RAND Journal of Economics
Vol. 42, No. 2, Summer 2011
pp. 266–291

Automobile replacement: a dynamic structural approach

Pasquale Schiraldi*

This article 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. 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-coefficient discrete-choice model that incorporates a dynamic optimal stopping problem. I apply this model to evaluate the impact of scrappage subsidies on the Italian automobile market.

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 United States; 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 good 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 cars (see Berry, Levinsohn, and Pakes, 1995, henceforth BLP; Bresnahan, 1981; Goldberg, 1995; Petrin, 2002). 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 characterizes markets for durable goods. The models either ignore the secondary market and its dynamics altogether, or lump used goods into a composite outside option. Despite their

*London School of Economics; p.schiraldi@lse.ac.uk.
This article has benefited greatly from discussions with Marc Rysman. Special thanks are also due to Victor Aguirregabiria, Michael Manove, Martin Pesendorfer, and Philipp Schimdt-Dengler and seminar participants at several institutions. I also thank Igal Hendel and two anonymous referees for their suggestions. All remaining errors are mine. Financial support from Bart Lipman, Michael Manove, the IED at Boston University, and the University of Bari is greatly appreciated.

importance, and although the auto market is one of the most studied in the literature, there have been relatively few empirical models of secondary markets for used goods.

A key 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 durable goods depreciate over time, arises because of the presence of transaction costs such as search costs, taxes, asymmetric information, switching costs, and so forth, which may vary over time. If there are no frictions, a consumer would choose a quality that maximizes her utility in each period and she would have no incentive to hold it across multiple periods once the quality of the good depreciates. Instead, the presence of these frictions makes replacement infrequent because consumers economize on transaction costs.

The model that I present explicitly accounts for dynamic consumer considerations and costly replacement decisions. More specifically, I assume consumers incur different kinds of costs upon the replacement of their automobile, which value depreciates over time. The quality of used goods is assumed to be common knowledge among the agents, and hence the model does not allow for the presence of adverse selection. With full information, the depreciation is captured by the decline in prices. The structural model explicitly takes this information into consideration and provides an estimation of the whole distribution of transaction costs, that is, the cost associated with each car type in each time period. Information about resales and prices, along with ownership data of used cars, provides a potential source of identification for 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 ownership for each individual car in the sample over a period of 11 years.[1] The data are from the Province of Isernia in Italy and are collected by the Motor Vehicle Department. The presence of these two market shares for each car type represents the main strength of my unique data set. In particular, the share of consumer holdings conveys information on the average time consumers keep a particular car model over time. Consequently, it conveys information on the relative level of transaction costs once the model accounts for the endogenous evolution of the consumer-type distribution across cars over time. Specifically, I estimate a discrete-choice logit model over a set of products allowing for preference heterogeneity on observable product characteristics that incorporates a dynamic optimal stopping problem in the spirit of Rust (1987) using market-level data. Like Berry (1994), I invert the market share of purchases and the market share of consumer holdings for each product in each period. When 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.

The contribution of this article to the durable goods literature is twofold. First, it is the first article 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. 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 to stimulate car sales. The model is used to investigate the impact of such policies on consumers' demand for new and used vehicles in the short and long run.

A number of recent articles (Melnikov, 2001; Carranza, 2006; Hendel and Nevo, 2006; Lee, 2009; Shcherbakov, 2008; Gowrisankaran and Rysman, 2009, henceforth GR) propose dynamic consumer choice models for aggregate data. As in many of these models, the major simplifying assumption here is that consumers perceive the evolution of product characteristics to be a simple

[1] 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.

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. GR is most similar to the present work. Differently from it, I allow for the presence of a second-hand market with transaction costs and good depreciation. There are recent studies that deal with the implications of durability on the dynamics of car demand (Esteban and Shum, 2007; Stolyarov, 2002; Adda and Cooper, 2000). My model is the first that implements a structural estimation of a dynamic framework allowing for persistent consumer heterogeneity and repeat replacement decision of new and used cars in the presence of transaction costs.[2]

2. The automobile market: data

■ The Italian automobile market is the fourth largest market in the world (after the United States, Japan, and Germany), with about 2 million cars sold every year. Most cars sold are manufactured by the FIAT Group, which controls the following brands: FIAT, Lancia, Alfa Romeo, Innocenti, Autobianchi, Ferrari, and Maserati. The FIAT Group's share was more than 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.[3] The average population in this province is 74,114 and is constant during the sample period, and the mean income is €21,547 in 1994 euros.

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 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, the vehicle's age $(0,\ldots,10)$, where 0 stands for a new car and 10 groups together all cars 10 years or older,[4] engine displacement (*small* if $cc <= 1300$, *medium* if $1300 < cc <= 1800$, *large* if $cc > 1800$), type of fuel: gasoline or diesel, and origin of manufacturers.[5] In particular, I consider three possible macrogroups of manufacturers: (i) the Italian FIAT Group, which controls the following brands: FIAT, Lancia, Alfa Romeo, Innocenti, Autobianchi, Ferrari, and Maserati; (ii) manufacturers located in Germany: BMW, Mercedes,

[2]Hendel and Lizzeri (1999a, 1999b) and Schiraldi and Nava (2010), among others, study vertical differentiated models in which durable goods live for just two periods, so that used goods of all ages are lumped together. See also Rapson and Schiraldi (2011) for another application of the present setting.

[3]www.bancaditalia.it and www.istat.it, respectively.

[4]I assume that a 10-year-old car no longer depreciates and provides the same utility to the consumer. I also assume that the price is the same across cars older than 10 years.

[5]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 were designed according to this characteristic (as explained later) and second, until 1999, property taxes were paid based on the size of the engine displacement.

TABLE 1 Descriptive Statistics

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

Note: Market shares are used to weight prices.

Opel, Volkswagen, Audi and Porsche; and (iii) 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).

In the empirical analysis, I focus on the market for passenger cars, excluding trucks, vans, minivans, SUVs, and luxury cars (such as Ferrari and Lamborghini). 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 1 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 number of new units purchased suddenly jumped in 1997 and 1998, when the government introduced the scrappage policy.[6] The purchases of used cars decreased in 1997 and 1998, followed by a steep increase in the following years.

□ **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 friction, consumers have no incentive to hold their durable goods across multiple quality levels. Each heterogeneous consumer will choose a durable good from the product spectrum so as to maximize her net surplus. Hence, we should observe a high turnover rate. Specifically, if the good depreciates every period, consumers will never hold their durable good more than one period but will always update to their preferred car quality (see, for example, Hendel and Lizzeri, 1999a; Hideo and Sandfort, 2002; Rust, 1985, among others). Whereas 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 associated transaction costs play a crucial role in explaining consumers' decisions to replace their car.

[6]See Section 5 for a more extensive discussion of the scrappage policy.

FIGURE 1

PURCHASE OF NEW AND USED CARS

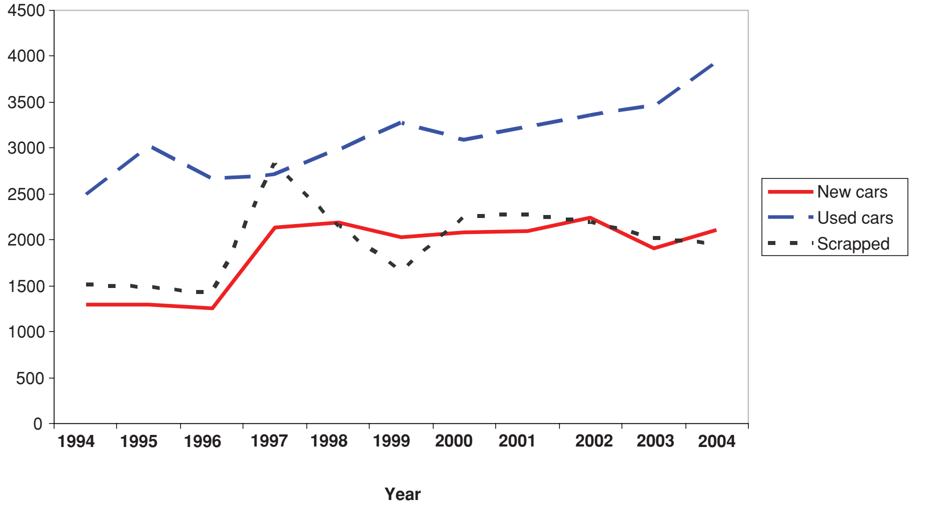

Each period, consumers assess the quality of the durable good they own. If the gain in utility from updating their holdings, net of prices, exceeds transaction costs, consumers sell their used goods in the second-hand market and replace them with durable goods of the preferred quality. It follows that a high level of transaction costs reduces the frequency of replacement. The two driving assumptions are: (i) cars depreciate after every year; and (ii) 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 and the durable good trading volume) across all cars in each year gives an idea about the presence and size of transaction costs. The average resale ratio per year in the data varies between 0.15 and 0.25.[7] This is substantially lower than 100% trade, as would be predicted by vertical differentiation models with consumers' heterogeneity but without frictions. Figure 2 reports the same ratios for a few models and 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.

Before turning to a more formal analysis, I present some simple regressions that show the importance of transaction costs in the second-hand market. In Table 2, I report the results of a regression in which the dependent variable is the resale ratios for different models/ages of cars, and the main explanatory variable of interest is the price depreciation. To control for unobservable quality, I also add model dummies and time dummies. Notice that the volume of trade is positively correlated with price depreciation, as we should expect in a model with transaction costs. The presence of transaction costs in the market could cause a potential endogeneity problem in the previous regression, as 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 it is to trade each given car on the secondary market. I collected these data from 2002 to 2004, and then ran the same regression as above, including this additional variable. More specifically, I included these ratings (with a negative sign) which should capture the presence of transaction costs related to specific car models. I report the results in Table 3. As

[7] The unit of observation is model/age for different car segments: subcompact cars, compact cars, medium size cars, medium/full-size cars, full-size cars.

FIGURE 2

RESALE RATIOS FOR DIFFERENT CAR TYPES

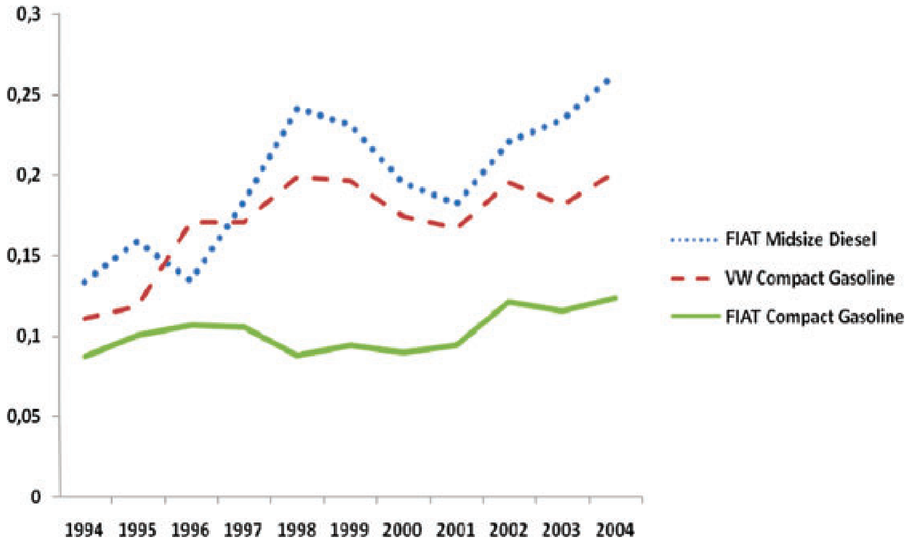

TABLE 2 **Estimation: Volume of Trade (1994–2004)**

	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	

Note: Standard errors are in parentheses; statistical significance at 5% level is indicated with **, and at 10% with *. R^2, 0.41. Observations 1648.

TABLE 3 **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	

Note: Standard errors are in parentheses; statistical significance at 5% level is indicated with **, and at 10% with *. R^2, 0.27. Observations 437.

expected, the volume of trade is negatively correlated with such a variable, whereas the sign of price depreciation does not change.

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. Because products degrade over time, a given consumer will desire to replace her durable good over time, either with a brand new durable good or with a second-hand one. In the model, consumers have perfect information about durable goods, 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 $J_t = \{j : j \in \{\cup_{\tau=1}^{t} \mathbf{j}_\tau\}\}$ denote 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, that is, $j = 0$, 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 good 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 markets in period t (including the outside option). In either case, she faces a similar (although not identical) decision problem in time $t+1$. Because 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 she purchases a car and that 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 heterogeneous in their price sensitivity and how intensively they prefer car characteristics. Consumers also have an idiosyncratic shock to their preferences for each good and in each period. Let $\epsilon_{it} = (\epsilon_{i0t}, \epsilon_{i1t}, \ldots, \epsilon_{iJ_t t})$ be the vector of idiosyncratic shocks of consumer i for period t, which are *i.i.d.* across (i, j, t) and capture horizontal differentiation.

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 represents 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, and hence both these elements capture the depreciation of durable goods over time. The depreciation is not deterministic because the unobserved product characteristic evolves stochastically over time.

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

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

where $\mu_t = \{d, j\}$, 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

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

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.

□ **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 dimensionality of the state space, and to make it computationally tractable. It is opportune to subtract the price of the car owned from equation (4) and redefine $V_i(k, \epsilon_{it}, \Omega_t) \equiv \hat{V}_i(k, \epsilon_{it}, \Omega_t) - \alpha_i^p p_{kt}$. Notice that $p_{0t} = 0$ implies $V_i(0, \epsilon_{it+1}, \Omega_{it+1}) = \hat{V}_i(0, \epsilon_{it+1}, \Omega_{it+1})$. Substituting for (1), (2), and

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

(3), we can rewrite equation (4) as follows[9]:

$$V_i(k, \epsilon_{it}, \Omega_{it}) = \max \Big\{\epsilon_{i0t} + \beta\mathbb{E}[V_i(0, \epsilon_{it+1}, \Omega_{it+1})|\Omega_{it}, \epsilon_{it}],$$
$$\max_{j\in J_t} x_{jt}\alpha_i^x + \xi_{jt} - \alpha_i^p p_{jt} - \tau_{jt} + \epsilon_{ijt} + \beta\mathbb{E}\left[\alpha_i^p p_{jt+1} + V_i(j, \epsilon_{it+1}, \Omega_{it+1})|\Omega_{it}, \epsilon_{it}\right],$$
$$1[k \neq 0]\cdot \left(x_{kt}\alpha_i^x + \xi_{kt} - \alpha_i^p p_{kt} + \epsilon_{ikt} + \beta\mathbb{E}\left[\alpha_i^p p_{kt+1} + V_i(k, \epsilon_{it+1}, \Omega_{it+1})|\Omega_{it}, \epsilon_{it}\right]\right)\Big\}. \tag{5}$$

This transformation is not strictly necessary in order to estimate the model, but it is convenient to clearly define the sufficient statistics used to reduce the dimensionality of the state space.

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, as do technological progress and changes in prices for existing products, according to optimal price decisions. Consumers are uncertain about future product attributes but rationally 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 $EV_i(j, \Omega_{it}) = \int_{\epsilon_{it}} V_i(j, \epsilon_{it}, \Omega_{it})dP_\epsilon \forall j \in J_{t-1} \cup \{0\}$ denote the expectation of the value function, integrated over the realization of ϵ_{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

$$\phi_{ijt} \equiv x_{jt}\alpha_i^x + \xi_{jt} - \alpha_i^p(p_{jt} - \beta\mathbb{E}_t[p_{jt+1}]), \tag{6}$$

where $(p_{jt} - \beta\mathbb{E}_t[p_{jt+1}])$ 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 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 later in Section 3.[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 take into account to make their decisions.

Finally, I use the aggregation properties of the type I extreme-value distribution of ϵ_{ijt} to express the expectation of the Bellman equation in a relatively simple form. In particular, the

[9] Notice that Ω_{it} already includes the vector of prices, and hence there is no need to explicitly include the resale price as a state variable once the change of variable is performed.

[10] The logit inclusive value (or social surplus) was first introduced by McFadden (1974) and then first used by Melnikov (2001) as a sufficient statistic that fully describes the continuation value of consumer utility.

[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 consumer-specific terms. In particular, given $\hat{\phi}_{jt}$, the simulated draws v_i and the vector of nonlinear parameters $\{\alpha^p, \sigma\}$: $\phi_{ijt} = \hat{\phi}_{jt} + \sum_j x_{jt}\sigma v_i - \alpha_i^p(p_{jt} - \beta\mathbb{E}_t[p_{jt+1}])$.

expected value of the best choice from several options in a logit model can be expressed as the logarithm of the sum of the mean expected discounted utility of each option and for each consumer i:

$$\delta_{it} = \ln\left(\sum_{j\in J_t} \exp(\phi_{ijt} - \tau_{jt} + \beta\mathbb{E}[EV_i(j, \Omega_{it+1})|\Omega_{it}])\right). \tag{7}$$

The logit inclusive value is the maximum expected utility from buying one of the J_t products present in the primary and secondary markets in period t.[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 $(\phi_{ikt}, \delta_{it})$ which can be summarized by a first-order Markov process:

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

where G_i is consumer specific.

Similarly to the inclusive value sufficiency assumption (Hendel and Nevo, 2006; GR), this assumption implies that the vectors $\{\phi_{ikt}, \delta_{it}\}$ contain all the relevant information in Ω_{it} to obtain the probability distribution of $\{\phi_{ikt+1}, \delta_{it+1}\}$ conditional on Ω_{it}.[14] 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 transition probabilities to the next-period state in the current framework. For the estimation of the model, I assume that the Markov processes take the following linear functional form:

$$\delta_{it+1} = \rho_{1i} + \rho_{2i}\delta_{it} + \eta_{it} \tag{8}$$

$$\phi_{ijt+1} = \gamma_{1i} + \gamma_{2i}\phi_{ijt} + \gamma_{3i}\delta_{it} + \mu_{it}, \tag{9}$$

where η_{it} and μ_{it} are independent and normally distributed, and ρ_{1i}, ρ_{2i}, γ_{1i}, γ_{2i}, and γ_{3i} are incidental parameters specific to each consumer i. The inclusive value in (9) captures the common component that drives the evolution of ϕ_{ijt}. As in GR, the specification of δ_{it} in equation (7) includes not only prices and characteristics of the products available but also future optimal decision making. A different car type purchased or owned in the current period is going to affect a future consumer's decision differently. Consumers purchasing product j will have different values of not replacing it because the alternative is to stay with the product currently owned. This definition of the inclusive value is different from the one used in Melnikov (2001) or Hendel and Nevo (2006), where only exogenous characteristics (the expected flow utility) define the inclusive value. In Melnikov, the static definition results from the assumption that there is no repeat purchase. In the no-repeat-purchase case, the dynamics involve a decision on when to buy. The decision of which product to buy, conditional on purchase, is static. In Hendel and Nevo,

---

[13]Notice that the transformation made at the beginning of this subsection allows me to uniquely define the net augmented utility flow and the logit inclusive value independently of the type of car owned.

[14]The introduction of the scrappage policy potentially requires the introduction of a third state variable in the model. I do not introduce it for two reasons: first, the policy was independently introduced in 1997 and 1998, with the requirement that consumers owned the scrapped car for at least one year; second, data availability and computational costs. An introduction of a third state variable would have required a richer specification than (8) with more parameters to estimate using only 11 points given the 11 years observed.

the static definition of the inclusive value derives from the assumption that products are perfect substitutes in consumption and storage. Hence, choices of products of the same size impact the dynamics in the same way. Moreover, 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 articles, includes only the price and characteristics of the product owned by consumers.

Using the previous assumptions, I can write $EV_i(\Omega_{it})$ as $EV_i(\phi_{ikt}, \delta_{it})$ and rewrite the Bellman equations (5) for consumer i as

$$\begin{aligned} EV_i(\phi_{ikt}, \delta_{it}) = \ln(&\exp(\delta_{it}) + 1[k \neq 0] \cdot \exp(\phi_{ikt} + \beta \mathbb{E}[EV_i(\phi_{ikt+1}, \delta_{it+1} | \phi_{ikt}, \delta_{it})] \\ &+ \exp(\beta \mathbb{E}[EV_i(0, \delta_{it+1} | \phi_{ikt}, \delta_{it}))). \end{aligned} \tag{10}$$

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

$$d_{it}^{kj} = \frac{\exp(\phi_{ijt} - \tau_{jt} + \beta \mathbb{E}[EV_i(\phi_{ijt+1}, \delta_{it+1}) | \delta_{it}, \phi_{ijt}])}{\exp(\delta_{it}) + \exp(\beta \mathbb{E}[EV_i(0, \delta_{it+1}) | \delta_{it}]) + 1[k \neq 0] \cdot \exp(\phi_{ikt} + \beta \mathbb{E}[EV_i(\phi_{ikt+1}, \delta_{it+1}) | \delta_{it}, \phi_{ikt}])}. \tag{11}$$

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:

$$\widetilde{d}_{it}^{k} = \frac{\exp(\phi_{ikt} + \beta \mathbb{E}[EV_i(\phi_{ikt+1}, \delta_{it+1}) | \delta_{it}, \phi_{ikt}])}{\exp(\delta_{it}) + \exp(\beta \mathbb{E}[EV_i(0, \delta_{it+1}) | \delta_{it}]) + 1[k \neq 0] \cdot \exp(\phi_{ikt} + \beta \mathbb{E}[EV_i(\phi_{ikt+1}, \delta_{it+1}) | \delta_{it}, \phi_{ikt}])}, \tag{12}$$

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 good 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 d_{it}^{kj} over all existing products:

$$s_{jt}^{D} = \int_{\upsilon_i} \sum_{k \in J_{t-1} \cup \{0\}} d_{it}^{kj} s_{ikt-1} dP_{\upsilon} \tag{13}$$

$$\widetilde{s}_{kt}^{H} = \int_{\upsilon_i} \widetilde{d}_{it}^{k} s_{ikt-1} dP_{\upsilon}, \tag{14}$$

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

$$s_{ijt} = \widetilde{s}_{ijt}^{H} + s_{ijt}^{D}, \tag{15}$$

where $\widetilde{s}_{ijt}^{H}$ and s_{ijt}^{D} are obtained as in equations (13) and (14) without integrating over the consumer heterogeneity. The market size M_t is observed and evolves deterministically over time.

□ **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. 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 of 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 modelled 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. Specifically, I assume that consumer preference for the age of the car α_i^{age} is independently distributed normally with mean $\alpha^{age}$and standard deviation $\sigma_{\alpha^{age}}$, that is, $\alpha_i^{age} = \alpha^{age} + \upsilon_i \sigma_{\alpha^{age}}$, where $\upsilon_i \sim N(0, 1)$.

Like Berry (1994), for any parameter vector θ, I recover the set of unobservable product characteristics and transaction costs (ξ_{jt}, τ_{jt}) that perfectly rationalize the model's predicted market shares,[15] 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 that these innovations in product unobservables, ξ_{jt} (and not the product unobservables themselves), are uncorrelated with a vector of instruments.[16] More specifically, I assume the following.

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

$$\mathbb{E}[Z_{jt}\varsigma_{jt}] = 0,$$

where Z_{jt} are instruments.

The drift of this process is set to 0, because 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 below. 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

$$\hat{\theta} = \arg\min_{\theta} G(\theta)' W^{-1} G(\theta), \tag{16}$$

where $G(\theta)$ is a vector of stacked moments and W is the weighting matrix.[17] The computation of the objective function requires knowledge of the weight matrix, W, that, 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. Given that θ_2 can be expressed as a function of θ_1, the nonlinear search is performed only over θ_1 using a non-derivative-based Nelder and Mead (1965) simplex algorithm.

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, I combine BLP's procedure to recover the unobserved product characteristics ξ_{jt} and transaction costs τ_{jt}

[15]See the discussion in the following subsection.

[16]This assumption is also discussed in BLP 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.

[17]The diagonal elements of the weighting matrix should be the inverse of the variance of the moment. For the 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 eligible 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 Berry, Levinsohn, and Pakes (2004) and Petrin (2002) for more details on calculating weighting matrices when combining micromoments with aggregate moments.

FIGURE 3

ESTIMATION OVERVIEW

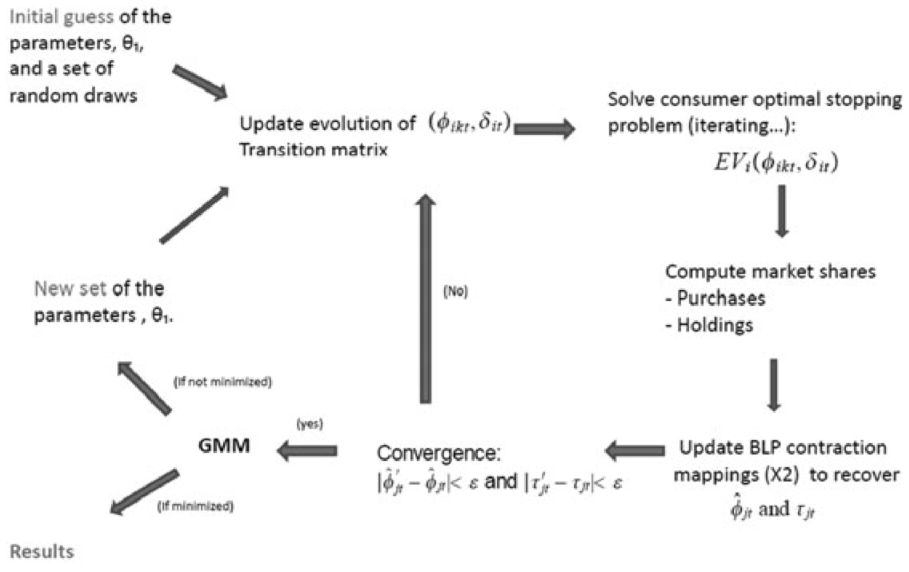

with Rust's (1987) fixed-point algorithm to solve consumers' dynamic optimization problems. Once ξ_{jt} is recovered, the objective function in (16) can be computed. Hence, the nonlinear search is performed to recover the parameters of the model. Figure 3 shows an overview of the computation algorithm 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 contraction mappings. 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 $\hat{\phi}_{jt}, \tau_{jt}$) and given the nonlinear parameters θ_1, the set of simulated draws is used to calculate the logit inclusive values as in equation (7) and net augmented utility[18] as in equation (6). Both these variables are used to estimate the coefficients of the Markov process regressions in (8) and (9). These coefficients are then used to construct the transition matrix and to calculate the expected value function (10) by iteration (Rust, 1987). Hence, individual probabilities are computed as in (11) and (12). The number and identity of consumers for each product available on the market evolve according to the individual probability of buying or keeping product j as predicted by the model. The individual probabilities are aggregated as in (13), (14), and (15) 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 point ξ_{jt} is recovered from the final value of $\hat{\phi}_{jt}$ via linear regression.[19] Then, the objective function in (16) is computed and the nonlinear search of θ_1 is performed.

I use two BLP-type contraction mappings to recover $\hat{\phi}_{jt}$ and τ_{jt}.[20] I observe $2K$ market shares for K used products: the market share of consumers' purchases and the market share of

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

[19]An alternative approach to utilizing multiple nested-fixed point routines includes mathematical program with equilibrium constraints (MPEC), as described in Dube, Fox, and Su (2008).

[20]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, (8) and (9).

consumers' holdings. Hence, I invert these shares and 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:

$$\hat{\phi}'_{jt} = \hat{\phi}_{jt} + \psi_1\big(\ln\big(\check{s}^H_{jt}\big) - \ln\big(\widetilde{s}^H_{jt}(\hat{\phi}_{jt}, \tau_{jt}, \theta)\big)\big). \tag{17}$$

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,

$$(\hat{\phi}_{jt} - \tau_{jt})' = (\hat{\phi}_{jt} - \tau_{jt}) + \psi_2\big(\ln\big(\check{s}^D_{jt}\big) - \ln\big(s^D_{jt}(\hat{\phi}_{jt}, \tau_{jt}, \theta)\big)\big). \tag{18}$$

Having computed $\hat{\phi}_{jt}$ from equation (17), I can recover τ_{jt} from the previous equation by difference. $s^D_{jt}(\hat{\phi}_{jt}, \theta)$ and $\widetilde{s}^H_{jt}(\hat{\phi}_{jt}, \theta)$ are computed from equations (13) and (14) and ψ_1 and ψ_2 are *tuning* parameters whereas $\check{s}^H_{jt}$ and $\check{s}^D_{jt}$ are the corresponding shares observed in the data. I set $\psi_2 = (1 - \beta)^2$ and $\psi_1 = 1 - \beta$ to increase the speed of convergence. 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.

To perform the iterative calculation, I discretize the state space (ϕ_{ikt}, δ_{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.

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 have tried different options,[21] and the results are quite similar.

Because 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.* pseudorandom 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.

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, and 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, α and σ, are identified analogously to BLP. Among people who choose to buy any car, I look at the share that chooses 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. 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, which 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 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 reoptimize and

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

choose 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 costs she will pay. With full information, prices reflect deterioration in quality. Conditional on replacing the car, $\xi_{jt} - \tau_{jt}$ is identified from demand among consumers who 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.[22] 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. Because all of the elements of the model are solved 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.

Instruments. Valid instruments must be correlated with regressors but uncorrelated with the time t unobservable innovation. Although observed product characteristics may be endogenous with respect to unobserved characteristics ξ_{jt}, 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 pseudodifferences 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, pseudodifferences in x_{jt} are used as instruments to identify α^x. These pseudodifferences are valid instruments if consumers cannot predict the future value of ς when making their decision at time t. Because 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 the price of used cars are, the 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.

4. Results and implications

□ **Parameter estimates: utility specification.** Table 4 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 micromoments; column 3 reports the estimates of a dynamic model without the transaction costs; column 4 reports 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 age of the car. The price coefficient is estimated nonlinearly and the magnitude

[22] 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.

[23] 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 alone will not be sufficient to explain consumers' holding behavior.

TABLE 4 Estimation Results: Utility

Parameters	Dynamic Model with Micromoments	Dynamic Model without Micromoments	Dynamic Model with No Transaction Costs	Static Model
Constant	4.21** (0.42)	3.30** (0.4)	−6.45** (0.57)	−4.25** (0.30)
Log(Age)	−2.21** (0.10)	−2.35** (0.09)	−0.76** (0.13)	−2.20** (0.09)
Engine size (cc)	2.73** (0.28)	3.08** (0.26)	−0.19 (0.24)	1.18** (0.15)
FIAT	0.06 (0.12)	0.14 (0.12)	0.48** (0.23)	0.62** (0.16)
German	0.51** (0.15)	0.55** (0.15)	−0.01 (0.23)	0.55** (0.15)
Diesel	−2.74** (026)	−2.33** (0.25)	0.41 (0.56)	0.11 (0.38)
Diesel ∗ Time trend	0.24** (0.03)	0.18** (0.03)	0.01 (0.05)	−0.01 (0.04)
Nonlinear Parameters				
(Price-Expected price), α_p	−52.37** (4.64)	−56.71** (9.14)	−2.04 (4.52)	–
Price, α_p				−9.32** (6.13)
Log(Age), σ_{age}	0.52** (0.11)	0.04 (0.19)	0.24 (0.88)	0.03 (4.64)
λ	0.35** (0.03)	0.30** (0.04)	0.80** (0.23)	0.75** (0.02)

Note: Standard errors are in parentheses; statistical significance at 5% level is indicated with **, and at 10% with *.

is −52.37. A consumer obtains a positive flow utility from owning a car (relative to the outside option) with a mean constant term of 4.21. 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 2.73 shows that consumers prefer cars with a higher cc engine. Dummies for location suggest consumers' preference for German cars. 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 captures 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.[24] 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 and λ is estimated to be 0.35, which shows a significant persistence in the unobservable product over time. The micromoments improve the identification of the consumer heterogeneity and the price coefficient: σ^{age} becomes significant in the specification with micromoments.

Column 3 provides estimates from the dynamic model where no transaction costs are paid to replace the automobile and hence there are no frictions in the market. The dynamic model without transaction costs has a simple analytical solution. If there are no transaction costs, the problem is no longer state dependent and $\mathbb{E}_t[EV_i(k, .)] = \mathbb{E}_t[EV_i(0, .) + \alpha_i^p p_{kt+1}]$.[25] The probability of purchasing any good j does not depend on the car owned. It is similar to the static model but for the presence of the expected price in the flow utility function as in the dynamic model with transaction costs. The very imprecise price coefficient and the unexpected sign on the characteristics that enter the mean utility suggest that the data cannot easily be explained 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 is −9.32. The estimates of 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 of the car becomes insignificant.

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

[25] Notice that $EV_i(0, .)$ no longer depends on the good held by consumers and is constant across products.

TABLE 5 Own-Price Elasticities for New Cars (1994)

Market Segment	Dynamic Model with Transaction Costs and Micromoments, Short Run	Dynamic Model with Transaction Costs and Micromoments, Long Run	Static Model	Dynamic Model, No Transaction Costs
Small car	**−8.81**	**−3.54**	**−2.39**	**−0.34**
Domestic	−8.82	−3.70	−2.45	−0.4
Foreign	−8.78	−3.26	−2.38	−0.33
Midsize car	**−9.31**	**−4.33**	**−2.87**	**−0.52**
Domestic	−8.83	−4.41	−2.73	−0.47
Foreign	−9.63	−4.28	−2.92	−0.54
Large car	**−10.23**	**−4.34**	**−3.9**	**−0.74**
Domestic	−10.20	−4.07	−4.01	−0.76
Foreign	−10.25	−4.69	−3.63	−0.68

Note: Market shares are used to weight the price elasticities.

TABLE 6 Within-Market Segment Cross-Price Elasticities (1994)

	Market Segment	New	2-Year-Old	5-Year-Old	7-Year-Old
Full dynamic model (Long run)	Small car	0.1558	0.0932	0.0526	0.0367
	Midsize car	0.1132	0.0675	0.0321	0.0179
	Large car	0.0797	0.0548	0.0202	0.0120
Full dynamic model (Short run)	Small car	0.2564	0.1026	0.0394	0.0206
	Midsize car	0.1466	0.0700	0.0150	0.0096
	Large car	0.0691	0.0241	0.0101	0.0050
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

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

Price elasticity. Ignoring the dynamics in the data and the presence of transaction costs can bias the estimate of the price elasticities. A static model and a dynamic model without transaction costs (hence without state dependence) do not account for the possibility of consumers timing their purchase and therefore do not allow for the distribution of consumer types to endogenously evolve over time. In Tables 5 and 6, I present the average[26] own- and cross-price short- and 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. The short-run price elasticities are computed for a temporary price increase in 1994. I calculate the long-run elasticity using a permanent change in the price of a product in 1994 and by looking at market-share changes in the following 3 years.[27] 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,[28] keeping the percentage of the depreciation in price of the same model across different ages unchanged.[29] Permanent changes capture the long-term effects of the change on the consumer's

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

[27]I have followed a similar approach as in GR, where the long-run elasticities are computed by looking at the market share change over 12 months.

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

[29]In computing the elasticities for the dynamic model without transaction costs, I also consider a change in price of the new model along with the change of the future expected price of the same model in the used market.

TABLE 7 Parameter Estimates Dependent Variable Transaction Costs

Parameters	
Constant	7.75** (0.25)
Log-Age	−0.06** (0.01)
Log-Engine size (cc)	−0.09** (0.04)
FIAT	−0.10** (0.01)
Diesel ∗ Time trend	−0.01** (0.002)
Initial share of car *j*	−7.46** (0.79)
Scrapped cars	−0.28 (0.21)
Log-Price	−0.20** (0.02)
Time trend	−0.01** (0.002)
ζ_{jt}	0.02** (0.008)

Note: Transaction costs in log. Standard errors are in parentheses; statistical significance at 5% level is indicated with **, and at 10% with *. R^2, 0.30. Observations 1648.

expectations. The elasticities were simulated for the dynamic model as follows. First, I use the observed quantities to solve the consumer problem and 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 markets. I then re-solve the dynamic model for the optimal consumers' behavior, allowing consumers to update their beliefs. Finally, I simulate 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 own-price elasticities for new products, distinguished by market segment and country of origin. The average cross-price elasticities for new products are reported in Table 6. The cross-price elasticities are reported within market segments and for different car ages.

Table 5 reports the average own-price elasticities. The average elasticities for the static model and the full dynamic model are, respectively, −3.05 and −4.07. The estimates show that the myopic model (and ignoring transaction costs) upwardly biases the estimates of the price elasticities. The static model generates lower elasticities than the dynamic model.[30] This is because in the dynamic model, a permanent price increase today leads consumers to expect prices to be high in the future, which makes owning the product in the future more attractive (relative to whether the price had stayed the same). Thus, a large response to a temporary price change can be consistent with the data containing relatively small changes in market share in response to price changes. The static model interprets the small change as a small elasticity. Thus, when we compute elasticities as temporary price changes, they are much larger when estimating with the full dynamic model than with the static model.

Table 6 shows the average value of the cross-price elasticities. The first row of Table 6 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 0.16% increase in demand of new cars within the same segment (first column), a 0.09% increase in demand of 2-year-old cars within the same segment (second column), a 0.05% increase in demand of 5-year-old cars (third column) in the same segment, and so on. It is interesting to observe that the older the cars within each segment, the lower the cross-price elasticities: older cars are poorer substitutes for new ones.[31]

[30] More generally, however, ignoring the dynamics and transaction costs could, in principle, bias the estimate of the price elasticities in opposite directions depending on the details of the data-generating process.

[31] 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 across periods. 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.

FIGURE 4

MEAN TRANSACTION COSTS OVER TIME

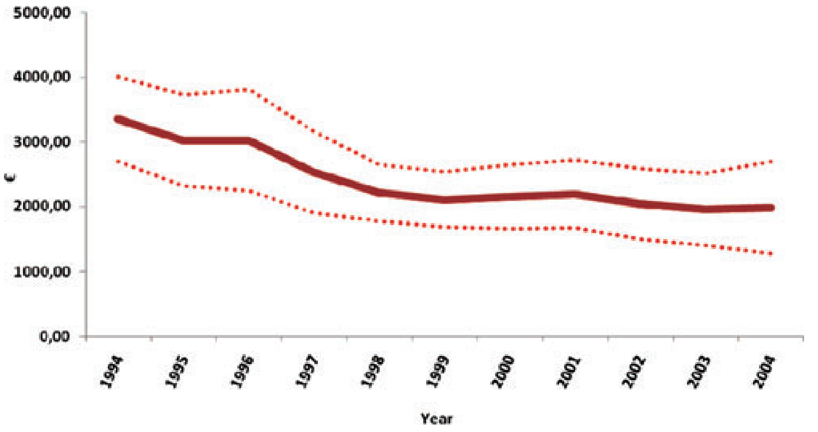

The comparison between the long-run and short-run elasticities shows that a temporary 1% increase in price determines a larger consumer response: sales will reduce on average by 10%, and consumers will purchase closer substitute models.

□ **Transaction costs.** In my model, transaction costs measure all possible frictions that 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 in Figure 4, the magnitude of transaction costs declines over time.[32] Among other factors, the effect is the result 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 €3200 in 1994 decreasing to €2000 in 2004, and the average standard deviation across time is about €600. The distribution of transaction costs is shown in Figure 5; it shows a peak in the level of transaction costs between €2100 and €2600. The minimum level of the cost is about €1000. Figure 6 shows the distribution of the transaction costs/price ratio across different models and different ages. 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 compute the standard deviation around the mean transaction costs for car models observed for more than 5 years. The average standard deviation across these models is about €400.

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.

Are these figures reasonable? According to the information published in the magazine *Quattroruote* in 1998, the explicit costs upon a transaction of a used car vary between €1000

[32] The estimate of the transaction costs is 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, which vary between €350 and €800 as specified by *Quattroruote*.

FIGURE 5

TRANSACTION COSTS DISTRIBUTION

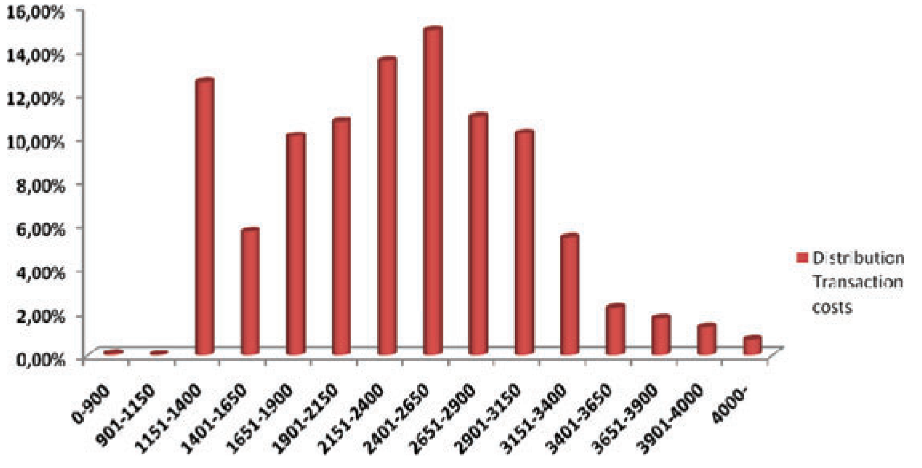

FIGURE 6

TRANSACTION COSTS/PRICE RATIO

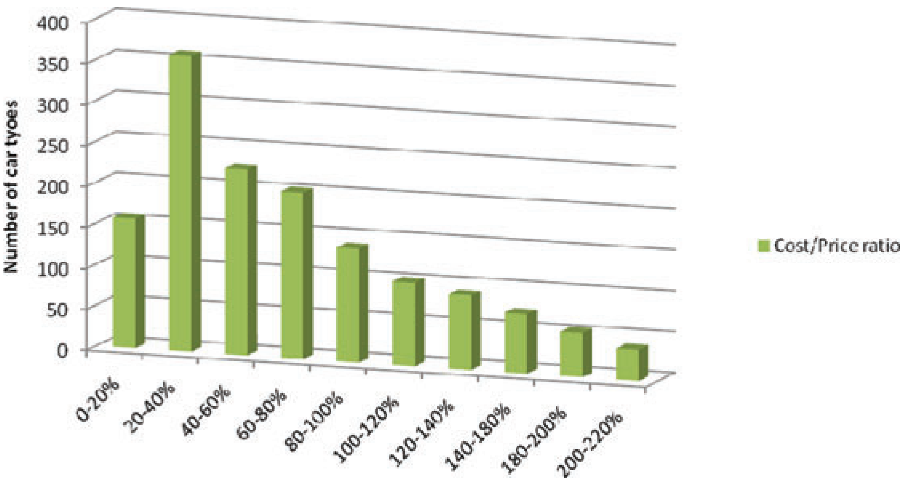

and €4000. The composition of these costs is the following: financial costs are about €400; *Quattroruote* reports that on average, the money borrowed to buy a used car in 1998 was €5000, and the spread over a safe interest rate was about 8%. The taxes and expenses to be paid upon transaction varied between €340 and €1600, according to the size and the type of car. 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 hidden costs such as search costs, asymmetric information, and so on (see also Table 8 for a comparison between the estimated transaction costs and dealer compensation as reported in *Quattroruote* relative to a few car models).

Decomposition of transaction costs. Next, I investigate the composition of the transaction costs in more detail. Table 7 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

TABLE 8 Transaction Costs: Examples

Model	Year	Age	Taxes + Dealer Compensation	Transaction Costs: Estimates
Alfa 156 1.6i	1999	1	€1675	€2000
BMW 318i	1999	3	€1700	€2400
FIAT Punto 1.9 D	2003	2	€950	€1460
Audi A3 1.6 D	2003	5	€1450	€1900

TABLE 9 Replacement Schemes in Italy: 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
Cars scrapped (Isernia)		8.40%	5.40%
Requirement	To scrap a car aged 10 years or older and buy a new one with an equal discount from the manufacturer. 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 manufacturer.	To scrap a car aged 10 years or older and buy a new one with an equal discount from the manufacturer. The discounts were awarded for a new car with average consumption < 7 l/km and average consumption <9 l/km.

in percentage terms is negative and highly significant. Specifically, an increase of 1% in the stock of cars available reduces transaction costs by €120. This relationship 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.

The variable *Diesel ∗ Time trend* captures the reduction of taxes over time relative to the car with diesel engines as discussed above. The costs increase with 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 be paid 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. It may also reflect lower maintenance costs associated with the national manufactured cars that reduce the risk and cost of buying a used vehicle.[33]

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

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

5. Scrappage policy evaluation

■ Scrappage subsidies have been particularly popular in the European Union, as well as in the United States and Canada. These policies are aimed at reducing pollution by forcing an early retirement of old and polluting cars and, in some cases, at stimulating national car industries. Typically, these subsidies are between €500 and €1500, and eligibility to participate in the program are a function of the vehicle's age.[34] 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 micromoments closely replicates 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 analysis quantifies 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.

To evaluate the replacement schemes, I perform the following analysis: I first simulate a baseline situation with no subsidies offered and 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 three different policy scenarios. The three 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 eligible 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 subsidy implemented by the Italian government for cars older than 10 years. The main finding is 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. However, different policies have different effects on new- and used-car sales.

The simulation shows that the replacement policies burst the aggregate sales of new cars followed by a contraction for few years (Figure 7). The bigger the short-run effect of the subsidy in expanding the demand, the bigger the contraction in sales that follows in the future. The model correctly accounts for the incentives provided by different policies to different consumers. In particular, the *car-for-replacement* schemes determine a contraction in new-car sales which is quite small in magnitude compared to the short-run expansion. The beneficiaries of the policies are consumers who most likely would have not bought new cars. Differently, the expansion and contraction of new-car sales are more similar in magnitude for the *cash-for-scrappage* scheme. In this case, the policy induces an early-replacement decision by consumers who would have most likely bought a new car in the future. These findings are reinforced by looking at the used-car market sales. Figure 8 shows that the *cash-for-replacement* schemes reduce sales in the used-car market, and that this negative effect on sales is more persistent than the contraction in the new-car sales. As discussed above, these policies are aimed at consumers with low income and/or low sensitivity to car age. Hence, they reduce the demand for used cars in the short and long run. As expected, the *cash-for-replacement* 8-year policy has a more negative effect than the 10-year one. 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 homogeneous and there is no active second-hand market, and hence the subsidy leads consumers to anticipate their replacement decision, causing a contraction in future sales of

[34] See European Conference of Ministers of Transportation (1999) and Hahn (1995) for comprehensive descriptions of the different scrappage subsidy programs in the United States and Europe.

FIGURE 7

EXPECTED NEW-CAR SALES RELATIVE TO THE BASELINE MODEL WITH NO SUBSIDY

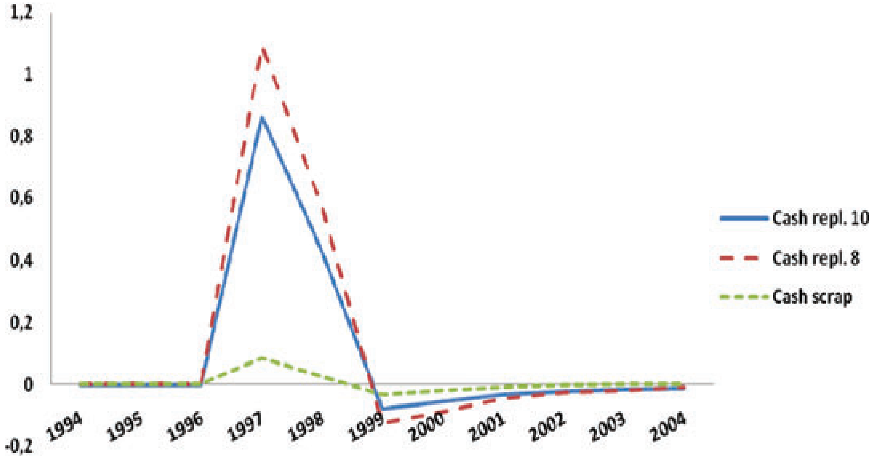

FIGURE 8

EXPECTED USED-CAR SALES RELATIVE TO THE BASELINE MODEL WITH NO SUBSIDY

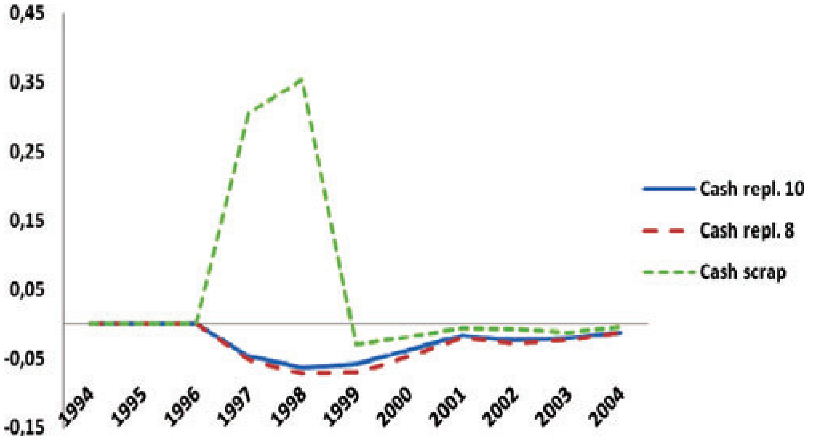

new cars. In my model, the subsidy will prevalently affect consumers who would have purchased a used car otherwise and, hence, the contraction in sales is more evident in the secondary market rather than in the primary market.

The second exercise is to study the effect of these policies on government revenues. The evaluation should account for both the short-term effect and for 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 value-added tax on the new-car sales, which was 19% in 1997 and 20% from 2000 until 2004. The province (the local government) collects a tax for each purchase of a new or 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. Figure 9 shows the net effect on government revenues of the three alternative replacement schemes compared with the baseline model. Figure 10 displays instead the net change in revenue collected by the province.

FIGURE 9

EXPECTED GOVERNMENT REVENUES RELATIVE TO THE BASELINE MODEL WITH NO SUBSIDY

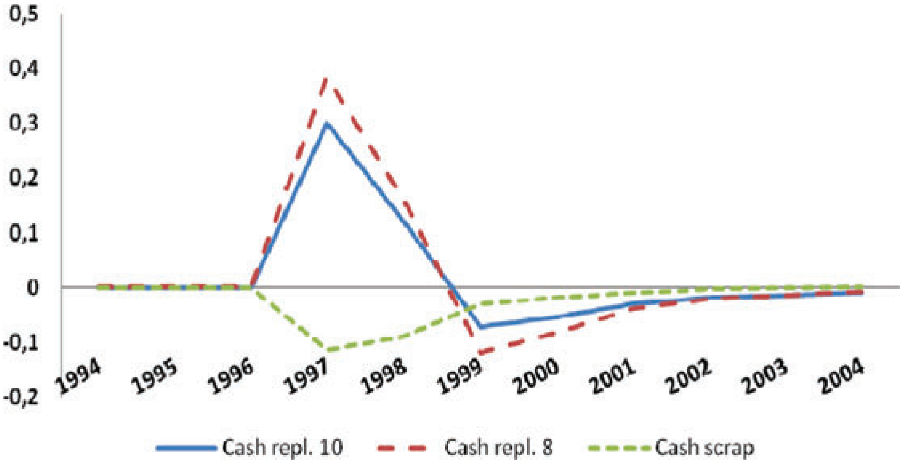

FIGURE 10

EXPECTED PROVINCE REVENUES RELATIVE TO THE BASELINE MODEL WITH NO SUBSIDY

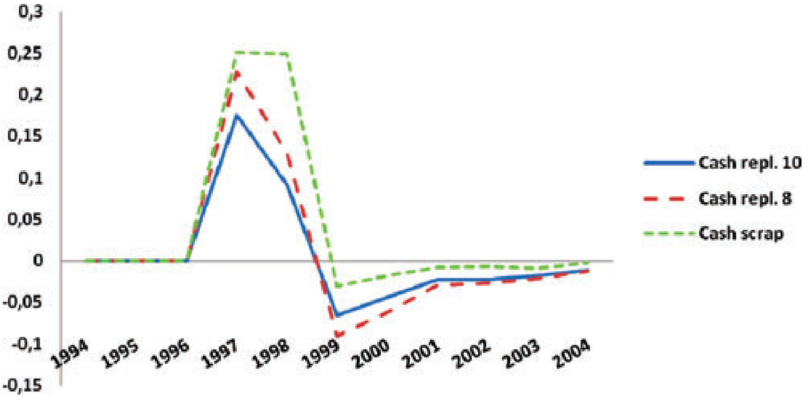

TABLE 10 Government and Province Revenue

	Baseline Model	Change after Cash for Replacement 10 Years	Change after Cash for Replacement 8 Years	Change after Cash for Scrappage
Government revenue	30,000	306	211	−807
Province revenue	4,797	−43	−54	118

Note: Value in 000s.

Finally, Table 10 reports the present discounted value of government revenues and of province revenues calculated for the baseline model, as well as the changes due to the implementation of the various policies. On average, *cash-for-replacement* policies increase the revenue of the central government, whereas *cash-for-scrappage* schemes caused 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-year policy causes a bigger increase in the revenue in the short run, compared with the *cash-for-replacement*10-year policy, the net present discounted value of revenues under the two policies are similar, due to a bigger contraction in sales under the 8-year policy. At the local government level, the impact is exactly the opposite, because the *cash-for-scrappage* policy causes an overall greater jump in total sales.

6. Concluding remarks

■ This article presents a structural model of dynamic demand for automobiles that explicitly accounts for the (costly) replacement decision of consumers in the presence of a second-hand market. The model estimates the distribution of transaction costs allowing for rational expectations about future product attributes, heterogeneous consumers with persistent heterogeneity over time, and endogeneity of prices. The analysis illustrates that taking these dynamics into account is feasible and instructive in order to account for the correct demand elasticity and to evaluate different policy designs that affect the new- and used-car markets and their dynamics.

References

ADDA, J. AND COOPER, R. "Balladurette and Juppette: A Discrete Analysis of Scrapping Subsidies." *Journal of Political Economy,* Vol. 108 (2000), pp. 778–806.

BERRY, S. "Estimating Discrete Choice Models of Product Differentiation." *RAND Journal of Economics,* Vol. 25 (1994), pp. 242–262.

——, LEVINSOHN, J., AND PAKES, A. "Automobile Prices in Market Equilibrium." *Econometrica,* Vol. 63 (1995), pp. 841–890.

——, ——, AND ——. "Estimating Differentiated Product Demand Systems from a Combination of Micro and Macro Data: The Market for New Vehicles." *Journal of Political Economy,* Vol. 112 (2004), pp. 68–105.

BRESNAHAN, T.F. "Departures from Marginal Cost Pricing in the American Automobile Industry: Estimates for 1997–1978." *Journal of Industrial Economics,* Vol. 17 (1981), pp. 201–277.

CARRANZA, J. "Estimation of Demand for Differentiated Durable Goods." Mimeo, University of Wisconsin–Madison, 2006.

DUBE, J.-P., FOX, J. T., AND SU, C.-L. "Improving the Numerical Performance of BLP Static and Dynamic Discrete Choice Random Coefficients Demand Estimation." Mimeo, University of Chicago, 2008.

ESTEBAN, S. AND SHUM, M. "Durable-Goods Oligopoly with Secondary Markets: The Case of Automobiles." *RAND Journal of Economics,* Vol. 38 (2007), pp. 332–354.

EUROPEAN CONFERENCE OF MINISTERS OF TRANSPORTATION. "Cleaner Cars. Fleet Renewal and Scrappage Scheme." European Conference of Ministers of Transport Publications, 1999.

GOLDBERG, P.K. "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry." *Econometrica,* Vol. 63 (1995), pp. 891–951.

GOWRISANKARAN, G. AND RYSMAN, M. "Dynamics of Consumer Demand for New Durable Goods." Mimeo, Boston University, 2009.

HAHN, R.W. "An Economic Analysis of Scrappage." *RAND Journal of Economics,* Vol. 26 (1995), pp. 183–196.

HENDEL, I. AND LIZZERI, A. "Interfering with Secondary Markets." *RAND Journal of Economics,* Vol. 30 (1999a), pp. 1–21.

—— AND ——. "Adverse Selection in Durable Goods Markets." *American Economic Review,* Vol. 89 (1999b), pp. 1097–1115.

—— AND NEVO, A. "Measuring the Implications of Sales and Consumer Stockpiling Behavior." *Econometrica,* Vol. 74 (2006), pp. 1637–1673.

HIDEO, K. AND SANDFORT, M.T. "Existence of Stationary Equilibrium in the Markets for New and Used Durable Goods." *Journal of Economic Dynamics and Control,* Vol. 26 (2002), pp. 1029–1052.

LEE, R.S. "Dynamic Demand Estimation in Platform and Two-Sided Markets." Mimeo, New York University, 2009.

MCFADDEN, D. "Conditional Logit Analysis of Qualitative Choice Behavior." In P. Zarembka ed., *Frontiers in Econometrics*. New York: Academic Press, 1974.

MELNIKOV, O. "Demand for Differentiated Products: The Case of the U.S. Computer Printer Market." Mimeo, Cornell University, 2001.

NELDER, J.A. AND MEAD, R. "A Simplex Method for Function Minimization." *Computer Journal,* Vol. 7 (1965), pp. 308–313.

NEVO, A. "A Practitioner's Guide to Estimation of Random Coefficients Logit Models of Demand." *Journal of Economics & Management Strategy,* Vol. 9 (2000), pp. 513–548.

PETRIN, A. "Quantifying the Benefits of New Products: The Case of the Minivan." *Journal of Political Economy*, Vol. 110 (2002), pp. 705–729.

RAPSON, D. AND SCHIRALDI, P. "Internet, Search Frictions, and the Efficiency of Decentralized Markets: Evidence from Automobiles." Mimeo, London School of Economics, 2011.

RUST, J. "Stationary Equilibrium in a Market for Durable Assets." *Econometrica,* Vol. 53 (1985), pp. 783–805.

—— "Optimal replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher." *Econometrica,* Vol. 55 (1987), pp. 999–1033.

SCHIRALDI, P. AND NAVA, F. "Second-Hand Markets and Collusion by Manufacturers of Durable Goods." Mimeo, London School of Economics, 2010.

SHCHERBAKOV, O. "Measuring Switching Costs in the Television Industry." Mimeo, Yale University, 2008.

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

SWEETING, A. "Dynamic Product Repositioning in Differentiated Product Industries: The Case of Format Switching in the Commercial Radio Industry." Mimeo, Duke University, 2007.

TAUCHEN, G. "Finite State Markov Chain Approximations to Univariate and Vector Autoregressions." *Economic Letters,* Vol. 20 (1986), pp. 177–181.