Are Lemons Really Hot Potatoes?*

Maxim Engers
University of Virginia

Monica Hartmann
University of St. Thomas

Steven Stern
University of Virginia

January 24, 2008

1 Introduction

Since Akerlof's (1971) seminal work on "lemons," economists have been examining how the presence of private information shapes market competition, particularly in the used car market. A vehicle is considered to be a lemon if its quality, as observed by the owner, is lower than the quality expected by a prospective buyer. A consumer who realizes she has purchased a lemon, other things equal, tends to sell it quicker than a car that is not a lemon. In this sense, a lemon is a "hot potato," something one does not hold onto for very long. Thus, one expects ownership spells to be shorter than average for vehicles that are lemons, and longer than average for nonlemons. The relationship between the lengths of ownership spell and price declines under adverse selection already has been derived by Hendel and Lizzeri (1999). Here we use correlation of ownership spell lengths across owners of the same vehicle in order to identify the type of adverse selection problem, owner or manufacturer-induced, (if any) that is present in the market.

We present a theoretical model that connects unobservable quality and ownership tenure length. Owners are modeled making maintenance choices each period that influence the car's quality and some component of this quality choice is unobserved by future potential buyers. The rate at which vehicles are sold is shown to decline with unobservable quality. Using data on vehicle ownership spell length, we construct a duration model to estimate the probability of selling a car conditional on its age, model year, brand, and length of ownership. This allows us to examine the correlation of spell lengths across owners of the same car to determine if an adverse selection problem is present, and, if so, what type.

*We would like to thank Amil Petrin, Bill Johnson, Jose Canals, Olivia Ceccarini, Tom Steinmeyer, and participants of seminars at Stanford Business School, the University of Colorado, and Duke University for helpful comments and the University of St. Thomas and the Bankard Fund at the University of Virginia for financial support. All errors are ours.

In the classic case, the lemons problem begins with the manufacturer. Any owner, including the first, who observes that her car is a lemon will tend to resell it more quickly than average. The reverse is true of cars whose quality turns out to be higher than expected, suggesting a positive relationship between ownership spell lengths across owners of the same vehicle. The data, however, indicate a negative relationship between the length of the first ownership spell and subsequent spells: if the first owner keeps the car for a long period, subsequent owners keep the car for a short period (this is true even after controlling for the vehicle's age at time of purchase). This implies that a classic lemons model cannot describe the asymmetric information problems present in the used car market. One possible explanation is that the lemons problem is created by the first owner. Actions taken by the first owner influence the car's unobservable quality and thus affect the rate at which subsequent owners sell their cars. While the classic lemons model is a pure adverse selection model, this alternative is a hybrid of moral hazard and adverse selection.

Our theoretical model provides one way to interpret the empirical results. Consumers who are enthusiastic about car quality sell their fairly new and defect-free cars in order to buy brand-new ones. In addition to having short ownership spells, these car enthusiasts take care of their vehicles because their utilities fall rapidly with a decline in automotive quality. On the other hand, consumers who care less about quality purchase new cars to avoid maintenance costs but do not maintain their cars as well as the car enthusiasts because their utility does not drop as much from lower quality. They sell their cars at a slower rate for the same reason. But once they sell their cars, the cars tend to be lemons because of the lack of maintenance, and future owners try to resell these cars quickly.

The layout of the paper is as follows: Section 2 summarizes the empirical literature on adverse selection. Sections 3 sets up a dynamic model of adverse selection, and Section 4 presents the implications of the model. Section 5 describes the data, while Section 6 presents the econometric methodology to estimate the model. Finally, sections 7 and 8 discuss the results and the implications for policymakers.

2 Literature on Adverse Selection

Classic theory suggests that markets plagued by information asymmetries function inefficiently or do not function at all unless institutions, such as warranties, brand names, or other signals of product quality, are developed to counteract the information problem. Yet, markets for goods that potentially suffer from adverse selection problems (e.g., used cars, health insurance) exist and thrive. Hendel and Lizzeri (1999) establish that the existence of markets for such goods follows from unobserved heterogeneity in preferences or technology across market participants.[1] They also demonstrate that one needs only information on

[1] Empirical papers, such as Rosenman and Wilson (1991), Chezum and Wimmer (1997, 2000) and Genesove (1993), use this heterogeneity to identify the presence of adverse selection.

price and volume of trade to identify if a market suffers from adverse selection.[2] It is present if the products with the greatest decline in prices over time are the ones that are traded the least often in the secondary market. This test, however, is based upon a model that assumes that goods are traded only once. Hendel, Lizzeri, and Siniscalchi (2005) present an alterative test, one that is based on a dynamic model with stochastic depreciation where consumers observe the sales histories.[3] They show that the number of times a car has been sold perfectly signals its unobservable quality. Unfortunately, buyers do not observe sales histories for many durable items, including used cars until quite recently.[4] Thus, consumers cannot infer information about unobservable quality from the number of times it has been sold and another empirical test is needed for these markets.

One such test is found in Emons and Sheldon (2003). They use Swiss data on lengths of ownership spells. They propose that unobserved car problems appear according to a Poisson process. Car owners use the revealed number of problems to estimate the actual number of problems not yet revealed. When the total number of estimated problems exceeds a threshold, the owner sells. A negative duration dependence in the hazard rate is estimated. But because the authors do not control for consumer-specific unobserved heterogeneity such as differences in owners' tolerance for car problems, there is a negative duration dependence bias. A test for adverse selection must also then account for consumer-specific unobservable heterogeneity in the timing of sales.

In our paper, we pose and implement a new empirical test for adverse selection. We present a dynamic model of buying and selling of durable goods in a market where sales histories are not observed. We control for how variation in the willingness to pay for quality affects the timing of sales and how this timing may be correlated across owners of the same vehicle. We also introduce another source of heterogeneity - heterogeneity in vehicle age, new versus used, in addition to the standard inclusion of heterogeneity in private information on product quality. Both Finkelstein and Poterba (2004) and Finkelstein and McGarry (2006) demonstrate this additional source of heterogeneity is necessary when measuring the effect of private information on market competition. Finally the test we propose not only addresses these parameterization issues concerning unobservable heterogeneity, but it also requires minimal data, just data on ownership spell lengths, for identification, thus, making it an attractive test for adverse selection.

---

[2] Porter and Satler (1999) have a model with deterministic exogenous depreciation leading to a correlation between price and the volume of trade. Gilligan (2004) applies Hendel and Lizzeri's model to the market for used business aircraft.

[3] Actually, cars are rented rather than bought and sold in this model.

[4] Sales histories generally were unobserved by both consumers and car dealers until the early 1990s, and observed only by dealers from the early 1990s to the late 1990s, and by both since the late 1990s with the introduction of services like CARFAX.com.

3 Model

The goal of this section is to set up a model of households deciding whether to keep a vehicle or sell it and buy another one. We show that their optimal selling behavior implies that lemons are hot potatoes. In other words, cars with bad privately observable characteristics are sold faster than cars with good privately observable characteristics. The theoretical model is then later used to provide insights into the empirical results from the hazard rate model; i.e. to aid in understanding the relationship between unobservable heterogeneity and ownership spell length. Once this relationship is understood, we can discuss how different spell length patterns across spells within a car can be used to learn something about the nature of how lemons are created. However, there are other stories that can explain the same facts about hazard rates, and some of these are discussed in Section 7.2. The primary benefit of the presented model is to motivate our focus on hazard rates and to aid in interpreting the results.

Each car is characterized by (b, v, a, t, e) where b is the brand of the car, v is the vintage, a is the age of the car when it was purchased, t (for tenure) is how long the car has been owned, and e is a scalar representing the effect of all car-specific variables that are observable only by the owner. When a car has a small value of e, we call it a lemon. The equilibrium price p of a car is a function of characteristics observable in the market, $p = p(b, v, a + t)$.[5] The utility flow that an owner gets from a car is $U(b, v, a, t, e)$, which can vary over individuals. It is assumed that $U_a \leq 0$, $U_t \leq 0$, and $0 \leq U_e \leq 1$, $U_{ee} < 0$, and U is nonnegative and bounded.[6] The first two assumptions follow if automobile maintenance costs increase with age (see Engers, Hartmann, and Stern, 2004) and automobile benefits decrease with age (see Engers, Hartmann, and Stern, 2007). The assumption that $0 \leq U_e \leq 1$ is innocuous, simply defining units of e and specifying e as a "good." The assumption that $U_{ee} < 0$ states that the marginal utility of increasing e is diminishing and, together with the boundedness assumption, implies that U_e tends to zero as e gets large.

The owner of a car with tenure t may increase the level of e from e_t to e_{t+1} where $e_{t+1} \geq e_t$ by incurring elective maintenance costs $C(e_t, e_{t+1})$.[7] We include a choice for the owner to allow for owner-created lemons. We assume that $C(e_t, e_t) = 0$, $C_1 < 0 < C_2$, $C_{12} \leq 0$, and that C is convex so that $C_{11} \geq 0$, $C_{22} \geq 0$, and $C_{11}C_{22} \geq C_{12}^2$. The assumption that $C_{12} \leq 0$ states that the marginal cost of increasing e_{t+1} does not rise as e_t rises. In the special case in which C depends only on the change in e, namely $e_{t+1} - e_t$, this amounts to assuming that marginal costs do not decrease as the change

[5] We assume there are no other characteristics of the car observable in the market but not by us.

[6] Although all variables except e are discrete, for brevity, we state monotonicity assumptions in terms of derivatives.

[7] We can generalize and allow for required maintenance costs as well with no change in the qualitative results. Also, we could add a second discretionary component of maintenance observable to potential buyers as long as the first component is not observable until after purchase.

in e rises. Convexity in this special case has the same interpretation. More generally, convexity requires that the costs of changing from e_t to e_{t+1} grow at a nondecreasing rate as e_t decreases and as e_{t+1} increases, and that the absolute value of the cross derivative C_{12} is not too large relative to C_{11} and C_{22}.[8]

Once the owner has paid $C(e_t, e_{t+1})$, she can keep the car or sell it and receive price $p(b, v, a+t)$. The value of keeping a car at t with state variables $S_t = (b, v, a, t, e_t)$ is

$$V(S_t) = \max_{e_{t+1}} \{U(b, v, a, t, e_t) - C(e_t, e_{t+1}) + \beta \max[V(S_{t+1}), W(S_{t+1})]\} \quad (1)$$

where $\beta \in (0,1)$ is the discount factor and $S_{t+1} = (b, v, a, t+1, e_{t+1})$. Similarly, the value of selling the car is

$$\begin{aligned} W(S_t) &= \max_{e_{t+1}} [-C(e_t, e_{t+1}) + p(b, v, a+t)(1-\psi)] \qquad (2) \\ &\quad + \max_{b',a',v'} \int [V(b', v', a', 0, e') - p(b', v', a')] \, dF_e(e' \mid b', v', a') \end{aligned}$$

where ψ denotes an *ad valorem* transactions cost.[9] The last term assumes that the owner purchases another car and immediately observes e.[10] In particular, she purchases a car that maximizes expected value given the conditional distribution of e', $F_e(e' \mid b', v', a')$. Because we assume that e is unobservable to potential buyers, it does not affect price. Note that $F_e(e' \mid b', v', a')$ captures any lemons effects. The owner sells if $W(S_t) > V(S_t)$. Note that the solution to the maximization problem (2) is to set $e_{t+1} = e_t$ (there is no point to increasing e if you're selling), and so

$$\begin{aligned} W(S_t) &= p(b, v, a+t)(1-\psi) \\ &\quad + \max_{b',a',v'} \int [V(b', v', a', 0, e') - p(b', v', a')] \, dF_e(e' \mid b', v', a'), \end{aligned}$$

which is independent of e_t. We assume that, as t gets large, $\beta^t W(S_t)$ tends to zero. This condition will hold if car prices and the distribution of new vehicles are stationary, or, more generally, as long as the rate of appreciation in both prices and the value derived from choosing a new car have an upper bound below β^{-1}.

4 Implications of the Model

The goal of this section is to derive the property that a car with a large value of e has a lower hazard rate than one with a small value of e. Instead of trying to

[8] One might worry that there is no explicit depreciation considered in our modelling of the costs of increasing e. However, to the degree that depreciation can be captured by aging or an interaction of aging with e, it is accounted for in $U(b, v, a, t, e)$.

[9] Instead of the transaction costs being proportional to car value, they could be fixed costs or have components of both with no significant change in model behavior.

[10] Emons and Sheldon (2003) assume instead that a new owner learns about the analog to e only over time.

characterize all aspects of equilibrium, we focus on this one goal because it is the critical connection necessary for interpreting the empirical results. Proofs are in the appendix. At the end of this section, we discuss properties of equilibria.

First, we show that value functions are well behaved.

Proposition 1 *The value function V is strictly increasing and continuous in e_t.*

Next, we show that behavior is characterized by a reservation value such that an owner sells her car iff e_t is less than the reservation value.

Proposition 2 $\exists e^* \in [0, \infty] : W(S_t) > V(S_t)\ \forall e_t < e^*$ *and* $W(S_t) < V(S_t)$ $\forall e_t > e^*$.

If $e_t = e^* > 0$, the owner is indifferent between holding on to the car and selling it. As a tiebreaking rule, we will suppose that the car is sold.

In order to understand how selling rates depend on e_t, we must understand first how e_t changes over time. First, we show that e_{t+1} is increasing with e_t.

Proposition 3 *The optimal choice of e_{t+1}, which we denote by $G(e_t)$, is increasing in e_t.*

Corollary 4 *Define $G_s(e_t)$ recursively by $G_0(e_t) = e_t$, and $G_s(e_t) = G(G_{s-1}(e_t))$. Then G_s is increasing for each s.*

Now we can show that time until sale is increasing in e_t, implying that hazard rates are decreasing in e_t.

Proposition 5 *Let $t^*(b, v, a, e)$ be the tenure when the car is sold. Then $t^*(b, v, a, e'') \geq t^*(b, v, a, e')\ \forall e'' > e'$.*

4.1 Stochastic Depreciation

We have assumed that, once e_t is given, the subsequent value e_{t+1} is selected deterministically as a result of the owner's optimization decision. An alternative is to introduce a stochastic element into the evolution of the e_t term. Stochastic depreciation is an important element of some models of adverse selection in durable goods markets. See, for example, Hendel, Lizzeri, and Siniscalchi (2005).[11]

A simple way to introduce stochastic depreciation is to maintain our current model, but to suppose that now, instead of deterministically choosing next period's value e_{t+1}, the owner instead chooses μ_{t+1}, and the actual value of e_{t+1} is drawn from a distribution $F(e_{t+1} \mid \mu_{t+1})$. We assume that $F(e_{t+1} \mid \mu_{t+1})$ is continuous in both arguments, has compact support, and, for $\mu_{t+1} > \mu'_{t+1}$,

[11] However, Porter and Satler (1999) have a model with nonstochastic depreciation generating interesting results. Nonstochastic depreciation is captured in our model by controlling for age and brand.

the distribution $F(e \mid \mu_{t+1})$ dominates $F(e \mid \mu'_{t+1})$ in the sense of (first order) stochastic dominance. The cost of maintenance function C is assumed to have exactly the same properties as before, but its second argument is now μ_{t+1} instead of e_{t+1}.

With these assumptions, the main implications of the model are preserved with the appropriate modifications. For example, an increase in e_t no longer ensures that the optimal e_{t+1} is no lower than before, but now ensures that the new value of e_{t+1} is drawn from a distribution that weakly stochastically dominates the old distribution. With stochastic depreciation, the tenure when the car is sold $t^* (b, v, a, e)$ is random, but an increase in e will shift its distribution upward in the sense of stochastic dominance. So the higher levels of e are still associated with longer ownership spells.

4.2 Persistence of Unobserved Quality Across Different Owners of the Same Car

Proposition 5 shows that lemons are "hot potatoes." However, the theory does not address what happens to the long-term unobserved quality of cars as they age and are transferred from one owner to another. Consider the stochastic e process described above with cost function $C\left(e_t, \mu_{t+1}\right)$ and distribution function $F(e_{t+1} \mid \mu_{t+1})$. It may be that there is no persistence in unobserved quality associated with a car. Consider the special case where the cost function has the property that, for some μ^* that might depend on (b, v, a, t),

$$\begin{aligned} C\left(e, \mu^*\right) - C\left(e, e\right) &< EU\left(b, v, a, t, \mu^* + \eta\right) - EU\left(b, v, a, t, e\right) \quad \forall e < \mu^*; \\ C\left(e, \mu^*\right) - C\left(e, e\right) &> \frac{EU\left(b, v, a, t, \mu^* + \eta\right) - EU\left(b, v, a, t, e\right)}{1-\beta} \quad \forall e > \mu^*. \end{aligned}$$

μ^* would correspond to the "appropriate" level of maintenance; it would be very worthwhile to invest in a car up to μ^* but not worthwhile beyond μ^*. In many ways, this seems the most reasonable specification for the cost function; when a muffler breaks, you fix it, but you rarely replace it with the super-duper gold-plated muffler. Then, any car purchased with $e < \mu^*$ would be a lemon in that it would require some repair by the new owner. But the new owner in fact would invest in the car up to μ^*. There still would be variation in how long owners held onto cars partially because of variation among car owners in preferences over car characteristics. However, spell lengths across owners of the same car would be independent of each other after controlling for observable characteristics, especially age.

Alternatively, it may be that unobserved quality is determined at the time of manufacture and is difficult to alter. This would require that the distribution of e for new cars $F_e\left(e \mid b, v, 0\right)$ (which is exogenous to the model) has a large standard deviation relative to $C_2^{-1}(e_t, \mu_{t+1})$, the inverse of the cost of increasing e. In the limiting case, as $C_2(e_t, \mu_{t+1})$ becomes very large, it becomes too expensive to repair a car with a small value of e. Thus, the car is stuck with the low e over its lifetime. However, since e is not observed by anyone other than the owner, the car is sold with high turnover over many owners; thus we would see positive correlation in spell lengths after conditioning on observable car characteristics. Alternatively, unobserved quality could be determined to a great degree by care provided by owners. This would require that the distribution $F_e\,(e \mid b, v, a)$ has a small standard deviation relative to $C_2^{-1}(e_t, \mu_{t+1})$. Whether it would lead to positive or negative conditional correlation would depend upon the properties of $C_2(e_t, \mu_{t+1})$ and the distribution of preferences in the population. In section 6, we describe an econometric methodology used to distinguish between these alternative stories of how unobserved quality of cars varies with age. More specifically, we test three models against each other:

1. There is no important persistence in unobserved quality in car tenure across owners. Note that such a case does not rule out the existence of unobserved quality associated with a car-owner combination.

2. There is persistence in unobserved quality variation across cars, and it is determined at the time of manufacture. This would imply that there is a car-specific unobserved heterogeneity constant across owners of the car.

3. There is persistence in unobserved quality variation across cars, and individual owners (especially the original owner) affect its size. This could be inferred if the unobserved heterogeneity were to change in a systematic way across owners of the car.

4.3 Eliminating Private Information

Now let us examine how robust these results are. For instance, if we eliminate the private information in the model, how does the equilibrium outcome change? Specifically, if the e term is observed by all, instead of being observed only by the owner, how are the results affected? There is now a motive for increasing the value of e even when selling a car, and (in comparison to the case in which e is not observed by all) it seems likely that there is a greater incentive to raise the level of e of one's car in the periods before sale. Thus one would expect that there would be higher average equilibrium levels of e than before. The fact that prices now can reflect differences in the value of e makes it hard to prove direct comparisons between the old and the new equilibrium. The value from holding onto a car is still increasing in e, but now the value from selling also can be increasing in e. So, given market prices, one might be more willing to sell a car with a higher value of e, unlike the case with private information, where there was a single threshold level of e below which the car would be sold and above which the car would be kept. Supply is more complicated now.

Intuitively, when e is common knowledge, the source of adverse selection, an impediment to trade, is removed. Thus one might expect that the equilibrium volume of trade would rise, which is to say that cars would turn over more quickly. Here again, while the statement seems reasonable, it is certainly not easy to prove. While one is less likely to hold onto a high-e car when one receives a premium price for it, the incentive to sell very low-e cars is reduced because their prices now reflect their low quality.

For the econometrician who can't observe e but who can observe prices, the fact that buyers can observe e increases the variance of prices for a given type of car. One crucial question is whether a model without private information can generate the pattern of lengths of ownership spells that we have found in the data. This is discussed in the Results section.

4.4 Ensuring Trade Occurs

Because of adverse selection, it is possible that no trade would occur in equilibrium. We can ensure that some trade occurs in equilibrium by assuming that there is sufficient variation in the individual-specific component of the utility flow provided by a car relative to other sources of variation. The appropriate condition needs to be assumed only for a subset of owners of positive measure and need not hold for all owners.

Let $U_i(b, v, a, t, e) + \epsilon_{ikt}$ denote the utility flow to agent i from owning car k with characteristics (b, v, a, t, e) where b, v, a, t, and e are as before. Let M_i be an upper bound on U_i for all cars k. Suppose that the ϵ_{ikt} are independent for all i, k, and t and, for given i, ϵ_{ikt} has distribution F_i with support S_i. Let $L_i = \sup S_i - \inf S_i$ be the (possibly infinite) length of the support of F_i.

Proposition 6 *Suppose that for each $M > 0$, there exists a set S of individuals i where S has positive measure and $L_i/M_i > M$ for all $i \in S$. Then in any equilibrium there will be a positive volume of trade in each car in each period.*

Note that if, for a positive measure of individuals i, $L_i = \infty$ and $M_i < \infty$, the sufficient condition is automatically satisfied.

5 Data

We use a random sample of 3543 cars in Virginia to distinguish among the three models of how unobservable heterogeneity affects the timing of car sales. For each car, we observe all sales from time 0 (September 13, 1993) to T (May 13, 2001) along with how long the owner at time 0 owned the car. Time T is the date the Department of Motor Vehicles (DMV) of Virginia collected the random sample for us, and time 0 is the earliest date for which they had electronic records available on May 13, 2001. Thus, we observe, for each car i, $\{t_{ij}, d_{ij}\}_{j=1}^{J_i}$ where t_{ij} is the length of the jth ownership of car i (starting at $j = 1$ with the owner at time 0) and d_{ij} is a generalized censoring dummy discussed later. Each record in this random sample also contains the VIN (vehicle identification number). The VIN is a 17 digit alphanumeric string that provides information on the make and manufacture year of the vehicle and other characteristics of the car (e.g., engine size, brake system, manufacturing plant, etc.). We can deduce from these data whether the first observed owner ($j = 1$) is the initial owner of the car by comparing the manufacture year to the date that the $j = 1$ owner purchased the car. Our dataset is similar to the one used by Porter and Sattler (1999). Unlike their study, we do not have information on the odometer reading at time of sale.

We have three types of ownership spells in our data as illustrated in Figure 1. We call ownership spells in the data "observed spells" and distinguish them from "real spells." Consider a car that was manufactured in 1990, sold in 11/92, and sold again in 1/98. Our data would have an observed "initial spell" t_{i1} equal to the number of days[12] from 6/90 until 11/92. During this period we know only that there was at least one owner (i.e., at least one real spell) and that the owner in 11/92 sold it. Next there would be an observed (and real) "middle spell" t_{i2} from 11/92 to 1/98 corresponding to an observed owner, i.e., one whose ownership is observed by DMV. Finally, there would be an observed and real "end spell" t_{i3} from 1/98 to 5/01. The observed initial spell would be censored in a unique way in that we know only that the length of the real initial spell is no more than t_{i1}. It could have been less than t_{i1} because, during that period, there could have been multiple owners. The second spell is uncensored, and the third is right censored in the usual way. Alternatively, consider a car that was manufactured and purchased in 1992, sold in 10/96, and then again in 8/00. Our data would have t_{i1}, the uncensored tenure of the initial owner (equals the number of days between the new purchase date and its sale in 10/96), t_{i2}, the uncensored tenure of the owner between 10/96 and 8/00, and t_{i3}, the censored tenure of the owner between 8/00 and 5/01. In general, we define a variable d_{ij} to indicate the type of observed spell: $d_{ij} = 1$ if the jth observed spell is an "initial spell," $d_{ij} = 2$ if the jth observed spell is a "middle spell," and $d_{ij} = 3$ if the jth observed spell is an "end spell."

Table 1 reports the number of observed spells of each type and the moments of each. There are 8584 observed spells. The average spell length is roughly two to four years, depending on the type of spell. In our empirical section, we condition on the brand. Table 2 reports the number of observations by brand. We have aggregated Cadillac and Lincoln into "American Luxury," Audi, BMW, Jaguar, Mercedes-Benz, Porsche, and Saab into "European Luxury," Infiniti and Lexus into "Japanese Luxury," and Acura with Honda. Some rarer brands such as Peugeot and Fiat are excluded. We include only cars that were manufactured after 1985 because, in related work, we merge these data with Kelley Blue Book price data, which are available only back to 1986.

Figure 2 shows the Kaplan-Meier survival curves for owners of new cars and owners of used cars. The Kaplan-Meier curve for new cars is biased upwards because many of the spells used to construct the estimate are initial spells ($d_{ij} = 1$) and may consist of many hidden spells. The curves show that 13% of new cars and 30% of used cars are sold within the first year. After 3 years, 65% of new cars and 29% of used cars are still owned by their owners. The Kaplan-Meier curve does not allow for observed (brands or age) or unobserved variation, and it does not correct for the "initial spells" problem.

[12] When date of purchase for the initial owner is not observed, we assume arbitrarily that it is 6/30/yy where yy is the manufacture year.

6 Econometric Methodology

The theoretical model in Section 3 establishes the relationship between ownership tenure and lemons (i.e., cars with low values of e). Here we present a hazard rate model to identify how the long-term unobserved quality affects ownership tenure over the life of a vehicle. The econometric model has many of the same features as the theoretical one. Vehicles are sold when the unobservable quality is below the household's reservation value. The rates at which vehicles are sold are a function of the car's observable as well as its unobservable characteristics. But, unlike the theoretical model, it can identify if unobservable car quality is constant across the same vehicle, indicating manufactured-induced unobserved quality, or varies across owners of the same vehicle, suggesting owner-induced unobserved quality.

6.1 Estimation

Let

$$h_{jvb}(t,a) = \exp\{\delta_b + \lambda(a,v,t) + u_j\}$$

be the hazard rate for the jth sale at tenure t for cars of vintage v and brand b that were purchased at age a. Brand affects the hazard rate through the brand effect δ_b. Tenure affects the hazard rate through the baseline hazard $\lambda(a,v,t)$ which can vary with age at purchase a. Finally, unobserved effects (possibly) specific to the jth owner affect the hazard rate through u_j. Such an effect typically is called "unobserved heterogeneity" in the econometric survival analysis literature and should be distinguished from randomness implicit in the hazard function itself since the hazard function is a conditional probability. Following and generalizing Meyer (1990), we specify $\lambda(a,v,t)$ as a spline function in levels in t and a. Define $k^*(t)$ to satisfy $\tau_{k^*} \leq t < \tau_{k^*+1}$ where τ_k, $k = 1,2,..,K$ are nodes in t. Similarly define $m^*(a)$ to satisfy $\gamma_{m^*} \leq a < \gamma_{m^*+1}$, where γ_m, $m = 1,2,..,M$ are nodes in a. Then the baseline hazard function is

$$\lambda(a,v,t) = \sum_{k,m} \lambda_{vkm} 1\left(\tau_k \leq t < \tau_{k+1}, \gamma_m \leq a < \gamma_{m+1}\right). \tag{3}$$

Following Heckman and Walker (1990), we model the unobserved heterogeneity as $u_j = \alpha_j u$ where α_j is allowed to vary over j and u is a car-specific unobserved heterogeneity term. To simplify calculation of the likelihood function, we assume that

$$\alpha_j = \begin{cases} 1 & \text{if } j = 1 \\ \alpha & \text{if } j > 1 \end{cases}. \tag{4}$$

This specification allows the unobserved heterogeneity term to differ between the initial spell (new cars) and subsequent spells (used cars), it reduces the order of the integral involved in evaluating the hazard and survivor functions, and it allows us to circumvent the issue that we do not know how many owners are in initial spells. Thus, the conditional hazard rate for tenure t cars of vintage v is

$$h_{jvb}(t, a \mid u) = \exp\left\{\delta_b + \sum_{k,m} \lambda_{vkm} 1\left(\tau_k \leq t < \tau_{k+1}, \gamma_m \leq a < \gamma_{m+1}\right) + \alpha_j u\right\}, \tag{5}$$

and the survivor function is

$$\begin{aligned} & S_{jvb}(t, a \mid u) \\ = \; & \exp\left\{-\int_0^t h_{jvb}(s, a)\, ds\right\} \\ = \; & \exp\left\{-\exp\{\delta_b + \alpha_j u\}\left[\sum_{k<k^*(t)} \exp\left\{\lambda_{vkm^*(a)}\right\}[\tau_{k+1} - \tau_k] + \exp\left\{\lambda_{vk^*(t)m^*(a)}\right\}\left[t - \tau_{k^*(t)}\right]\right]\right\}. \end{aligned}$$

Let $K_v = \{i : i\text{'s vintage is } v\}$. Define t_{ij} as the observed length of the jth spell of observation i, and d_{ij} as an indicator of what "type" of spell it is ($d_{ij} = 1$ if initial, $d_{ij} = 2$ if middle, $d_{ij} = 3$ if censored). Then the log likelihood function for cars of vintage v is

$$\begin{aligned} L_v \; = \; & \sum_{i \in K_v} \ln \int \left[1 - S_{jvb}(t_{ij}, a \mid u)\right]^{1(d_{ij}=1)} \cdot \\ & \left[S_{jvb}(t_{ij}, a \mid u)\, h_{jv}(t_{ij}, a \mid u)\right]^{1(d_{ij}=2)} \left[S_{jvb}(t_{ij}, a \mid u)\right]^{1(d_{ij}=3)} f(u)\, du \end{aligned} \tag{6}$$

where $f(u)$ is the density of the unobserved heterogeneity. Note that the contribution of an "initial spell" ($d_{ij} = 1$) is $1 - S_{jvb}(t_{ij}, a \mid u)$. This is because all we know about the initial spell is that its length is *at most* t_{ij}. In the results, we present estimates for normally distributed unobserved heterogeneity. The log likelihood function is

$$L = \sum_v L_v. \tag{7}$$

We also tried using a Heckman-Singer (1984) specification of unobserved heterogeneity and got similar and imprecise estimates.

We also might want to allow for plant or plant-year specific effects. But, the asymptotics for plant effects require that the number of plants (as opposed to the number of cars per plant) go to infinity which is not a good assumption. Furthermore, the small size of the estimated standard deviation of the car-specific unobserved heterogeneity argues against using a richer specifications of unobserved heterogeneity.

One might be concerned that the specification of the joint density of the unobserved heterogeneity in equation (4) is too restrictive. Later, in Section 7.2, we discuss how to estimate the covariance matrix of generalized residuals and show that the specification in equation (4) is not driving our results.

Let $\theta = (\delta, \lambda, \alpha, \sigma_u)$ be the set of parameters to estimate, where δ is the vector of brand effects in equation (5), λ are the spline levels in equation (3), α is the vector of coefficients on the unobserved heterogeneity in equation (5), and σ_u is the standard deviation of the unobserved heterogeneity. Note that this implies that the τ-nodes and γ-nodes in equation (3) are chosen prior to estimation. The value of θ that maximizes the log likelihood function in equation (7) is the maximum likelihood estimator. It is consistent, efficient, and asymptotically normal with the standard asymptotic covariance matrix.

6.2 Interpretation and Identification

The model requires that a car is sold the first time that $e_t < e_t^*$. The distribution of e_0 is $F_e(e \mid b, v, a)$, and, conditional on e_0, e_t evolves deterministically according to the optimal program. Recall that, given e_t, the optimal value of e_{t+1} is denoted by $G(e_t)$. Let $e_t = G_t(e_0)$ where G_t is defined in Corollary 4. Thus, the distribution of e_t (evaluated at e_t^*) among those cars still with the same owner at t is

$$\begin{aligned} \Pr[e_t < e_t^*] &= \Pr\left[G_t(e_0) < e_t^* \mid e_{t-1} \geq e_{t-1}^*\right] \\ &= \Pr\left[G_t(e_0) < e_t^* \mid G_{t-1}(e_0) \geq e_{t-1}^*\right] \\ &\quad \Pr\left[e_0 < G_t^{-1}(e_t^*) \mid e_0 \geq G_{t-1}^{-1}\left(e_{t-1}^*\right)\right] \\ &= \frac{F_e\left(G_t^{-1}(e_t^*) \mid b, v, a\right) - F_e\left(G_{t-1}^{-1}\left(e_{t-1}^*\right) \mid b, v, a\right)}{1 - F_e\left(G_{t-1}^{-1}\left(e_{t-1}^*\right) \mid b, v, a\right)}. \end{aligned} \tag{8}$$

The probability of a sale (i.e., the hazard rate) is denoted by $Q_{te}\left(e_t^*, e_{t-1}^* \mid b, v, a\right) = \Pr[e_t < e_t^*]$.

In equation (5),[13]

$$\partial \log h_{jvb}(t, a \mid u) / \partial b = \delta_b.$$

Using equation (8),

$$\delta_b = \partial \log Q_{te}\left(e_t^*, e_{t-1}^* \mid b, v, a\right) / \partial b$$

with $\partial Q_{te}\left(e_t^*, e_{t-1}^* \mid b, v, a\right) / \partial b$ depending on direct effects such as $\partial F_e\left(G_{t-1}^{-1}\left(e_{t-1}^*\right) \mid b, v, a\right) / \partial b$ and indirect effects such as

$$\frac{\partial F_e\left(G_t^{-1}(e_t^*) \mid b, v, a\right)}{\partial G_t^{-1}(e_t^*)} \left[\frac{\partial G_t^{-1}(e_t^*)}{\partial b} + \frac{\partial G_t^{-1}(e_t^*)}{\partial e_t^*} \frac{\partial e_t^*}{\partial b}\right]$$

and a similar term involving $G_{t-1}^{-1}\left(e_{t-1}^*\right)$. The direct effects measure the effect of brand on how adverse selection changes the distribution of cars for sale with characteristics (b, v, a). The indirect effects measure the effect of brand on maintenance expenditures up to tenure t. The terms in brackets can be decomposed into a brand effect on $G_t^{-1}(e_t^*)$ holding e_t^* constant and a brand effect on e_t^*. Without maintenance data, we can identify only the sum of the three effects; we can not separately identify them.

[13] We are using partial derivatives even though "brand" is a dummy variable.

The analysis for tenure effects is similar. In particular,[14]

$$\partial \log h_{jvb}\left(t, a \mid u\right) / \partial t = \Delta \lambda_{vkm}.$$

Equation (8) requires that $\partial Q_{te}\left(e_t^*, e_{t-1}^* \mid b, v, a\right) / \partial t$ depends on

$$\frac{\partial F_e\left(G_t^{-1}\left(e_t^*\right) \mid b, v, a\right)}{\partial G_t^{-1}\left(e_t^*\right)}\left[\frac{\partial G_t^{-1}\left(e_t^*\right)}{\partial t} + \frac{\partial G_t^{-1}\left(e_t^*\right)}{\partial e_t^*}\frac{\partial e_t^*}{\partial t}\right]$$

and a similar term involving $G_{t-1}^{-1}\left(e_{t-1}^*\right)$. Again, without maintenance data, we can identify only the sum of the effects.

However, we can consider some special cases of interest. Let T_1 and T_2 be two adjacent spell lengths for the same car. If $\sigma_u = 0$ in equation (5), then T_1 and T_2 are independent of each other after conditioning on (b, v, a_1) where a_1 is the age of the car at the beginning of the first ownership. If there is any variation in e across cars that remains with a car as it changes owners (persistence in unobserved quality), then T_1 and T_2 would be dependent. Thus $\sigma_u = 0$ implies that there is no lemons problem. If $\sigma_u = 0$, then the variation in spell length across owners can not be a result of unobserved car characteristics; otherwise it would carry over from one owner to the next. Instead, the variation in spell length is caused by variation in preferences across owners. More precisely,

$$F_e\left(e \mid b, v, a\right) = 1\left(e \geq e_0\right).$$

Variation in e_t^* across owners caused by variation in preferences causes variation in spell length.

Consider the particular specification of α_j in equation (4). We use this specification in the empirical analysis because it allows us to distinguish between the behavior of new car owners and used car owners while not relying on nonlinear functional form assumptions to identify α. Instead, the correlation of spell length across ownership spells within the same car with multiple owners identifies α. If $\alpha > 0$, then a long first spell length increases the probability of long subsequent spell lengths. This suggests that part of e is specific to the car and possibly is determined at the time of manufacture. If $\alpha = 1$, then all owners behave the same with respect to a car's e. This suggests that unobserved quality at manufacture, owners find it hard to change e, and there is a lack of interaction in utility between e and a. On the other hand, if $\alpha < 0$, then an explanation is that the absolute value of the derivatives of utility with respect to e and t are positively correlated among new car owners. This correlation leads some new car owners to invest a lot in maintaining their car and still decide to sell quickly and others to invest little and hold the car for a while. The variation in investment affects tenure of subsequent owners and causes the negative correlation.[15] Note that such an effect would have to dominate the natural tendency of an owner to invest little in a car she expects to sell soon. Also note that it would be difficult for any important long-term unobserved quality effects due to the manufacturer to exist with $\alpha < 0$.

[14] For this to be nonzero, t must change by enough so that the relevant spline cell changes.

[15] This idea was first suggested to us by Shannon Mitchell as she observed how poorly her husband, Maxim Engers, maintained his car, which he then sold after 12 years.

7 Results

7.1 Estimates

We test three alternative models of the role unobservable heterogeneity plays in the timing of used car sales. In the first model, we assume there is no unobserved quality associated with the timing of sales. In the second model, we specify the persistent unobserved quality as determined at manufacture, while in the third model, persistent unobserved quality is assumed to be determined by the first owner.

First, we estimate a model with no unobserved heterogeneity in the hazard function (we fix $\sigma_u = 0$). Table 3 presents results when one assumes there is no unobserved quality associated with the timing of sales. Because there are too many variables to report, only coefficients for the brand and age variables (but not the vintage variables) are listed. Although all of the age effects are significant, the brand effects are not. This suggests that brand has little effect on the reservation e^* in Proposition 2. In the case when $\sigma_u = 0$, the only source of variation in e^* is due to preferences. The empirical results suggest that owners sort themselves so that the variation in e^* is small.

The baseline hazard in equation (5) is specified to have nodes in tenure at each year up to five years and a node in age at 2. Figure 3 shows the survival curves for cars disaggregated by age at purchase. One sees that the survivor curve is much steeper at first for new cars than used cars. This is in contrast to the Kaplan-Meier curves in Figure 2. Most of the difference in the survivor curves is captured by the difference in the baseline hazard in the first year of ownership. Part of the difference is due to the disaggregation mechanism; Figure 2 disaggregated by "new" or "used," while Figure 3 disaggregates by the age of the car when first purchased. Figure 2 indicates that 24% of new cars are sold by the end of 2 years. Thus, a significant proportion of the cars with $a \leq 2$ are used cars. However, the steepness of the $a \leq 2$ curve in Figure 3 cannot be explained by that alone. Rather, the correct accounting for "initial spells" captured in Figure 3 explains much of the change.

Next, we estimate a model with unobserved heterogeneity but with α restricted to be one. This corresponds to the classic lemons problem where the unobserved quality is determined at manufacture. The estimate of σ_u is 0.00014, and the other parameter estimates change very little. This suggests that, if there is unobserved heterogeneity, it can not be car-specific. It is therefore unlikely that unobserved quality caused by the manufacturer is a significant problem.

Finally, we estimate a model with both σ_u and α unrestricted. This allows for the unobserved quality to be created by the first owner. Only the brand and age coefficients for this specification are shown in Table 4. The estimate of σ_u is quite large relative to the variation in the effect of brand dummies, and the estimate of α is a large negative number. This suggests that there is variation in new car owners. Automobile aficionados want a plum of a car; they take care of their cars because $\partial U/\partial e$ is large but sell them quickly because

$\partial U/\partial t$ is also large and negative. Meanwhile, car pococurantes may buy new cars to avoid maintenance costs. They do not invest much in the car because $\partial U/\partial e$ is not large, and they sell them slowly because $\partial U/\partial t\mid_{a=0}$ is also not large in absolute value. However, once they sell the car, it is likely to be a lemon because of lack of maintenance and then becomes a hot potato. The magnitude of this effect in the data is reflected in Figure 4. The four survivor curves correspond to a new Ford. The two curves with "$u > 0$" correspond to a car with $u = \sigma_u$, and the two curves with "$u < 0$" correspond to a car with $u = -\sigma_u$. One can see that a new car with a $u > 0$ is turned over quickly (lowest curve) and then held for a long time by subsequent owners (second highest curve), while a car with a $u < 0$ is held for a long time (highest curve) and then turned over quickly by subsequent owners (second lowest curve).[16] The negative relationship exhibited is consistent with Proposition 5 in that variation in the unobservable heterogeneity u in the hazard rate in equation (5) is caused by variation in the utility function of initial owners of cars. Thus, a vehicle's unobservable quality not only depends on past levels as suggested by theory, but it also varies in a systematic way across first and subsequent owners of the same vehicle.

Once we allow for an unrestricted α, the brand effects become more significant and negative. However, this is relative to the brand effect for the excluded brand, Hyundai. More importantly, there is more variation in the brand effects. Also, the tenure effects become smaller but have a pattern similar to the earlier models. One can test whether the change in brand coefficients between the two specifications is significant using the Wald statistic[17]

$$\widehat{\theta}' D' \left[DCD'\right]^{-1} D\widehat{\theta} \sim \chi^2_m$$

where m is the number of brand estimates (22),

$$\widehat{\theta} = \begin{pmatrix} \widehat{\theta}_R \\ \widehat{\theta}_U \end{pmatrix},$$

$\widehat{\theta}_R$ is the restricted set of brand estimates in Table 3, $\widehat{\theta}_U$ is the unrestricted set of brand estimates in Table 4,

$$D = (I_m \mid 0 \mid -I_m \mid 0)$$

differences the brand estimates from the restricted specification from the brand

[16] These results are conditional on the age at which the vehicle is purchased. Therefore, the fact that an owner of car with $u < 0$ sells her vehicle quickly after the initial owner held onto it for a long time is not because the car is near the end of its working life.

[17] This test statistic follows from taking a Taylor series approximation of the score statistic for each specification, and writing

$$\sqrt{n}\left(\widehat{\theta}_j - \theta_j\right) = -A_j^{-1}\left(\frac{1}{\sqrt{n}}\sum_i \frac{\partial \log L_{ji}}{\partial \theta}\right), \quad j = R, U.$$

estimates from the unrestricted specification and nullifies the other estimates,

$$
\begin{aligned}
C &= plim\ n\left(\widehat{\theta}-\theta\right)\left(\widehat{\theta}-\theta\right)' \\
&= \begin{pmatrix} A_R^{-1}B_{RR}A_R^{-1} & A_R^{-1}B_{RU}A_U^{-1} \\ A_U^{-1}B_{UR}A_R^{-1} & A_U^{-1}B_{UU}A_U^{-1} \end{pmatrix},
\end{aligned}
$$

$$
\begin{aligned}
A_j^{-1} &= plim\ \left[\frac{1}{n}\sum_i \frac{\partial^2 \log L_{ji}}{\partial\theta_j \partial\theta_j'}\right], \quad j=R,U, \text{ and} \\
B_{jk} &= plim\ \left[\frac{1}{n}\sum_i \frac{\partial \log L_{ji}}{\partial\theta_j}\frac{\partial \log L_{ki}}{\partial\theta_k'}\right], \quad j,k=R,U.
\end{aligned}
$$

The χ_m^2 statistic is equal to 17.25 (not significant) when we do not condition on the other estimates, and it is equal to 53.89 (significant at the 1% level) when we do condition on them.

The reduction in the brand estimates with the introduction of unobserved heterogeneity (UH) can be explained in terms of a modified UH duration dependence bias argument. First consider a single brand with UH that is ignored (as in the restricted specification). Then a large realization of u will cause short ownership spells, and a small one will cause long ownership spells. When UH is ignored, then all of the spells are given equal weight because they are treated as independent. When UH is not ignored, multiple short spells within a car are not given as much weight because the likelihood function takes into account that they were all caused by the same large value of u. Thus, the constant is biased upwards when UH is ignored. Brand estimates are constants specific to a brand.

In our example, we observe the brand estimates declining relative to the base brand, Hyundai, when we add UH. This may mean that UH is more important for other brands than for Hyundai, causing declines in the brand estimates relative to Hyundai. In fact, Hyundai was offering strong and long warranties during this period, and this may have caused the UH associated with Hyundai to be small relative to other brands. This argument raises the issue whether there was heteroskedasticity in the UH across brands. We can test this hypothesis by constructing a Lagrange Multiplier (LM) statistic that considers the generalization that σ varies by brand under the alternative and is constant under the null. The LM test statistic is distributed χ_{22}^2 and is equal to 62.48 (significant at the 1% level). Thus, there is some evidence for heteroskedasticity across brands.

7.2 Alternative Interpretations of Results

In this section, we examine four alternative interpretations of the correlation of unobservable heterogeneity across owners of the same car. First, the estimated correlation of spell lengths is negative because we observe cars for only a fixed length of time. Second, the spells are correlated only because of the structure imposed on α_j. Third, the unobservable heterogeneity terms pick up vehicle characteristics that are observed by buyers - brand name and location of manufacturing plant - but have not been controlled for in the model. Finally, we address concerns over lack of data on leasing and the condition of the car's exterior (e.g., dents and scratches) and how this affects our ability to identify adverse selection in the used car market.

The first alternative interpretion is that our estimated negative correlation is due to observing cars for only a fixed period of time. Therefore, a car with a long initial spell must have shorter subsequent spells. However, we control for the age of the car when it was purchased by each sequential owner, and we control for right censoring. The negative correlation occurs even after controlling for age at time of purchase. The consistency of the parameter estimates follows from the standard theory of maximum likelihood estimation. However, just to be sure, we performed a Monte Carlo study using the parameter estimates from Table 4 with the exception of setting $\sigma = \alpha = 0$. If the procedure is biased, then the mean and median of the parameter estimates of α should be significant and negative. In fact, the mean (median) estimate of α was 0.029 (-0.034) and statistically insignificant ($t = -0.58$), and the mean (median) estimate of σ was 0.052 (0.062).[18]

The second alternative interpretation of our results is that the spells among owners after the first are correlated only because of the structure imposed on α_j in equation (4). One may be concerned about how robust our results are to a more general specification. The existence of initial spells, however, makes it computationally too difficult to allow for a more general specification for α_j. Alternatively, we can check robustness by simulating residuals in the sense of Gourieroux et al. (1987) and measure their correlation once grouped by order of ownership. This would allow us to determine if there is a more general relationship between ownership tenure spells than the one we assumed. The simulated correlation matrix is

$$\begin{array}{cccccccc}
1.00 & & & & & & & \\
-0.77 & 1.00 & & & & & & \\
-0.49 & 0.00 & 1.00 & & & & & \\
-0.47 & 0.05 & 0.34 & 1.00 & & & & \\
-0.46 & 0.08 & 0.29 & 0.35 & 1.00 & & & \\
-0.33 & 0.12 & 0.13 & 0.15 & 0.11 & 1.00 & & \\
-0.56 & 0.44 & 0.12 & 0.24 & 0.07 & 0.57 & 1.00 & \\
-0.47 & 0.35 & 0.01 & 0.20 & 0.13 & 0.45 & -0.01 & 1.00
\end{array} . \quad (9)$$

Note that all of the correlations of the first spells with other spells are negative and almost all of the other correlations are positive.[19] This indicates that using

[18] The t-statististic for the estimate of σ does not have a t-distribution because H_0 is at the boundary of the parameter space.

[19] One should be aware that the first spell is generally an initial spell and therefore may include more than one actual spell.

a more general correlation structure would provide similar results, and thus our results are not sensitive to the simplifying assumption for α.

The third alternative interpretation of our results is that a significant part of the unobserved heterogeneity in the model is car characteristics observed in the market but not used in the analysis. The most obvious example of such a characteristic is variation in brand beyond the 23 brands used in the paper. For example, there are six different sub-brands within Buick[20] and seventeen within "European Luxury." Consider representing the generalized residual (see Gourieroux et al. 1987) for the kth car of sub-brand j of brand i as

$$\begin{aligned} r_{ijk} &= \varsigma_{ij} + \eta_{ijk}, \\ \eta_{ijk} &\sim iidN\left(0, \sigma_\eta^2\right). \end{aligned} \tag{10}$$

The null hypothesis of interest is

$$H_0 : \varsigma_{ij} = \varsigma_{ij'} \left(= \zeta_i\right) \ \forall j, j', i \ \text{ vs } H_A : \exists i, j, j' : \varsigma_{ij} \neq \varsigma_{ij'}.$$

Given the normality assumption in equation (10), we can construct a Lagrange multiplier test statistic for each brand i of the form[21]

$$= \frac{1}{\widehat{\sigma}_\eta^2} \sum_{j=1}^{J_i} K_{ij} \left(\frac{1}{K_{ij}} \sum_{k=1}^{K_{ij}} \left(r_{ijk} - \widehat{\zeta}_i \right) \right)^2 \sim \chi^2_{J_i - 1} \tag{11}$$

where J_i is the number of sub-brands within brand i, K_{ij} is the number of observed cars of sub-brand j within brand i,

$$\widehat{\zeta}_i = \frac{\sum_{j=1}^{J_i} \sum_{k=1}^{K_{ij}} r_{ijk}}{\sum_{j=1}^{J_i} K_{ij}},$$

[20] These are Century, Le Sabre, Riviera, Regal, Roadmaster, and Skyhawk.

[21] Under H_0, the log likelihood function is

$$L = -\sum_{i=1}^{I} \sum_{j=1}^{J_i} \left[K_{ij} \log \sigma_\eta + \sum_{k=1}^{K_{ij}} \frac{\left(r_{ijk} - \zeta_i\right)^2}{2\sigma_\eta^2} \right].$$

The score statistic is

$$\left.\frac{\partial L}{\partial \varsigma_i}\right|_{\varsigma_{ij} = \zeta_i} = \frac{-1}{\sigma_\eta^2} \begin{pmatrix} \sum_{k=1}^{K_{i1}} \left(r_{i1k} - \zeta_i\right) \\ \sum_{k=1}^{K_{i2}} \left(r_{i2k} - \zeta_i\right) \\ \vdots \\ \sum_{k=1}^{K_{iJ_i}} \left(r_{iJ_ik} - \zeta_i\right) \end{pmatrix},$$

and its covariance matrix is

$$\frac{1}{\sigma_\eta^2} \begin{pmatrix} K_{i1} & 0 & \cdots & 0 \\ 0 & K_{i2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & K_{iJ_i} \end{pmatrix}.$$

The rest is matrix multiplication.

and

$$\widehat{\sigma}_\eta^2 = \frac{\sum_{i=1}^{I} \sum_{j=1}^{J_i} \sum_{k=1}^{K_{ij}} \left(r_{ijk} - \widehat{\zeta}_i\right)^2}{\sum_{i=1}^{I} \sum_{j=1}^{J_i} K_{ij}}.$$

These can be added up over (independent) brands. Note that the $\chi^2_{J_i-1}$ distribution in equation (11) followed from the assumption of normality in equation (10). However, even if the normality assumption is incorrect,

$$\frac{1}{\sqrt{K_{ij}}} \sum_{k=1}^{K_{ij}} \left(r_{ijk} - \widehat{\zeta}_i\right) \sim iidN\left(0, \sigma_\eta^2\right)$$

under H_0 and much weaker distributional assumptions from a central limit theorem. Thus, even if the normality assumption in equation (10)) is incorrect, asymptotically (as $K_{ij} \to \infty\ \forall i, j$)[22], the distribution in equation (11) is correct. Furthermore, above we provided some evidence of heteroskedasticity across brands. If we change the distributional assumption in equation (10) to

$$\eta_{ijk} \sim indN\left(0, \sigma_i^2\right),$$

then equation (11) changes only by replacing $\widehat{\sigma}_\eta^2$ with

$$\widehat{\sigma}_i^2 = \frac{\sum_{j=1}^{J_i} \sum_{k=1}^{K_{ij}} \left(r_{ijk} - \widehat{\zeta}_i\right)^2}{\sum_{j=1}^{J_i} K_{ij}}.$$

When we assume homoskedastic errors, the χ^2 statistics are significant for only three brands (Dodge, Pontiac, and Saturn), and the sum over brands is $\chi^2_{116} = 134.14$ which is not significant. When we allow for heteroskedastic errors, the χ^2 statistics are significant for only two brands (Dodge and Saturn), and the sum over brands is $\chi^2_{116} = 129.75$ which is also not significant. Thus there is no evidence that sub-brand heterogeneity is an important omitted source of heterogeneity.

Another possibility is that some other car characteristic such as manufacturing plant location (e.g., domestic or foreign) is a significant omitted source of heterogeneity. There are two reasons why the plant location may affect a car's quality and consequently the rate at which it is sold. First, vehicles produced in Germany by BMW, Mercedes-Benz, and Volkswagen have been found to be of higher quality than the same model cars produced in North American plants.[23] This also was found to be true for Japanese-produced Hondas, Mitsubishis, and Toyotas. The differential could be due to the fact that "different [car] models [are] produced in North America and these designs may be more difficult

[22] The range of K_{ij} is 2 (for two different sub-brands of Lexus and one sub-brand of Mitsubishi) to 168 for Honda Accord, and the average is 25.4.

[23] Each year J.D. Power & Associates surveys new car owners about their ownership experience and constructs a car quality rating based on the number of car problems reported per 100 vehicles (J.D Power & Associates, 2003).

to build," as suggested by an executive at J.D. Power & Associates that constructed the quality index (Porretto, 2003, pg. A3). Second, the car's quality may influence the firm's decision to produce it domestically or abroad. For instance, a manufacturer may choose to produce only its low-end vehicles abroad. Buyers in this market segment are more price-sensitive, more willing to trade off some quality for a lower price. If the firm transfers the production of the low-end vehicles abroad where labor is cheaper, this will allow the firm to price the cars lower without worrying about losing sales because the cars no longer possess the "domestic" car quality on which the manufacturer's reputation is based. This production strategy is particularly evident in explaining which cars Volkswagen and BMW choose to continue producing in Germany versus abroad.

Consider an alternative specification for residuals,

$$\begin{aligned} r_{ijmk} &= \varsigma_{ijm} + \eta_{ijmk} \\ \eta_{ijmk} &\sim iidN\left(0, \sigma_{\eta}^{2}\right) \end{aligned} \tag{12}$$

where now i indexes brand, j indexes whether the car was produced in a domestic plant or foreign plant, m indexes the year of production, and k indexes the particular car. The "domestic/foreign" dummy variable equals 1 iff the car is domestically produced (i.e., produced in the same country as the manufacturer's headquarters) and zero otherwise. For example, American-built Hondas are considered foreign production for this study. Define ξ^1 as a constant that does not vary over i, j, or m, ξ_i^3 as a constant that varies only over i, u_j^1 as a constant that varies only over j, u_{jm}^2 as a constant that varies only over j and m, and u_{ij}^3 as a constant that varies only over i and j. Consider the following hypotheses:

1. $H_0 : \varsigma_{ijm} = \xi^1$ vs. $H_A : \varsigma_{ijm} = u_j^1$ [No domestic/foreign effect vs. effect exists];
2. $H_0 : \varsigma_{ijm} = \xi^1$ vs. $H_A : \varsigma_{ijm} = u_{jm}^2$ [No domestic/foreign effect vs. brand*production year*domestic/foreign effect exists]; and
3. $H_0 : \varsigma_{ijm} = \xi_i^3$ vs. $H_A : \varsigma_{ijm} = u_{ij}^3$ [Brand effect vs. brand*domestic/foreign effect exists].

The first hypothesis tests the null that there is no "domestic/foreign" effect against the alternative that there is. The second tests the same null against the more general alternative allowing for an interaction between year and "domestic/foreign." The third tests the null that there is a brand effect against the alternative that the brand effect interacts with "domestic/foreign." Given the normality assumption in equation (12), Lagrange Multiplier statistics for the three hypotheses are respectively[24]

$$\frac{1}{\widehat{\sigma}_\eta^2}\sum_{j=0}^{1}\frac{\sum_{i=1}^{I}\sum_{m=1}^{M}\sum_{k=1}^{K_{ijm}}\left(r_{ijmk}-\widehat{a}^1\right)^2}{\sum_{i=1}^{I}\sum_{m=1}^{M}K_{ijm}} \sim \chi_1^2; \qquad (13)$$

$$\frac{1}{\widehat{\sigma}_\eta^2}\sum_{j=0}^{1}\sum_{m=1}^{M}\frac{\sum_{i=1}^{I}\left(r_{ijmk}-\widehat{a}^1\right)^2}{\sum_{i=1}^{I}K_{ijm}} \sim \chi_{2M-1}^2; \text{ and}$$

$$\frac{1}{\widehat{\sigma}_\eta^2}\sum_{i=1}^{I}\sum_{j=0}^{1}\frac{\sum_{m=1}^{M}\sum_{k=1}^{K_{ijm}}\left(r_{ijmk}-\widehat{a}_i^3\right)^2}{\sum_{m=1}^{M}K_{ijm}} \sim \chi_I^2$$

where $\widehat{a}^1$, $\widehat{a}_i^3$, and $\widehat{\sigma}_\eta^2$ are the restricted MLEs of a^1, a_i^3 (defined in the hypotheses), and σ_η^2 (defined in equation (12)).[25] The χ^2-statistics are reported in the first column of Table 5. The first two hypotheses are not rejected, but the third is. However, once we decompose the third into its χ_1^2 components, we see, among other things, that many of the largest χ_1^2 contributions to the test statistic are associated with brands that are produced only domestically or locally (not both). This suggests that possibly the existence of heteroskedastic residuals biases the test statistic.

We can generalize equation (12) to

$$\begin{aligned} r_{ijmk} &= \varsigma_{ijm} + \eta_{ijmk} \\ \eta_{ijmk} &\sim indN\left(0, \sigma_{ijm}^2\right), \end{aligned}$$

and consider the same hypotheses. To conserve degrees of freedom and avoid small sample problems, we assume σ_{ijm}^2 varies with respect to only one of i, j, or m. Consider the case when $\sigma_{ijm}^2 = \sigma_i^2$. Then the Lagrange Multiplier statistics are respectively

$$\sum_{j=0}^{1}\frac{\sum_{i=1}^{I}\frac{1}{\widehat{\sigma}_i^4}\sum_{m=1}^{M}\sum_{k=1}^{K_{ijm}}\left(r_{ijmk}-\widehat{a}^1\right)^2}{\sum_{i=1}^{I}\frac{1}{\widehat{\sigma}_i^2}\sum_{m=1}^{M}K_{ijm}} \sim \chi_1^2;$$

$$\sum_{j=0}^{1}\sum_{m=1}^{M}\frac{\sum_{i=1}^{I}\frac{1}{\widehat{\sigma}_i^4}\sum_{k=1}^{K_{ijm}}\left(r_{ijmk}-\widehat{a}^1\right)^2}{\sum_{i=1}^{I}\frac{1}{\widehat{\sigma}_i^2}K_{ijm}} \sim \chi_{2M-1}^2; \text{ and}$$

$$\sum_{i=1}^{I}\frac{1}{\widehat{\sigma}_i^2}\sum_{j=0}^{1}\frac{\sum_{m=1}^{M}\sum_{k=1}^{K_{ijm}}\left(r_{ijmk}-\widehat{a}_i^3\right)^2}{\sum_{m=1}^{M}K_{ijm}} \sim \chi_I^2$$

where $\widehat{\sigma}_i^2$ is the restricted MLE of σ_i^2.[26] The χ^2-statistics are reported in the

[24] For the third hypothesis, degrees of freedom are less than indicated in equation (13) because some cells are empty. For example, no Saturns were produced in foreign plants.

[25] $\widehat{a}^1$, $\widehat{a}_i^3$, and $\widehat{\sigma}_\eta^2$ are all first and second moments of the generalized residuals aggregated appropriately.

[26] For the models with homoskedastic residuals, the restricted MLEs could be easily evaluated. For the models with heteroskedastic residuals, the log likelihood function must be numerically maximized because the varying variance parameters affect the weighting of data in computation of first moments.

second column of Table 5.[27] Now neither reported hypothesis is rejected. If, instead, we specify the heteroskedasticity as $\sigma_{ijm}^2 = \sigma_j^2$, then the Lagrange Multiplier test statistics are respectively

$$\sum_{j=0}^{1} \frac{\sum_{i=1}^{I}\sum_{m=1}^{M}\sum_{k=1}^{K_{ijm}}\left(r_{ijmk}-\widehat{a}^1\right)^2}{\widehat{\sigma}_j^2\sum_{i=1}^{I}\sum_{m=1}^{M}K_{ijm}} \sim \chi_1^2;$$

$$\sum_{j=0}^{1}\frac{1}{\widehat{\sigma}_j^2}\sum_{m=1}^{M}\frac{\sum_{i=1}^{I}\sum_{k=1}^{K_{ijm}}\left(r_{ijmk}-\widehat{a}^1\right)^2}{\sum_{i=1}^{I}K_{ijm}} \sim \chi_{2M-1}^2; \text{ and}$$

$$\sum_{j=0}^{1}\frac{1}{\widehat{\sigma}_j^2}\sum_{i=1}^{I}\frac{\sum_{m=1}^{M}\sum_{k=1}^{K_{ijm}}\left(r_{ijmk}-\widehat{a}_i^3\right)^2}{\sum_{m=1}^{M}K_{ijm}} \sim \chi_I^2.$$

The χ^2-statistics are reported in the third column of Table 5. Again, no hypotheses are rejected. Thus, it appears that the brand-"domestic/foreign" interactions were heteroskedastic interactions rather than interactions affecting the first moments. For the specification where $\sigma_{ijm}^2 = \sigma_j^2$, the restricted MLEs of σ_j are 0.208 for domestic cars and 1.596 for foreign cars.

The final alternative interpretation of our results is that we lack data on leasing and vehicle condition to be able to identify adverse selection in the used car market. One might consider an alternative story in which there is no significant asymmetric information, and the unobserved heterogeneity observed in the duration data reflects characteristics of individual cars observable to the agents in the market but not to us. A car characteristic such as a paint scratch, a dented fender, or many miles on the odometer[28] might decrease the value of the car to a car aficionado but not be important to a potential used car owner. Thus, such a car might be sold early by the initial owner and then held by the first used owner as if it had no scratches or dents. Such a story would lead to a negative correlation between the first and second ownership spell. However, it is not so obvious why it would lead to the high correlations observed among subsequent spells in (9). The high correlations in (9) suggest that the car characteristic must affect the tenure decisions of all future owners. Asymmetric information is the most likely cause of such a correlation matrix because it can surprise owner after owner of a given car. However, another possibility is that e is observed by all, and those used-car owners with utility functions such that they prefer to buy low-e cars also prefer to sell them relatively quickly. Such a story is empirically indistinguishable from our asymmetric information story, but both imply a similar model of new car owner behavior.

We do not observe leasing arrangements in our data, and leasing may have a significant effect on turnover. Hendel and Lizzeri (2002) argue that leasing options segment the new car market with aficionados becoming lessees. They sell quickly and sell above-average quality cars. This is consistent with our empirical results and our story.

[27] No test statistic is reported for the second hypothesis because many of the cells are too small.

[28] Recall that, while we do not observe the odometer reading, Porter and Sattler (1999) do.

Akerlof's model is one of adverse selection: the private information concerns something that is exogenously determined. If an agent's private information concerns something that the agent chooses, the model is one of moral hazard. For example, if initial car quality were fully under the manufacturer's control and this quality were private information to the manufacture, the model would be one of moral hazard. The model in this paper is a hybrid, since the initial value of e is exogenous, but the owner has control over the subsequent values of e. A model in which the evolution of e is purely stochastic and unaffected by the owner's decisions would be a pure adverse selection model. The fact that lemons are not hot potatoes shows that the data are inconsistent with any simple lemons model. Whether it is an adverse selection or a moral hazard version is really a red herring.

8 Density of Lemons Among Cars on the Market

Given our estimates in Table 4, we can compute the density of $u_j = \alpha_j u$ among new cars and among cars participating in the used car market. By assumption, $u \sim iidN\left(0, \sigma^2\right)$, and the estimates of α and σ are -0.77 and 2.39 respectively. Suppose $0 < \Delta \leq t$, and let

$$f(u) = \frac{\sum_j \int \int h_{jv}^{\Delta}(t, a \mid u) \psi(t, a, j, u) \, dtda}{\sum_j \int \int \int h_{jv}^{\Delta}(t, a \mid u') \psi(t, a, j, u') \, dtdadu'}$$

be the density of u among cars on the used car market where $\psi(t, a, j, u)$ is the joint density of (t, a, u, j) and

$$h_{jv}^{\Delta}(t, a \mid u) = \frac{S_{jv}(t - \Delta, a \mid u) - S_{jv}(t, a \mid u)}{S_{jv}(t - \Delta, a \mid u)}.$$

Note that $\lim_{\Delta \to 0} h_{jv}^{\Delta}(t, a \mid u) = h_{jv}(t, a \mid u)$. Consider the function $n(t, a, j, u)$ that solves the system of equations,[29]

$$n(t, a, j, u) = \frac{S_{jv}(t, a \mid u)}{S_{jv}(t - \Delta, a \mid u)} n(t - \Delta, a, j, u); \quad t \geq \Delta; \qquad (14)$$

[29] Note that equation (14) converges to

$$\frac{\partial \log n(t, a, j, u)}{\partial t} = -h_{jv}(t, a \mid u),$$

equation (15) converges to

$$n(0, a, j+1, u) = \int_0^a \int h_{jv}\left(t, a' \mid u\right) n\left(t, a', j, u\right) dtda',$$

and equation (17) converges to

$$\phi(t, a, j, u) = \frac{n(t, a, j, u)}{\sum_{j'} \int \int \int n(t', a', j', u') \, dt'da'du'}$$

as $\Delta \to 0$.

$$n(0,a,j+1,u)=\sum_{a'<a}\sum_{t=0}^{a-a'} h_{jv}^{\Delta}(t,a'\mid u)\, n(t,a',j,u)\,; \tag{15}$$

$$n(0,0,1,u)=\frac{1}{\sigma}\phi\left(\frac{u}{\sigma}\right) \tag{16}$$

where $\phi(\cdot)$ is the standard normal density function. Then the density $\psi(t,a,u)$ satisfies

$$\psi(t,a,j,u)=\frac{n(t,a,j,u)}{\sum_{t'}\sum_{a'}\sum_{j'}\sum_{u'} n(t',a',j',u')}. \tag{17}$$

Note that the analysis can be done independently across brands.

The limits of equations (14), (15), (16), and (17) are too difficult to solve analytically. But we can approximate the solution by iteratively solving the system of equations (14), (15), (16), and (17) for some finite value of Δ. First we solve equation (16). Then, we solve equations (14) and (15) for those values of t and a such that $t+a=1$. At any point in the algorithm where we have solved for those values of t and a such that $t+a\leq k-1$, we solve equations (14) and (15) for those values of t and a such that $t+a=k$. Once we have done this for all $k\leq\tau$, we can solve equation (17).

Figure 5 shows the initial density of u, $N\left(0,\sigma^2\right)$. The other curves to the left are the density of u among cars being sold on the used car market, each curve corresponding to a different brand. All of the brand curves have a similar shape, and all are significantly to the left of the initial density of u. These correspond to cars where the initial owner held them for an unusually long time and subsequent owners are selling them quickly. Thus, to a significant degree, the used car market is dominated by hot potatoes.

Figure 6 displays the density of u among cars selling in the used car market and how it changes with the age of the car. It is clear that, once a car is 25 months old, it is likely to have a moderately negative value of u if being sold in the used car market. It is only for new cars (<25 months) that a significant proportion of the cars being sold have a more representative distribution of u.

There are two different ways of interpreting the downward shift in the distribution of the u terms as depicted in Figure 6. The different interpretations depend on what it means for a car to have a low u value. One possibility is that is that low-u cars tend to be those whose quality e has fallen in a way that is unobservable to buyers (by initial owners who held onto the car for a long time), so that these cars are lemons and hence become hot potatoes with high values of αu (because $\widehat{\alpha}<0$). Another possibility is that these low-u cars tend to be those whose attributes have changed in a way that is observable to buyers, so that the cars become an attractive choice for buyers (like college students) who plan to keep their cars for only a short time.

9 Conclusion

Our model allows us to distinguish among different types of adverse selection effects by observing the type of unobserved heterogeneity across owners of the same car. Our empirical results strongly suggest that there is a lemons effect because there is significant unobserved heterogeneity. However, they also suggest that the lemons effect is caused by the first owner rather than the manufacturer. Had the manufacturer created the lemon, the unobserved heterogeneity would be positively correlated over all owners of a given car. Instead we observe a negative correlation between the unobserved heterogeneity term for the first owner and the unobserved heterogeneity term for subsequent owners.

Our brand-effect estimates are small and generally insignificant. This is in contrast to the empirical results in Hendel and Lizzeri (1999) which relies on variation across brands in frequency of sales. However, our evidence in favor of heteroskedasticity of unobserved heterogeneity across brands suggests that the frequency of lemons problems varies across brands. This might be caused by the manufacturer, or it might be caused by correlation among new owners between brand preference and tendency to create lemons.

A recent development has affected the information available to buyers of used cars. In 1996, CARFAX, a corporation that had been collecting information on vehicle histories for car dealers, introduced CARFAX.com, a web-based service that allows consumers access to the same information. This service allows anyone to check the sales history of vehicles for a modest fee, currently below $20.

Suppose that everyone were to take advantage of this.[30] Now price can depend on sales history, and the previous equilibrium in which cars with different sales histories appeared identical to buyers, and hence sold at identical prices, will no longer remain an equilibrium.

Intuitively, what will happen now is that, if a used car is resold quickly, it is clear it is a lemon and will be sold at very low price. Thus, people buying cars realize that they won't be able to get rid of cars relatively quickly and inexpensively, and thus are going to be more cautious about avoiding bad cars than in the original equilibrium. In the original model, we posited an equilibrium in which some consumers who bought new cars kept them longer than average and maintained them more poorly than average. If the buyer observes the length of the initial ownership spell, the longer initial spell length would indicate poorer maintenance and hence reduces the sales price of such cars more than previously. Consumers now would have to decide whether to sell their cars earlier so as to conceal that they maintain their cars poorly, or to keep them until they are scrapped.

[30] Because the information is costly, and we model the used car market as competitive, there is a Grossman-Stiglitz (1980) problem. Since the price will now reflect relevant vehicle history information a consumer can free ride and skip the CARFAX fee if everyone else (or enough other people) are checking histories. But if everyone free rides, the information will not be reflected in price. The actual equilibrium has some but not all consumers paying the fee. But, as the fee goes to zero, the outcome should converge to the competitive one.

10 Appendix: Proof of Propositions

Proof. (Proposition 1) Because the single period net utility flow $U - C$ is bounded above and $\beta^N W(S_{t+N})$ converges to zero as N tends to infinity, the function $V(S_t)$ is the solution to the problem of choosing the number of periods N (possibly infinite) to keep the car before selling and the sequence $e_{t+1}, e_{t+2}, ..., e_{t+N}$, so as to maximize the objective

$$\sum_{s=0}^{N-1} \beta^s [U(b, v, a, s+t, e_{t+s}) - C(e_{t+s}, e_{t+s+1})] + \beta^N W(S_{t+N}).$$

For fixed finite N, we are guaranteed a solution to this optimization problem because the objective function is continuous, and we can restrict our choices to a compact set, as we now show. Let $\overline{U}$ be an upper bound of U. Fix e and suppose $e_{t+s} \le e$. Because C is convex and increasing in its second argument, there is a finite e^u such that $e_{t+s+1} > e^u$ implies that $C(e_{t+s}, e_{t+s+1}) > \overline{U}/(1-\beta)$, which means that choosing $e_{t+s+1} > e^u$ would necessarily make the entire objective function negative. Thus we need only consider values of e_{t+s+1} that are less than or equal to e^u. Proceeding in this way, we obtain, period by period, a sequence of upper bounds on the level of e_{t+s}. This means that we need search only over the product of a sequence of compact sets, which is itself a compact set. For each N, denote the maximum value attained in the optimization problem by $V_N(e_t)$. Thought of as a sequence in N, this either has a maximum at some finite N^0 or, if not, then the optimal selection is to choose N to be infinite The continuity of each V_N follows from Berge's Maximum Theorem (see, for example, Aliprantis and Border, 1999, p. 539). Let $V_N^{'} = \max\{V_1, ..., V_N\}$. Then the $V_N^{'}$ are continuous and, because $U - C$ is bounded above and $\beta^N W(S_{t+N})$ converges to zero, the $V_N^{'}$ converge uniformly to V. Hence V is continuous.
For monotonicity, suppose that $e_t' > e_t$, and let $S_t = (b, v, a, t, e_t)$ and $S_t' = (b, v, a, t, e_t')$. Starting at e_t, suppose that the optimal plan is to hold for N periods (possibly infinitely many) and to choose $e_{t+s} = \overline{e}_{t+s}$, for $s = 1, ..., N$, which yields $V(S_t)$. Then, starting at $e_t' > e_t$, the plan of also holding for N periods and choosing $e_{t+s} = \max\{e_t', \overline{e}_{t+s}\}$ for $s = 1, ..., N$ gives a strictly larger discounted payoff than $V(S_t)$ because, for each period s, the utility term is at least as high and the cost term no higher than in the optimal plan starting at e_t. The strict inequality follows because the utility is strictly greater and the cost strictly lower in the current period $s = 0$. It follows that $V(S_t') > V(S_t)$. Our assumptions are sufficient to ensure that, for a given level of e_t and given number of holding periods N, the optimal sequence $e_{t+1}, e_{t+2}, ..., e_{t+N}$ is unique. This follows because the objective function above is strictly concave. ■

Proof. (Proposition 2) As observed above, W is independent of $e_{t.}$. But, by Proposition 1, V is strictly increasing in e_t. The result follows. ■

Proof. (Proposition 3) Suppose that $e_t' > e_t$ and the corresponding optimal choices are $e_{t+1}' = G(e_t')$ and $e_{t+1} = G(e_t)$. If, given e_t, the owner sells the car, then $e_{t+1}=e_t$, and $e'_{t+1}>e_{t+1}$ follows immediately because $e'_{t+1}\geq e'_t$. Suppose instead that, given e_t, the owner keeps the car. Because $e'_t>e_t\geq e^*$ the owner, given e'_t, also will keep the car. If $e_{t+1}<e'_t\leq e'_{t+1}$, the result follows immediately. If not, since e'_{t+1} is optimal given e'_t, we have

$$\begin{aligned}&U\left(b,v,a,t,e'_t\right)-C(e'_t,e'_{t+1})+\beta V\left(S'_{t+1}\right)\geq\\&U\left(b,v,a,t,e'_t\right)-C(e'_t,e_{t+1})+\beta V\left(S_{t+1}\right)\end{aligned}$$

or equivalently

$$C(e'_t,e_{t+1})-C(e'_t,e'_{t+1})\geq\beta V\left(S_{t+1}\right)-\beta V\left(S'_{t+1}\right).$$

Similarly,

$$C(e_t,e'_{t+1})-C(e_t,e_{t+1})\geq\beta V\left(S'_{t+1}\right)-\beta V\left(S_{t+1}\right).$$

Adding yields

$$C(e'_t,e'_{t+1})-C(e'_t,e_{t+1})\leq C(e_t,e'_{t+1})-C(e_t,e_{t+1})$$

or

$$\int_{e_{t+1}}^{e'_{t+1}}C_2(e'_t,e_2)de_2\leqslant\int_{e_{t+1}}^{e'_{t+1}}C_2(e_t,e_2)de_2$$

or equivalently

$$\int_{e_{t+1}}^{e'_{t+1}}\int_{e_t}^{e'_t}C_{12}(e_1,e_2)de_1de_2\leqslant 0.$$

Since $C_{12}\leq 0$, it follows that $e'_{t+1}\geq e_{t+1}$.
Corollary 4 follows by induction and the fact that G is increasing. ■

Proof. (Proposition 5) If $e''_t>e'_t$ then, in every period $t+s$ that car e'_t is unsold we have $e'_{t+s}=G^s(e'_t)>e^*$. By the corollary above $e''_{t+s}=G^s(e''_t)\geq G^s(e'_t)>e^*$, and so car e''_t will remain unsold too. The result follows immediately. ■

Proof. (Proposition 6) Consider any equilibrium price vector which determines the price of each observationally distinct class of cars in each period. The condition guarantees that, for each such class, in each period, there will always be a positive measure of individuals who wish to buy at that price or a positive measure of individuals who wish to sell at that price. Thus there must be a positive volume of trade in equilibrium. ■

11 References

References

[1] Akerlof, George (1970). "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics*. 84(3): 488-500.

[2] Aliprantis, C. D. and K. Border (1999). *Infinite Dimensional Analysis: A Hitchhiker's Guide*. 2nd Edition.

[3] Bond, Eric (1982). "A Direct Test of the 'Lemons' Model: The Market for Used Pickup Trucks." *American Economic Review*. 72(4): 836-840.

[4] Bond, Eric (1984). "Test of the Lemons Model: Reply." *American Economic Review*. 74(4): 801-804.

[5] Chezum, Brian and Bradley Wimmer (1997). "Roses or Lemons: Adverse Selection in the Market for Thoroughbred Yearlings." *Review of Economics and Statistics*. 79(3): 521-526.

[6] Chezum, Brian and Bradley Wimmer (2000). "Evidence of Adverse Selection from Thoroughbred Wagering." *Southern Economic Journal*. 66(3): 700-714.

[7] Desai, Preyas and Devavrat Purohit (1998). "Leasing and Selling: Optimal Marketing Strategies for a Durable Good Firm." *Management Sciences*. 44(11): S19-S34.

[8] Emons, Winand and George Sheldon (2003). "The Market for Used Cars: A New Test of the Lemons Model." Unpublished manuscript.

[9] Engers, Maxim, Monica Hartmann, and Steven Stern (2004). "Automobile Maintenance Costs, Used Cars, and Adverse Selection." Unpublished manuscript.

[10] Engers, Maxim, Monica Hartmann, and Steven Stern (2007). "Annual Miles Drive Used Car Prices." *Journal of Applied Econometrics*. forthcoming.

[11] Finkelstein, Amy and Kathleen McGarry (2006). "Multiple Dimensions of Private Information: Evidence from the Long Term Care Insurance." *American Economic Review*. 96(4): 938-958.

[12] Finkelstein, Amy and James Poterba (2004). "Adverse Selection in Insurance Markets: Policy Holder Evidence from the UK Annuity Market." *Journal of Political Economy*. 112(1): 183-208.

[13] Genesove, David (1993). "Adverse Selection in the Wholesale Used Car Market." *Journal of Political Economy*. 101(4): 644-665.

[14] Gilligan, Thomas (2004). "Lemons and Leases in the Used Business Aircraft Market." *Journal of Political Economy.* 112(5): 1157-1180.

[15] Gourieroux, Christian, Alain Monfort, Eric Renault, and Alain Trognon (1987). "Simulated Residuals." *Journal of Econometrics.* 34(1/2): 201-252.

[16] Grossman, Sanford and Joseph Stiglitz (1980). "On the Impossibility of Informationally Efficient Markets." *American Economic Review.* 70(3): 393-408.

[17] Heckman, James and Burton Singer (1984). " A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data." *Econometrica.* 52(2): 271-320.

[18] Heckman, James and James Walker (1990). "The Relationship Between Wages and Income and the Timing and Spacing of Births: Evidence from the Swedish Longitudinal Data." *Econometrica.* 58(6): 1411-1441.

[19] Hendel, Igal and James Cardon (2001). "Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey." *RAND Journal of Economics.* 32(3): 408-27.

[20] Hendel, Igal and Allessandro Lizzeri (1999). "Adverse Selection in Durable Goods Markets." *American Economic Review.* 89(5): 1097-1114.

[21] Hendel, Igal and Allessandro Lizzeri (2002). "The Role of Leasing Under Adverse Selection." *Journal of Political Economy.* 110(1): 113-143.

[22] Hendel, Igal, Allessandro Lizzeri, and Marciano Siniscalchi (2005). "Efficient Sorting in a Dynamic Adverse Selection Model." *Review of Economic Studies.* 72(2): 467-497.

[23] Johnson, Justin and Michael Waldman (2003). "Leasing, Lemons, and Buy-Backs." *RAND Journal of Economics.* 34(2): 247-265.

[24] Lacko, James (1986). "Product Quality and Information in the Used Car Market." Staff Report. Washington: Federal Trade Commission, Bureau of Economics.

[25] Meyer, Bruce (1990). "Unemployment Insurance and Unemployment Spells." *Econometrica.* 58(4): 757-782.

[26] Porretto, John (2003). "Toyota Still Top Quality." *Deseret News.* May 7.

[27] Porter, Robert and Peter Sattler (1999). "Patterns of Trade in the Market for Used Durables: Theory and Evidence." NBER Working Paper 7149.

[28] Power, J.D. (2003). *2003 Initial Quality Survey.* West Lake Village, California.

[29] Rosenman, Robert and Wesley Wilson (1991). "Quality Differentials and Prices: Are Cherries Lemons?" *Journal of Industrial Economics.* 39(6): 649-658.

12 Tables and Figures

Table 1 Spell Moments			
Spell Type	# Obs	Mean	Std. Dev.
Initial	3197	1608.0	1161.9
Middle	2379	726.0	670.0
End	3008	1164.1	865.2
Total	8584		

Note: Spells are measured in days.

Table 2
Number of Observations by Brand

Variable	# Obs		Variable	# Obs
Buick	145		Oldsmobile	154
Chevrolet	346		Plymouth	69
Chrysler	82		Pontiac	183
Dodge	126		Saturn	51
Ford	533		Subaru	60
Geo	29		Toyota	356
Honda	405		Volkswagen	77
Hyundai	62		Volvo	47
Mazda	154		American Luxury	133
Mercury	138		European Luxury	114
Mitsubishi	60		Japanese Luxury	34
Nissan	185		Total	3543

Table 3
Parameter Estimates for Nonparametric Model
Without Unobserved Heterogeneity

Variable	Estimate		Variable	Estimate
Buick	-0.397** (0.157)		Oldsmobile	-0.121 (0.146)
Chevrolet	-0.162 (0.133)		Plymouth	-0.035 (0.159)
Chrysler	-0.270 (0.173)		Pontiac	-0.056 (0.136)
Dodge	-0.244* (0.145)		Saturn	-0.338** (0.169)
Ford	-0.092 (0.131)		Subaru	-0.110 (0.167)
Geo	0.017 (0.189)		Toyota	-0.229* (0.133)
Honda	-0.078 (0.131)		Volkswagen	0.045 (0.149)
Hyundai	0.000		Volvo	-0.337 (0.231)
Mazda	-0.212 (0.143)		American Luxury	-0.121 (0.155)
Mercury	0.014 (0.144)		European Luxury	-0.223 (0.158)
Mitsubishi	-0.136 (0.160)		Japanese Luxury	-0.409 (0.275)
Nissan	-0.075 (0.138)			

Parameter Estimates (Continued)				
Age:$\leq$ 2 years			Age:$\geq$ 2 years	
Tenure	Estimate		Tenure	Estimate
$\leq$ 1 year	-5.675** (0.128)		$\leq$ 1 year	-8.068** (0.178)
1 - 2 years	-7.086** (0.119)		1 - 2 years	-7.078** (0.140)
2 - 3 years	-7.692** (0.167)		2 - 3 years	-8.157** (0.193)
3 - 4 years	-7.304** (0.129)		3 - 4 years	-7.153** (0.150)
4 - 5 years	-7.992** (0.177)		4 - 5 years	-7.675** (0.162)
$\geq$ 5 years	-7.340** (0.136)		$\geq$ 5 years	-7.132** (0.150)

Notes:

1. Numbers in parentheses are standard errors.
2. Double starred items are significant at the 5% level.
3. The “Hyundai” dummy variable is the base and is restricted to zero.
4. The log likelihood value is -26640.1.

Table 4
Parameter Estimates for Nonparametric Model
With Unobserved Heterogeneity and Unrestricted α

Variable	Estimate		Variable	Estimate
Buick	-0.658** (0.257)		Oldsmobile	-0.360 (0.251)
Chevrolet	-0.374* (0.226)		Plymouth	-0.101 (0.294)
Chrysler	-0.522* (0.282)		Pontiac	-0.195 (0.238)
Dodge	-0.326 (0.258)		Saturn	-0.379 (0.294)
Ford	-0.230 (0.222)		Subaru	-0.283 (0.296)
Geo	0.082 (0.409)		Toyota	-0.385* (0.225)
Honda	-0.201 (0.224)		Volkswagen	-0.097 (0.316)
Hyundai	0.000		Volvo	-0.751** (0.416)
Mazda	-0.341 (0.247)		American Luxury	-0.339 (0.263)
Mercury	-0.123 (0.254)		European Luxury	-0.419* (0.267)
Mitsubishi	-0.239 (0.274)		Japanese Luxury	-0.794** (0.395)
Nissan	-0.216 (0.240)			
Sigma(σ)	2.390** (0.147)		Alpha (α)	-0.774** (0.047)

Parameter Estimates (Continued)				
Age:≤ 2 years			Age:≥ 2 years	
Tenure	Estimate		Tenure	Estimate
≤ 1 year	-5.575** (0.233)		≤ 1 year	-6.340** (0.256)
1 - 2 years	-7.093** (0.206)		1 - 2 years	-6.211** (0.223)
2 - 3 years	-6.265** (0.243)		2 - 3 years	-6.349** (0.275)
3 - 4 years	-6.984** (0.210)		3 - 4 years	-5.988** (0.235)
4 - 5 years	-6.388** (0.254)		4 - 5 years	-5.629** (0.256)
≥ 5 years	-6.740** (0.216)		≥ 5 years	-5.569** (0.247)

Notes:

1. Numbers in parentheses are standard errors.
2. Double starred items are significant at the 5% level.
3. The "Hyundai" dummy variable is the base and is restricted to zero.
4. It is assumed that the unobserved heterogeneity is distributed normally. "Sigma" is the standard deviation.
5. The log likelihood value is -22069.6.
6. The likelihood ratio test statistic for $H_0 : \alpha = 0$ against $H_A : \alpha \neq 0$ is 1141.2. The 1% critical value for a χ_1^2 is 6.63.

Table 5

χ^2 Test Statistics for Domestic/Foreign Plant Effects

Hypothesis	df	Homoskedastic Errors	Heteroskedastic Errors $\sigma^2_{ijm} = \sigma^2_i$	$\sigma^2_{ijm} = \sigma^2_j$
H_0 : No domestic/foreign effect vs. H_A : Effect exists	1	2.06	0.95	1.22
H_0 : No domestic/foreign effect. vs. H_A : Production year*dom./foreign effect	45	20.95		1.78
H_0 : Brand effect. vs. H_A : Brand*domestic/foreign effect	19	187.04**	5.53	10.70

1. Notes:
2. In σ^2_{ijm}, i, j, and m represent the ith brand, jth production location - domestic or foreign, and mth model year, respectively.
3. For the third hypothesis, degrees of freedom are less than indicated in equation (13) because some cells are empty.
4. Double-starred items are significant at the 5% level.

Figure 1: Examples of Sample Ownership Spells

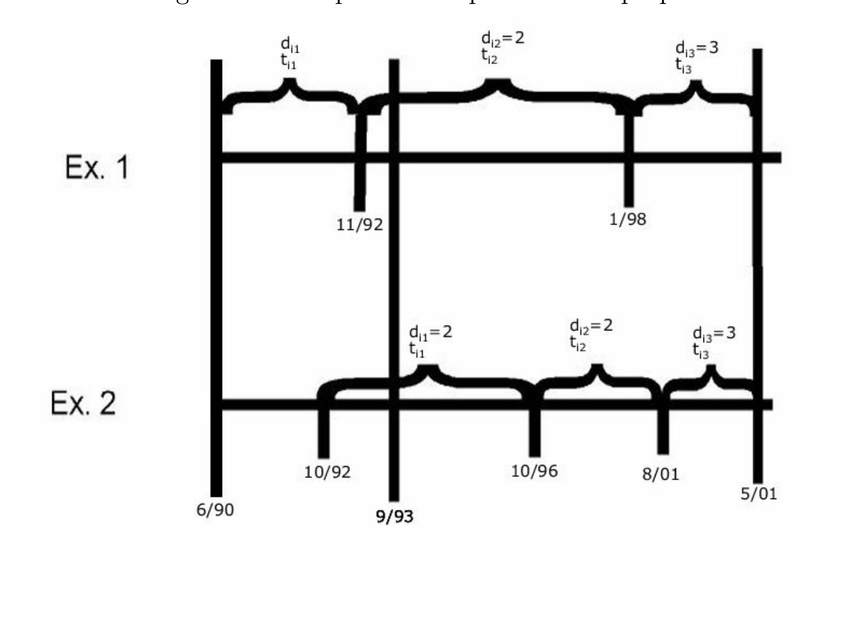

Figure 1

Figure 2: Kaplan-Meier Survival Curves

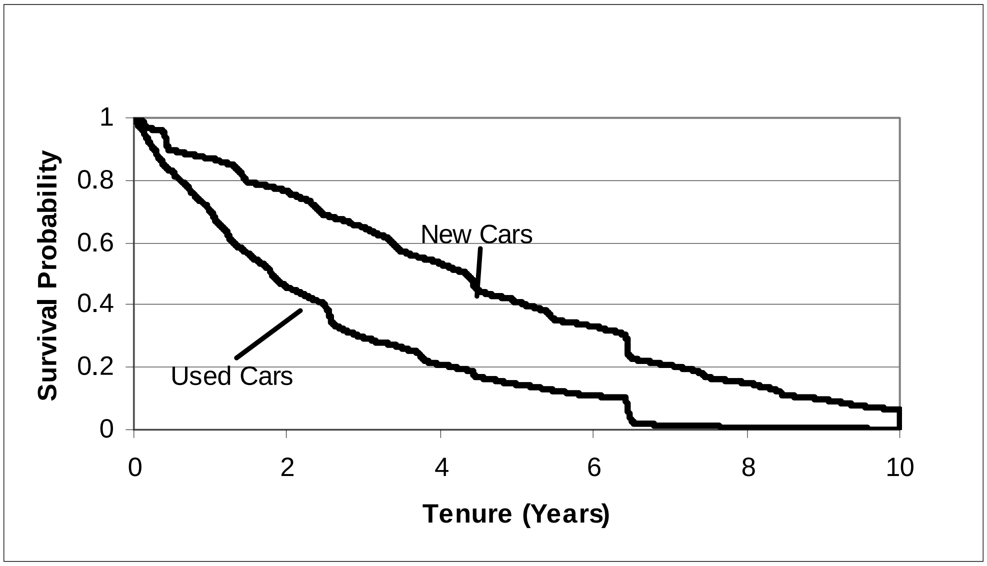

Figure 2

Figure 3: Survival Curves without Unobserved Heterogeneity

Figure 3

Figure 4: Survival Curves with Correlated Unobserved Heterogeneity

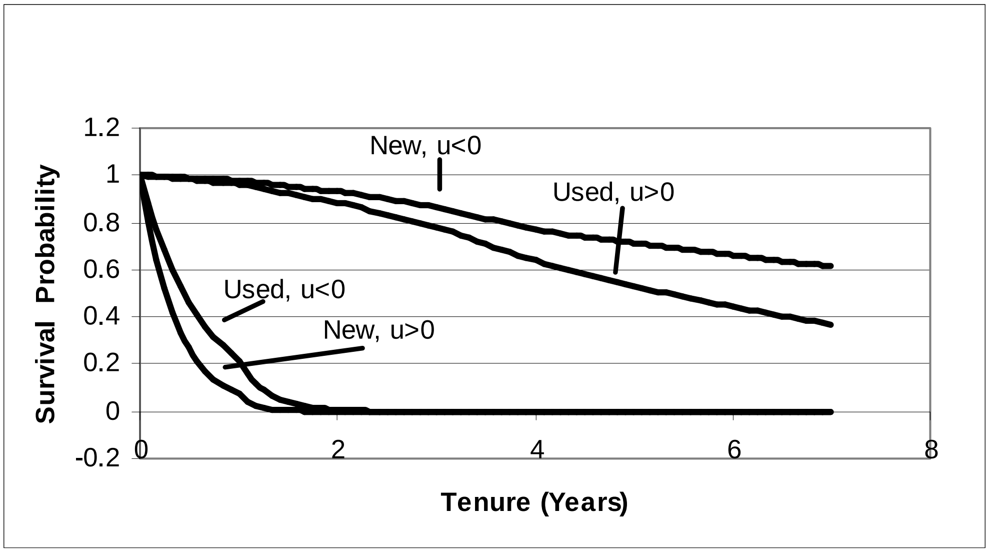

Figure 4

Figure 5: Variation in Density of Lemons Across Brands

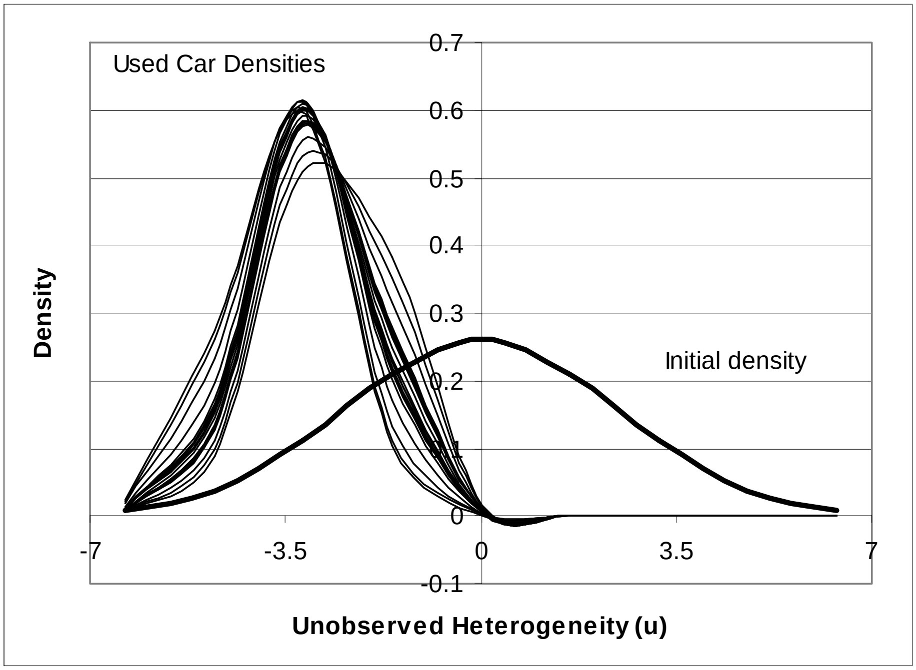

Figure 5

Figure 6: Density of Unobserved Heterogeneity by Age

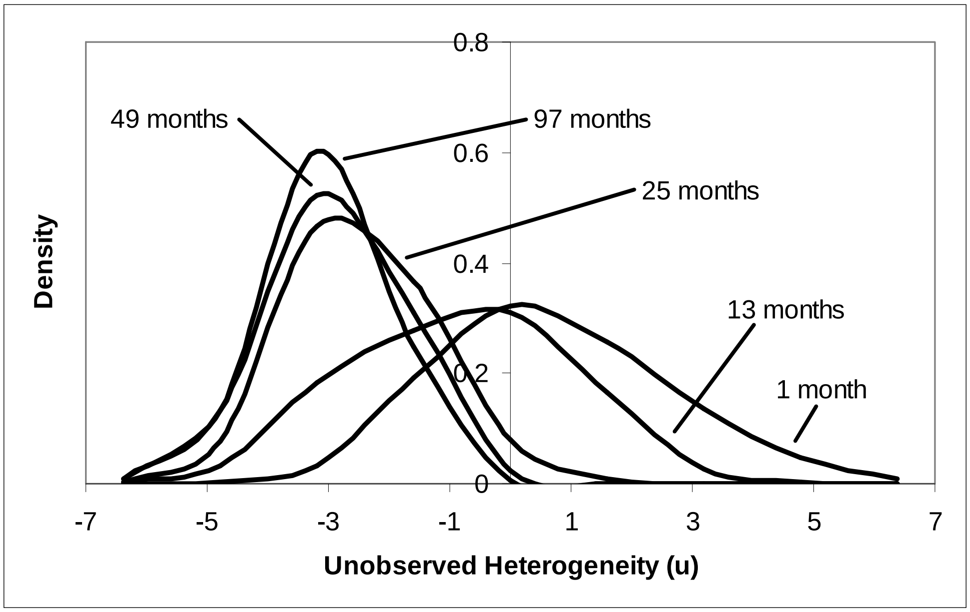

Figure 6