5% level indicated with **, and at 10% with *

Table 6: Parameter Estimates

Own-Price Elasticities for New Cars (Year, 1995)

Market Segment	Dynamic Model w/Transaction Costs and micro moments		Static Model		Dynamic Model NO Transaction Costs	
Small car	**-3.75**		**-2.39**		**-0.34**	
Domestic		-3.73		-2.45		-0.40
Foreign		-3.78		-2.38		-0.33
Midsize car	**-4.14**		**-2.87**		**-0.52**	
Domestic		-3.99		-2.73		-0.47
Foreign		-4.24		-2.92		-0.54
Large car	**-4.60**		**-3.90**		**-0.74**	
Domestic		-4.40		-4.01		-0.76
Foreign		-4.95		-3.63		-0.68

Table 7: Own-price elasticity - Market shares are used to weight the price-elasticities.

Within Market Segment Cross-Price Elasticities (Year, 1995)

	Market Segment	**New**	**3-year old**	**6-year old**	**9-year old**
Full Dynamic Model	Small car	0.2226	0.1253	0.0383	-0.0559
	Midsize car	0.1110	0.0683	-0.0372	-0.0841
	Large car	0.0985	0.0809	0.0320	0.0093
Static Model	Small car	0.0225	0.0176	0.0142	0.0111
	Midsize car	0.0135	0.0104	0.0071	0.0048
	Large car	0.0154	0.0103	0.0066	0.0035
Dynamic Model no Transaction Costs	Small car	0.0012	0.0012	0.0012	0.0011
	Midsize car	0.0006	0.0006	0.0006	0.0005
	Large car	0.0003	0.0003	0.0003	0.0003

Table 8: Within segment cross-price elasticity of cars of different ages are reported. The elasticities are computed after a permanent change in the price of a new car belonging to the same segment.

Transaction Costs

PARAMETERS	
Constant	5.15** (0.12)
Log-Age	-0.24** (0.03)
Log- Engine size (CC)	0.28** (0.8)
Fiat	-0.24** (0.03)
Diesel*Time trend	-0.02** (0.004)
Log - Initial Stock	-0.14** (0.02)
Log - Price	-0.61** (0.03)
Time trend	-0.06** (0.01)
ξ_{jt}	0.09** (0.01)

Standard errors in parentheses; statistical significance at 5% level indicated with **, and at 10% with *. R-squared 0.45. Obs.1483

Table 9: Parameter Estimates – Dependent variable log of transaction costs.

Model	Year	Age	Taxes+Dealer Compensation	Transaction Costs: Estimates
Alfa 156 1.6i	1999	1	€ 1675	€ 2400
BMW 318i	1999	3	€ 1700	€ 1700
Fiat Punto 1.9 D	2003	2	€ 950	€ 1400
Audi A3 1.6 D	2003	5	€ 1450	€ 1940

Table 10: Transaction Costs - Examples

Scrappage Scheme

Starting Date	*January 1997*	*October, 1997*	*February, 1998*
Time in force	8 months	4 months	6 months
Total discount	€775 + €922 €1033+€1229	€775 + €922	€775+€922 €620+€738
Requirement	To scrap a car aged 10 years or older and buy a new one with an equal discount from the manufacturers. The first discount was awarded for a new car with cc<1300 and the second for cc >1300	To scrap a car aged 10 years or older and buy a new one with an equal discount from the manufacturers	To scrap a car aged 10 years or older and buy a new one with an equal discount from the manufacturers. The discounts were awarded respectively for a new with average consumption <7 l/km and average consumption <9 l/km

Table 11: Replacement Schemes in Italy

	Baseline model	Change after cash for replacement 10	Change after cash for replacement 8	Change after cash for scrappage
Government revenue	**20,641**	555	576	-1,776
Province Revenue	**4,786**	-28	-36	134

Value in 000s

Table 12: Government and Province revenues

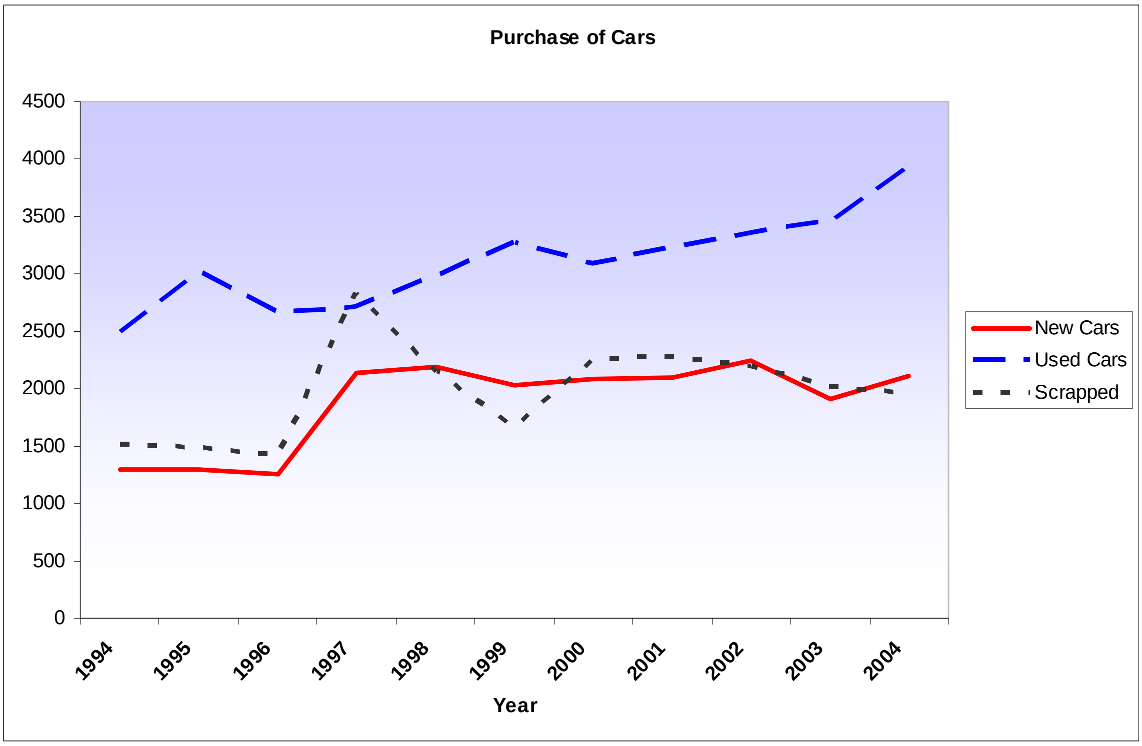

Figure 1 – Purchase of New and Used Cars

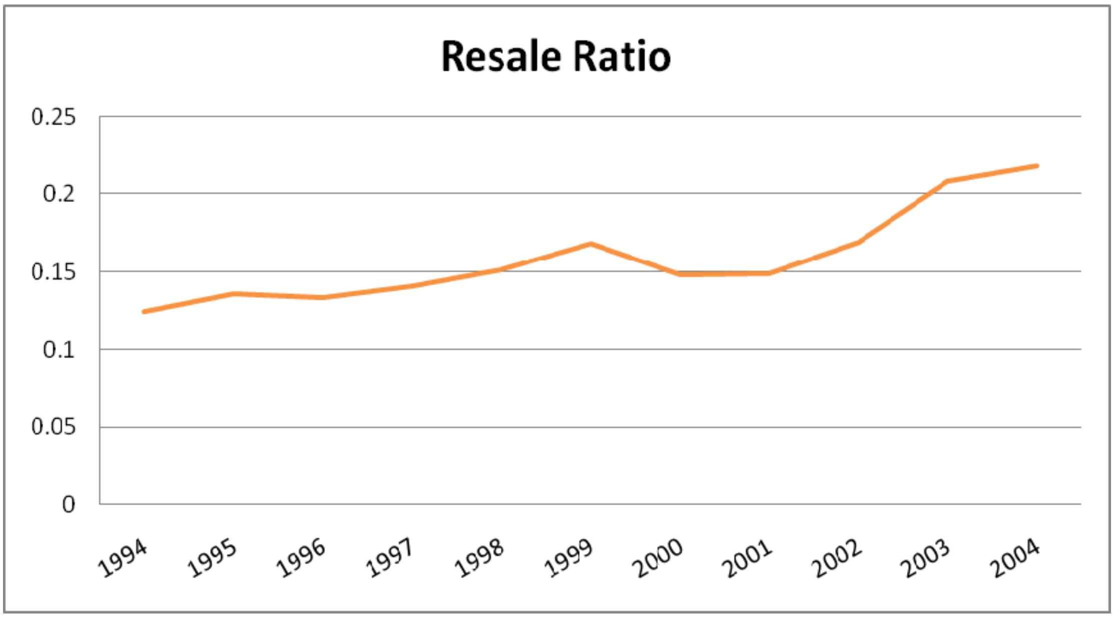

Figure 2 – Resale Ratio

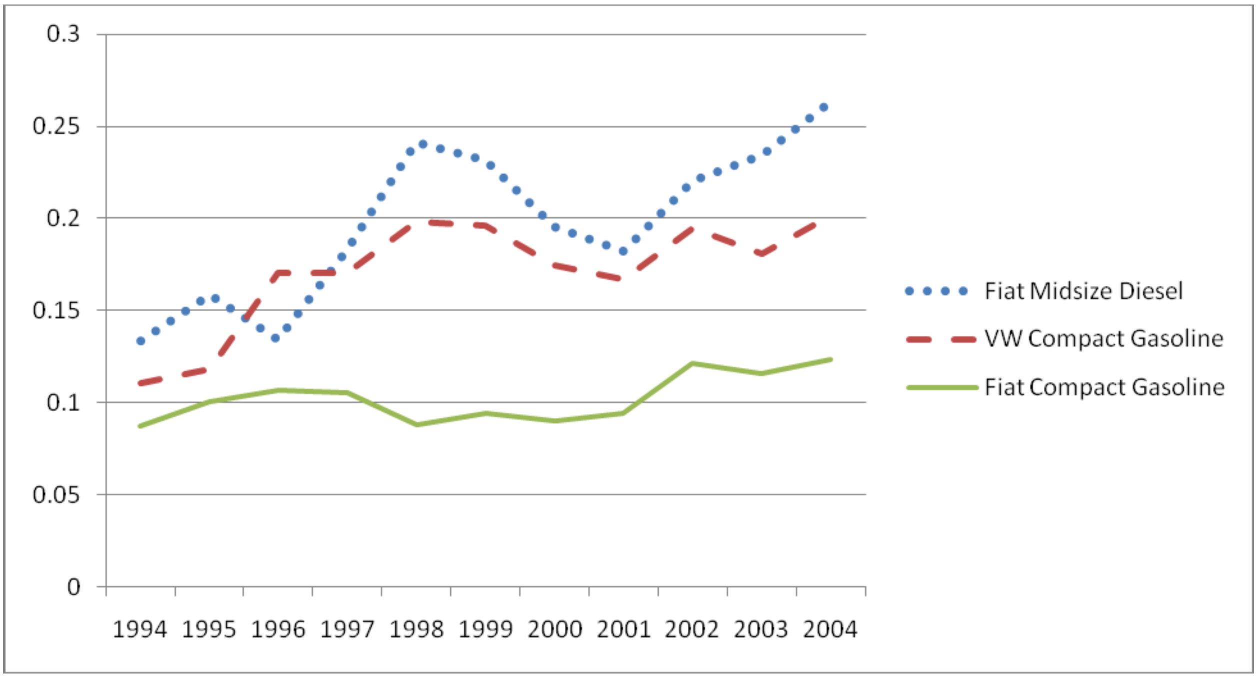

Figure 3 - Resale ratios for different car types

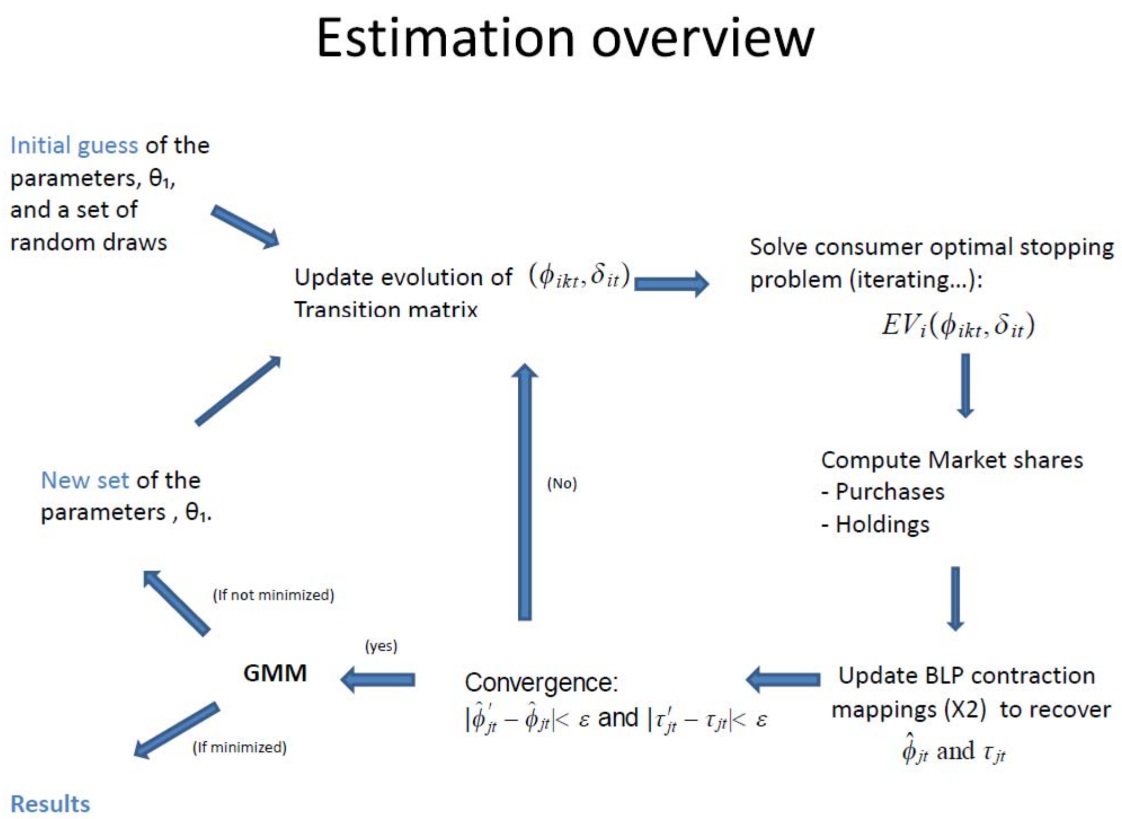

Figure 4 – Estimation Overview

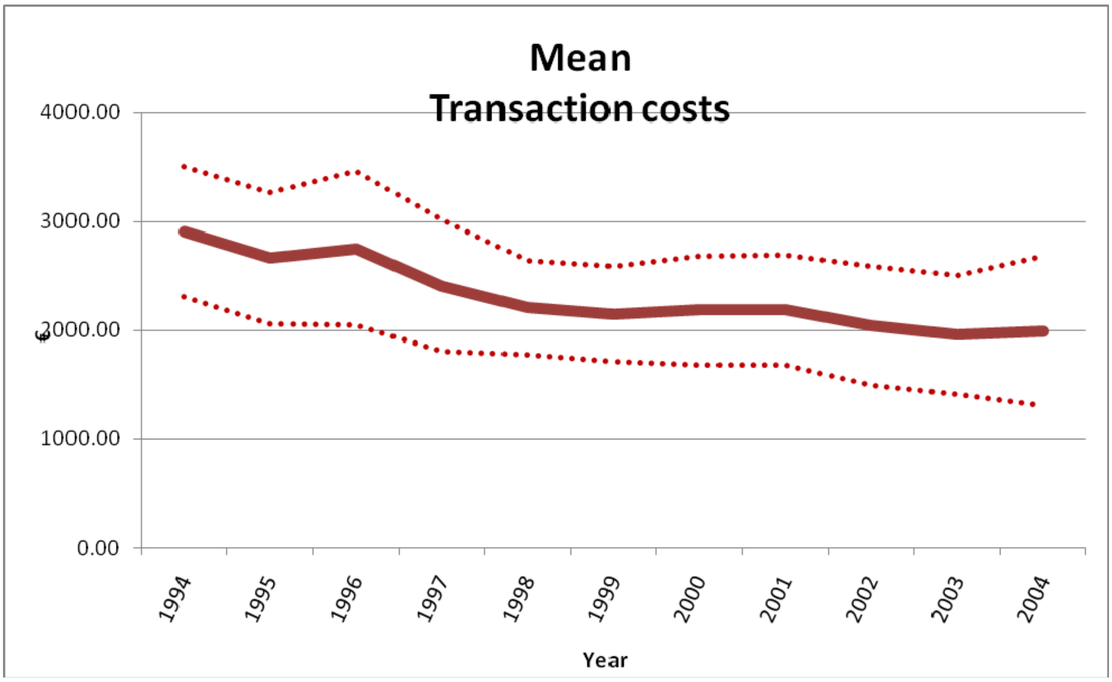

Figure 5 – Mean Transaction Costs over Time

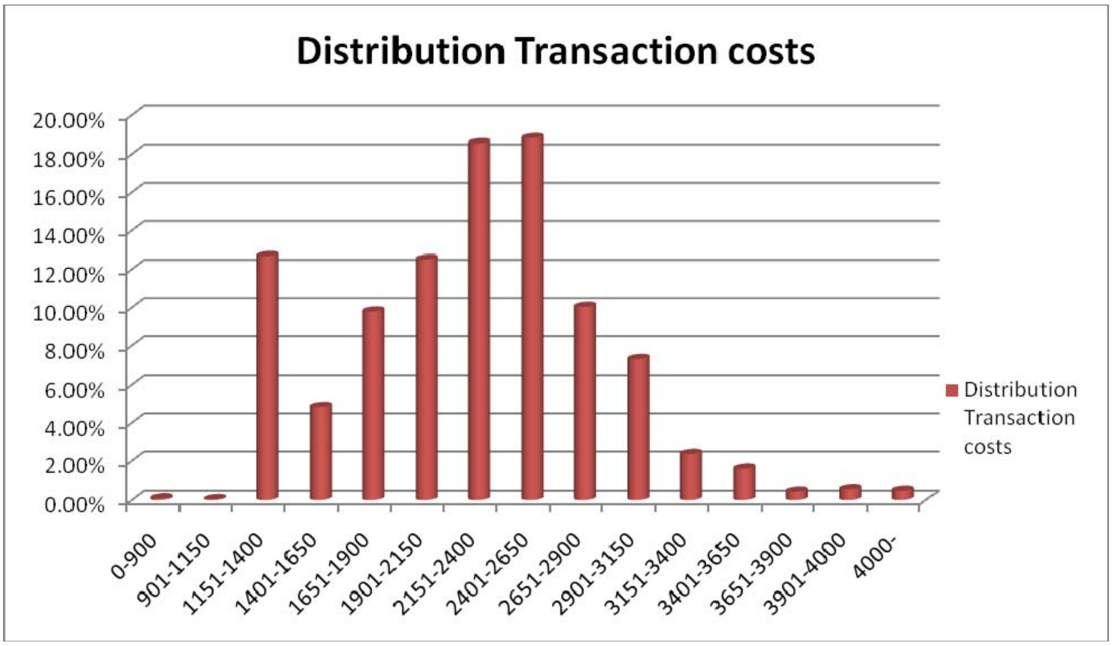

Figure 6 – Transaction Costs Distribution

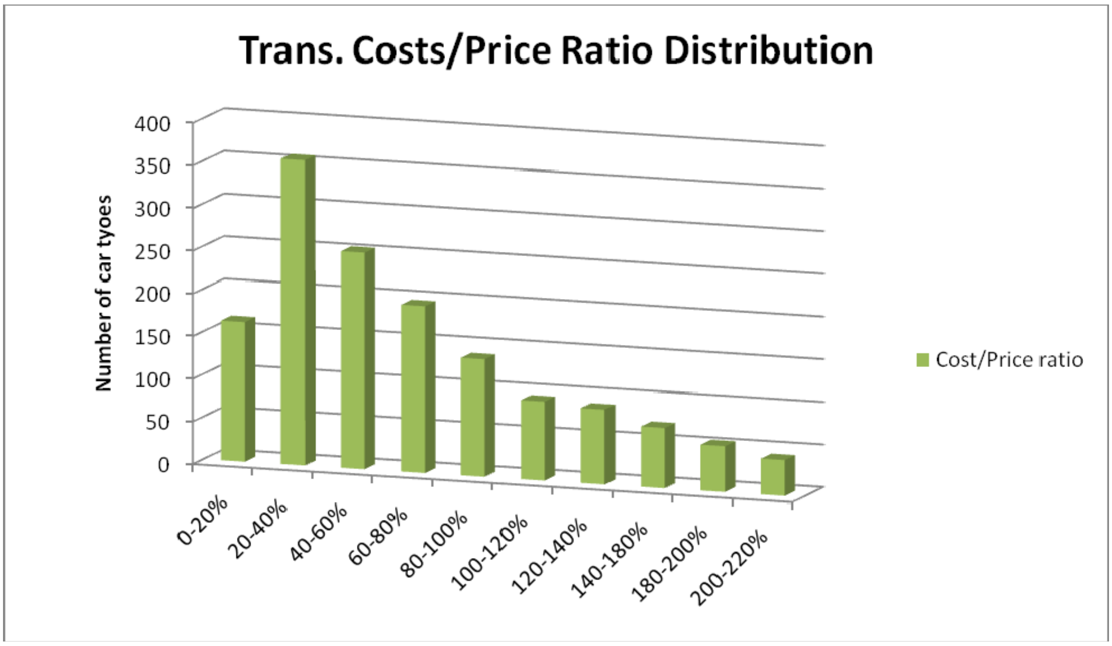

Figure 7 – Transaction Costs/Price Ratio

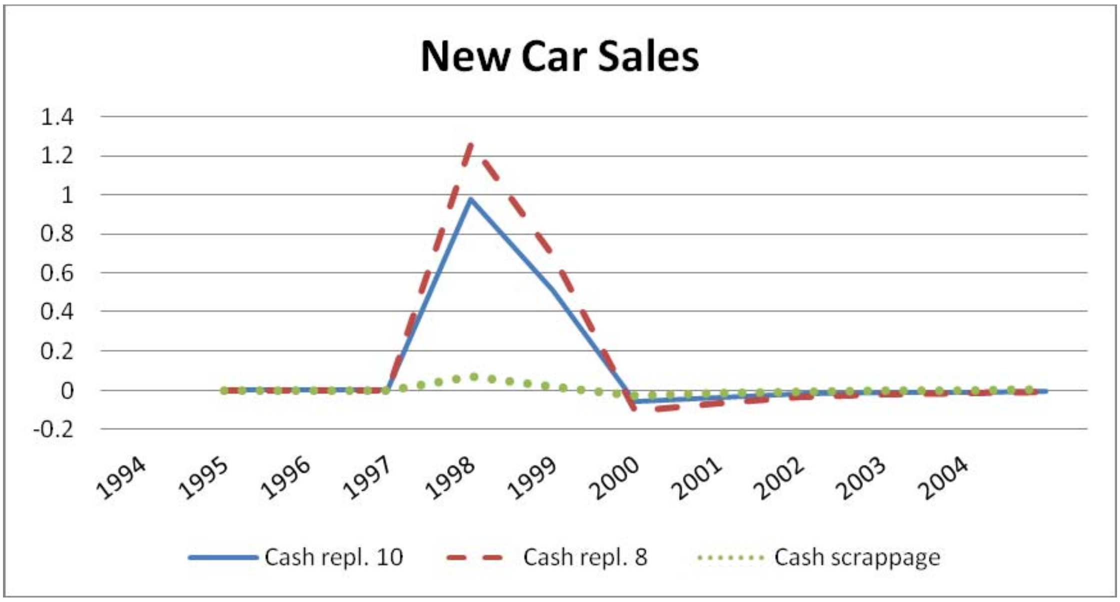

Figure 8 – Expected new-car sale relative to the baseline model with no subsidy

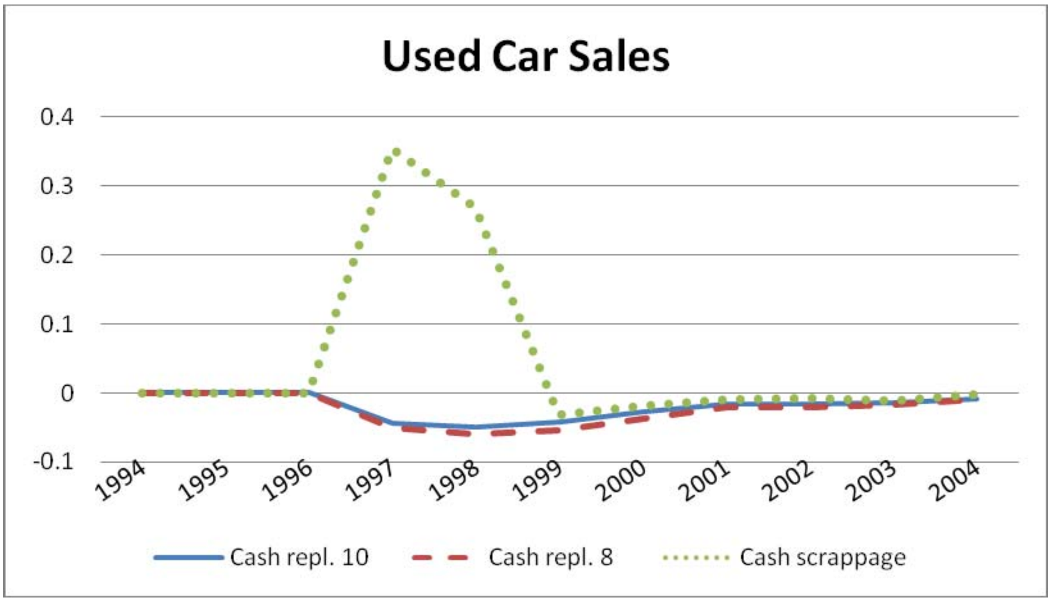

Figure 9 – Expected used-car sale relative to the baseline model with no subsidy

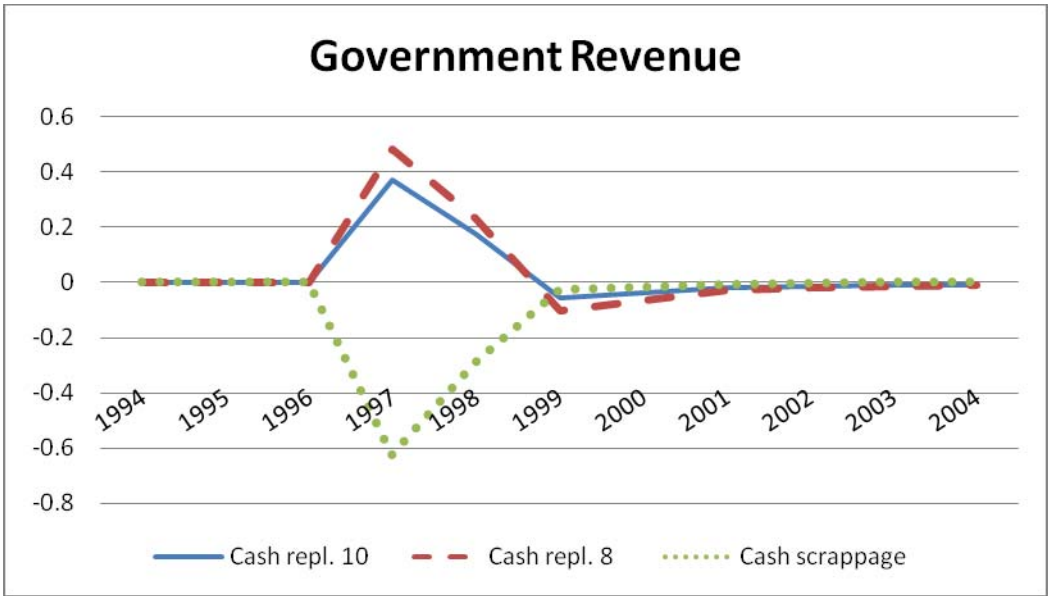

Figure 10 – Expected government revenues relative to the baseline model with no subsidy

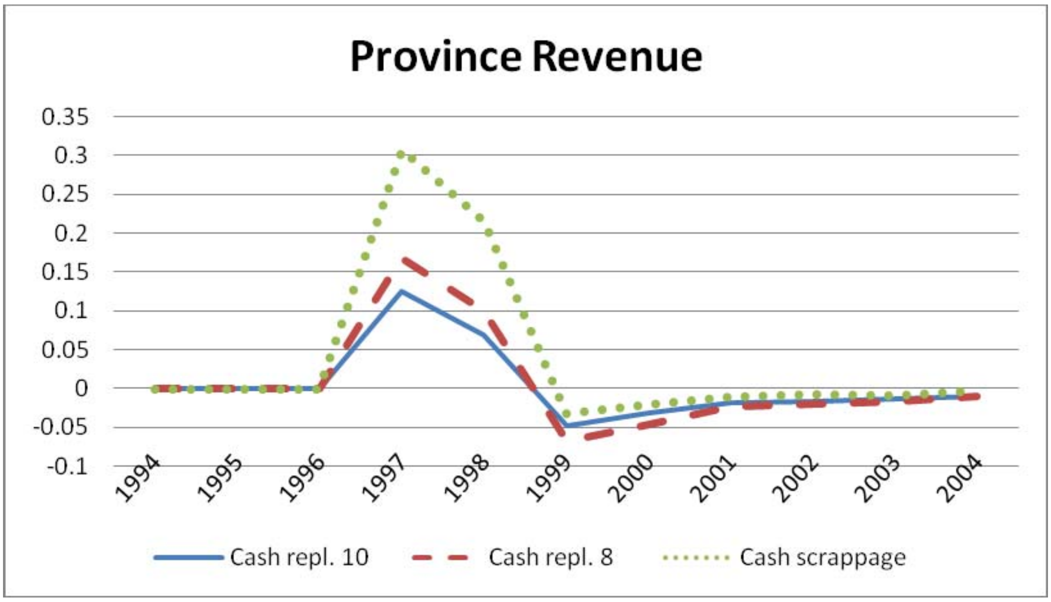

Figure 11 – Expected province revenues relative to the baseline model with no subsidy