Estimating the Effects of Adverse Selection in Used Car Markets

Henry Schneider*
Yale University

November 14, 2005

Abstract

In this paper, I address the long-standing question of whether adverse selection prevents used cars from reaching owners who value them most highly. In doing so, I confront the challenge of identifying the effects of adverse selection separately from the effects of efficient sorting of vehicles based on their conditions. This latter process would usually occur simultaneously to adverse selection and also affects the distribution of vehicles that trade. Using the prediction in Hendel and Lizzeri (1999), that adverse selection and efficient sorting both increase the rate of price depreciation, I propose to use their joint effect as an upper bound on the effect of adverse selection. My estimate of this joint effect, based on proprietary data on one million dealer used car sales and trade-ins, is close to zero, a result that indicates that adverse selection is unimportant. Using Consumer Expenditure Survey data, I provide additional support for this conclusion by showing that vehicles that were recently purchased from a dealership received approximately the same number of repairs as comparable continuously-held vehicles. I conclude with a discussion of the role that sellers' concerns for their reputations may play in limiting information-based inefficiencies.

*I am grateful to Steven Berry and Fiona Scott Morton for discussions and guidance. I also thank Romana Primus at Whaling City Ford and Paul Trembley at East Rock Auto Repair for background information on the used car and car repair markets. Philip Schmidt-Dengler and various seminar participants provided useful comments.

1 Introduction

In this paper, I address the long-standing question of whether adverse selection prevents used cars from reaching owners who value them most highly. This information asymmetry arises from buyers' inabilities to evaluate vehicle condition and is predicted to introduce potentially serious inefficiencies into the used car market. Figures 1 and 2, which show car owners' expenditures on vehicle repairs and is based on Consumer Expenditure Survey data, provide a picture of the distribution over these conditions. Figure 1 shows that average annual repair expenditures are modest: Even seven-year-old American-made vehicle had less than \$400 in repairs. Figure 2, however, shows that the most troublesome vehicles had expenditures in excess of \$1100. These amounts represent significant fractions of the values of these vehicles, and provide their owners with incentives to sell them to uninformed buyers. Despite the risks of being on the receiving ends of these transactions, buyers still purchased 41 million used cars in 2003.[1] The extent of these activities reveals that adverse selection fails to prevent at least some gains from trade from being realized. The possibility remains distinct, however, that adverse selection still introduces inefficiencies.

Hendel and Lizzeri (1999) analyze two mechanisms that may affect the distribution of vehicles that trade. One mechanism is efficient sorting, a process whereby vehicles that have deteriorated since they were originally purchased trade to new owners who value their current conditions more highly. This process is driven by the gains from trade that arise from differences in tastes for vehicle conditions.[2] The second mechanism is adverse selection, an information asymmetry that threatens the efficiency of this sorting process. Since sellers receive a price that is consistent with average unobserved condition, owners of cars with good conditions would receive lower prices, and owners of cars with poor conditions would receive higher prices, than they would receive under complete information. The incentives that arise from these price disparities may influence these transaction decisions, and affect the trade volumes, prices, and qualities of the vehicles that end up trading.

The task of identifying the effects of adverse selection separately from the effects of efficient

[1] This number is substantially larger than the 17 million new cars that were sold in 2003. These trade volumes are reported in http://www.edmunds.com/advice/buying/articles/45310/article.html (November 5, 2005).

[2] Efficient sorting may also be driven by differences in other aspects of vehicle quality besides vehicle condition. I focus on vehicle condition because it is this aspect of quality that also leads to adverse selection. Hendel and Lizzeri (1999) take the same approach.

sorting has proven challenging because the underlying causes of these mechanisms - the amount of information asymmetry and the rate of quality deterioration - both stem from reliability.[3] As such, adverse selection and efficient sorting would usually occur simultaneously and affect the same market outcomes. This challenge is exacerbated by the absence of precise measures for either information asymmetry or quality deterioration.

I address this identification challenge, which has prevented clear conclusions from emerging in the empirical literature, by introducing a simple approach for estimating the importance of adverse selection. Using the prediction in Hendel and Lizzeri, that adverse selection and efficient sorting both increase the rate of price depreciation, I propose to use their joint effect as an upper bound on the effect of adverse selection. Finding a small joint effect would reveal that adverse selection has little impact on this market. Finding a large effect is consistent with adverse selection being important, but does not actually identify its effect separately from the effect of efficient sorting.

The transfer of vehicle repair costs from manufacturers to owners upon warranty expiration provides an opportunity to estimate this joint effect. Warranties effectively equalize the condition of cars by covering the costs of repairing defects. As such, they remove the incentives that arise from adverse selection for owners to sell defective vehicles to uninformed buyers. For the same reason, warranties also limit the gains from trade that lead to efficient sorting on vehicle condition. Both mechanisms come into play when warranties expire, and if either is important, the rate of price depreciation should increase.[4]

I examine proprietary data on over one million used car transactions that occurred at franchised dealerships in California between 1996 to 2002, and find that the rate of price depreciation remained unchanged upon warranty expiration. I find a similar pattern for vehicles that were traded in to dealers, as measured by the actual cash values dealers booked into inventory. These results indicate that adverse selection is unimportant in these segments of the market.

I also analyze vehicle repair histories contained in the Consumer Expenditure Survey (CES) and find additional evidence that adverse selection has little effect on the distribution of vehicles that trade: Traded vehicles received about the same number of repairs in the 3 months following

[3]For example, unreliable cars will have more information asymmetry and also faster quality deterioration.

[4]The existence of this warranty effect relies on the assumption that customers do not view warrantied repairs as indicators of future costly repairs. In my analysis of robustness, I support this assumption by showing that current repairs are uncorrelated with future repairs once vehicle make and model are accounted for.

transactions relative to comparable continuously-held vehicles.

California regulates its used car markets minimally. The state's lemon laws apply only to cars under warranty, and most vehicles are sold "as is" unless extended warranties are purchased separately. The absence of regulatory structures suggests that market-based mechanisms alone have mitigated inefficiencies that would have resulted from asymmetric information. Two mechanisms of this kind are particularly salient. First, buyers are able to obtain assessment services such as independent vehicle appraisals from garages and background vehicle reports from Carfax. Second, buyers can use their ability to withhold repeat visits and referrals to discourage dealers from selling defective vehicles. Buyers are able to effectively employ these reputation-based incentives - so called because they provide incentives for dealers to pursue honest reputations - because they learn about the qualities of their purchases fairly quickly. I discuss further the role that reputation may play in limiting information-based efficiencies in the conclusion.

I also examine maintenance and repair histories of privately-traded vehicles and find that they received 22% fewer oil changes before transaction and 64% more repairs after transaction relative to comparable continuously-held vehicles. These findings indicate that privately-traded vehicles are inferior to comparable continuously-held vehicles. This result suggests that adverse selection is important in this market segment, but fails to rule out the possibility that efficient sorting is the cause instead. In the conclusion, I argue that the one-time nature of these transactions, which limits buyers' abilities to employ reputation-based incentives, allows this outcome to occur.

The rest of the paper is structured as follows. In section 2, I discuss the challenges to estimating adverse selection and how I propose to overcome them. Section 3 describes previous work on adverse selection in durable goods markets, and highlights places in the empirical literature where identification is unclear. Section 4 contains a description of the data. Section 5 describes the main results of this paper. In section 6, I analyze vehicle maintenance and repair histories from the CES, and find additional support for the results in section 5. I conclude in section 7 with a discussion of the role that reputation may play in limiting the harmful effects of adverse selection.

2 Identifying adverse selection effects

The most direct test of adverse selection would examine how trade volumes and the distributions over prices and unobserved conditions respond to changes in information asymmetry.[5] Researchers have been unable to effectively implement this method for two reasons. First, variation in information asymmetry, which is used to measure the effects of adverse selection, is usually highly correlated with variation in quality deterioration, which creates the gains from trade that drive efficient sorting. Thus, even when data that contain variation in information asymmetry are available, researchers may be unable to identify these effects separately from the effects of efficient sorting. Second, this difficulty is exacerbated by the absence of a precise measure for either quantity.

To illustrate this inherent identification problem, consider a stylized car whose quality is defined by the condition of its n parts. Each part has good condition with probability p and poor condition with probability $1-p$. Furthermore, the conditions of these parts are independent events, and are known only to the seller. The expectation and variance of quality in the population of vehicles, then, are np and $np(1-p)$. The high correlation of these two quantities, which are the underlying determinants of efficient sorting and adverse selection, generates this identification problem. In fact, in this case, as in many real-world cases, identification may rely on functional form assumptions alone.

Given these challenges to separately identifying the effects of adverse selection and efficient sorting, I propose to estimate them jointly. I find the outcomes from this approach to be highly informative about the effects of adverse selection.

I draw heavily on predictions from the model in Hendel and Lizzeri (1999) to justify this approach. The authors construct a dynamic model of a market in which a constant number of new cars goes on sale every period, and becomes used the following period. Cars deteriorate so that used cars have lower quality than new cars. Differences in tastes for quality lead individuals with high-valuations to prefer new cars, and individuals with low-valuations to prefer used cars.

[5] I will often use the term information asymmetry as a shorthand to represent the part of vehicle quality that is known to owners but would not be to buyers were these owners actually to bring these vehicles to market. Applying this term to all cars in a population instead of just cars that have been brought to market represents an abuse of terminology since actual buyers and sellers, between which information asymmetry would exist, may not actually be present. Hendel and Lizzeri (1999) and Gilligan (2004) refer to this quantity as quality uncertainty and reliability, though neither of these terms precisely reflects its true meaning either.

This heterogeneity in tastes leads to the efficient sorting of used cars that have deteriorated at different rates to new owners who value these conditions more highly.

The model in Hendel and Lizzeri predicts that this process of efficient sorting causes vehicles with faster quality deterioration to trade more often, and to have faster rates of price depreciation. The intuition for this result is that buyers of new cars, who have higher valuations for quality, will seek to replace cars that deteriorate quickly more often. Buyers will willingly purchase these vehicles because they find their mix of lower quality and affordable price appealing.

The authors also introduce into their model asymmetric information between buyers and sellers over vehicle quality. Their model predicts that vehicles with more information asymmetry have faster price depreciation and lower trade volumes. This outcome occurs because adverse selection is predicted to decrease the number of high quality vehicles in the distribution of vehicles that trades.

Since adverse selection and efficient sorting should both increase the rate of price depreciation, their combined effect represents an upper bound on the effect of adverse selection. Finding a small joint effect is most informative because it reveals that adverse selection is unimportant. Finding a large joint effect raises the possibility that adverse selection is present, but does not identify its effect separately from the effect of efficient sorting.[6]

As discussed in the introduction, warranty expiration represents a point in a vehicle's life when adverse selection becomes relevant. There are two reasons why the rate of quality deterioration may also increases at this point. The first is the inherent connection between the mean and variance of condition that is discussed above. The second is the likely decline in the frequency of repairs that would arise when price elastic owners suddenly bear these cost.[7] The simultaneity of these two processes shows that warranty expiration may be no better at separately identifying the effects of adverse selection than some previous approaches. However, I use this point in the vehicle's life as an opportunity to estimate their joint effect, which bounds the effect of adverse selection between this amount and zero.

[6]Since the model in Hendel and Lizzeri predict that efficient sorting and adverse selection affect the volume of trade in opposite directions, these outcomes are unhelpful for constructing bounds.

[7]The possibility exists that the rate of price depreciation would also increase upon warranty expiration because prices no longer incorporate the insurance benefit of paying for repairs that may be needed in the future. I limit my sample to vehicles between 2.5 and 3.5 years old with warranties that expire after 3.0 years. Since these vehicles have warranties that cover less than six months of future repairs, and these cars are young enough that few repairs will be needed anyway, this should be a minor factor.

3 Literature

Akerlof (1970), Wilson (1980), and Kim (1985) are founding papers in the theoretical literature on adverse selection. Porter and Sattler (1999) and Hendel and Lizzeri (1999, 2000) are more recent studies that make explicit dynamic considerations such as the interaction between the new and used car market. Hendel and Lizzeri (1999) is the most relevant to the estimation strategy in this paper, and was discussed in section 2.

The empirical literature on adverse selection in durable goods markets begins with Bond (1982, 1984). He attempts to test for differences in unobserved quality between traded and continuously-held trucks that have the same observed quality, and asserts that such differences would be attributable to adverse selection. Bond estimates the following regression model,

$$repairs_i = \beta_0 + \beta_1 age_i + \beta_2 mileage_i + \beta_3 traded_i + \epsilon_i$$

where $repairs_i$ is annual repair expenditures, age_i is truck age, $mileage_i$ is odometer reading, and $traded_i$ indicates whether the truck was traded in the previous year. If the estimate of β_3 is positive, then traded trucks require more repairs than continuously-held trucks.

A number of factors prevent this test from effectively measuring the effects of adverse selection. First, except for vehicle mileage and age, this specification fails to control for observed vehicle quality. According to the model in Hendel and Lizzeri, efficient sorting will cause traded vehicles to have systematically lower observed quality than continuously-held vehicles (even after controlling for age and mileage). Since observed quality will be captured by the error term, and will be correlated with the variable *traded*, its effect will be misattributed to adverse selection. Second, even if the unobserved qualities of traded and continuously-held vehicles were found to differ, efficient sorting, which is predicted to lower the quality (observed and unobserved together) of traded vehicles, could have been the cause. In order to attribute differences in unobserved quality to adverse selection, information asymmetry must be shown to be the cause. Still, Bond only finds a difference in repairs between traded and continuously-held trucks for trucks that are over 10 years old.

Genesove (1993) also tests for adverse selection by examining whether traded vehicles have lower unobserved quality than continuously-held vehicles. Specifically, Genesove examines whether car dealers sell their lowest unobserved quality cars at auction first by testing for differences in prices obtained by new car dealers (NCDs) compared to used car dealers (UCDs). Since NCDs sell a higher proportion of their vehicles at auction than UCDs, NCDs should receive a higher average price if both dealer types sell their lowest unobserved quality cars first.[8] Like Bond, Genesove has assumed that a finding of lower unobserved quality among traded cars compared to continuously-held cars reveals that adverse selection is present. As already discussed, however, efficient sorting could also cause the quality distribution (observed and unobserved together) of traded cars to differ from that of continuously-held cars - and hence average prices for UCDs to be lower than those for NCDs - regardless of adverse selection.

Gilligan (2004) tests the prediction in Hendel and Lizzeri (1999) that the trade volumes and the rates of price depreciation are directly related under complete information, but are inversely related under adverse selection, using data on used airplane sales. His measure of adverse selection is airworthiness directives, which are regulatory instructions issued to particular airplane models when safety concerns arise. Gilligan finds a direct relationship between the trade volumes and the rates of price depreciation for airplanes with few airworthiness directives, but an inverse relationship for airplanes with many airworthiness directives. These results would seem to validate Hendel and Lizzeri's model.

I would argue, however, that airworthiness directives may proxy for observed quality better than information asymmetry. First, these directives are as observable to buyers as they are to sellers. Second, even if airworthiness directives proxy for quality uncertainty - an assertion Gilligan supports by showing that the correlation between airworthiness directives and accidents is 0.85 - this uncertainty will likely be shared by both buyers and sellers. As such, airworthiness directives may have a larger effect on expected quality than on information asymmetry. For these reasons, this test measures the joint effect of adverse selection and efficient sorting. Still, Gilligan finds that the trade volumes and the rates of price depreciation are directly related for most airplane models, a result that suggests that efficient sorting usually dominates adverse selection.[9]

[8] This test relies on the implicit assumption that the distributions over unobserved conditions of vehicles in the NCDs' and UCDs' inventories are the same.

[9] Other empirical studies include Pratt and Hoffer (1984), Lacko (1986), Winard and George (2002), and Engers, Hartmann, and Stern (2003).

4 Data

An auto industry data company has provided proprietary data on over one million franchised dealer used car transactions that occurred in California between January 1996 and January 2002. This segment of the used car market, together with the independent dealer segment, accounts for about three-quarters of used trades. Transactions between private individuals comprises the remaining quarter.[10] Data on two sets of transactions are recorded: Dealer sales to buyers, and trade-ins from buyers back to dealers. For both sets of transactions, the data contain the actual transaction price, the actual cash values dealers booked into inventory, the financing terms, and various other attributes of the transaction. The data also contain a description of the vehicle, including its age, mileage, and trim level, but do not include assessments of vehicle condition, nor do they report the cost of any repairs that may have been conducted by the dealership.

I combine these records with car reliability data that are published in the April issues of Consumer Reports magazine. These data are derived from responses to The Consumer Union's annual survey of magazine subscribers, which draws in feedback on about 500,000 vehicles each year. For each vehicle model, these data record the frequency of serious problems in 14 trouble spots, such as engine, transmission, and brakes.[11] Consumer Reports magazine records reliability on a scale of 1 to 5, where 1 indicates that more 14.8% of respondents reported serious problems, 2 indicates that 9.3-14.8% reported serious problems, 3 signifies that 5-9.3% had problems, 4 shows that 2-5% had problems, and 5 indicates that less 2% had problems. In my calculations, I approximate these ranges with the values, 19%, 12.05%, 7.15%, 3.5%, and 1%. I use data from the surveys they conducted between 1995 and 2003, which provide reliability records on vehicles with model years between 1988 and 2003. In total, these data cover 26,239 model, model year, vehicle age combinations.

I also analyze repair and maintenance expenditure recorded in the Consumer Expenditure Survey for survey years 1995 through 2004. The survey is a five quarter rolling panel dataset that records expenditures for 9 major repair categories, such as engine, transmission, and cooling. I also examine consumer expenditures on the maintenance item oil changes.

[10] 73% of used trades recorded in the CES occurred through dealerships, 24% occurred through private channels, and 3% occurred through other channels.

[11] Respondents are instructed to define serious in terms of "cost, failure, compromised safety, and downtime."

5 Results

The primary empirical results of this paper are joint estimates of the effects of adverse selection and efficient sorting on price. Before reporting the results of more formal statistical tests, however, I present price depreciation curves for 1995 and 1997 Honda Civics and Accords, and 1995 and 1997 Ford Escorts and Taurus's. These models were among the most popular sold in the United States in those years. Figures 3 and 4 contain price curves for 2 sets of transactions: Used vehicle sales from franchised dealers to customers, and also vehicle trade-ins from customers back to these dealers. The dependent variable in these figures is the final negotiated transaction price expressed in 2003 dollars. For sales to customers, this price includes discounts from rebates, costs from any after market options that were purchased, and adjustments to account for any trade-in allowances. For customer trade-ins back to dealers, this price is the actual cash value the dealer booked into inventory. The independent variable is vehicle mileage at the time of transaction. These figures contain prices for cars that are 2.0 to 2.9 years old and have mileages between 26,000 and 46,000. These vehicles all have warranties that expire after 36,000 miles and 3 years have elapsed. Age is restricted to be less than 3 years to make clear how prices change upon warranty expiration at 36,000 miles. The two line segments superimposed on each set of transactions are linear fits of these transaction prices before and after the warranty has expired. Since these data contain relatively few trade-in transactions, the trade-in splines should be viewed more cautiously. The wiggly curve that is superimposed on these prices is a median cubic spline fit to the transaction prices.

From these figures, it is apparent that warranty expiration had little effect on the rate of price depreciation. For all 8 models, the linear fit to transaction prices changes remarkably little upon warranty expiration. Furthermore, there is no consistent shift in the intercept upon warranty expiration. These graphs suggest that neither adverse selection nor efficient sorting on vehicle condition has much effect on transaction price.

Franchised dealer-to-customer transactions

Next, I test formally whether the rate of price depreciation increases upon warranty expiration. I restrict the sample to vehicles with three-year, 36,000 mile manufacturer "bumper-to-bumper" warranties. 88% of models had these warranties, although drivetrain warranties that extend for additional miles and years often applied.[12] I examine cars that were 2.5 to 3.5 years old with mileages between 26,000 to 46,000 at the time of transaction. Except for the shifts in adverse selection and quality deterioration that form the basis of my experiment, the effect of age and mileage on price should be smooth around 3.0 years and 36,000 miles. I estimate the price depreciation curves by fitting transaction prices with model-specific intercepts and model-specific linear age and mileage slope terms. I also allow these age and mileage effects to shift upon warranty expiration in order to measure the joint effect of adverse selection and efficient sorting. The regression model is,

$$\begin{aligned} p_{mi} &= \sum_{m=1}^{M} \beta_{0m} I_{mi} + \sum_{m=1}^{M} \beta_{1m} age_{mi} + \sum_{m=1}^{M} \beta_{2m} miles_{mi} \\ &+ \sum_{m=1}^{M} \beta_{3m} age_{mi} nw_{mi} + \sum_{m=1}^{M} \beta_{4m} miles_{mi} nw_{mi} + \epsilon_{mi} \end{aligned} \tag{1}$$

where p_{mi} is price for transaction i of model and model year m expressed in 2003 dollars; I_{mi} indicates if vehicle i is of model and model year m; age_{mi} and $miles_{mi}$ are age and mileage at the time of transaction; and nw_{mi} indicates that the warranty has expired because the age or the mileage limit has been reached.

Since estimation of this model produces a large number of coefficient estimates, I report only summary outcomes. An F-test shows that the estimates of the model-specific fixed effects, β_0, are jointly statistically different than zero (p=0.00). Estimates of the model-specific vehicle age effects, β_1, are also jointly significant (p=0.00). However, the estimates of the effects of warranty expiration on the age effects, β_3, are generally small in magnitude and only marginally jointly significant (p=0.08). The estimates of β_2 and β_4, the model-specific mileage effect before and after warranty expiration are also generally small in magnitude and jointly insignificant (p=0.41 and p=0.15, respectively). This regression has an adjusted R-squared of 0.79, and is based on data from 14,772 vehicle transactions.[13]

Using coefficient estimates from equation 1, I calculate the average rate of price depreciation

[12] Bumper-to-bumper warranties cover all repairs that are not routine maintenance. Drivetrain warranties cover transmission and engine failures only. My outcomes remained unchanged when I restricted my sample to vehicles with drivetrain warranties that also expired after 36,000 miles and 3 years had elapsed.

[13] I also estimated a specification that included model-specific warranty expiration fixed effects to allow the intercept of the price curve to shift upon warranty expiration. These coefficient estimates were jointly insignificant (p=0.44).

for 1 year and 12,000 miles of driving using a weighted average of these rates across all vehicle models. I use for the weights the number of observations for that model. My estimator of the average rate of price depreciation, then, is,

$$\sum_{m=1}^{M} n_m(\widehat{\beta}_{1m} + \widehat{\beta}_{2m}12000)$$

where n_m is the number of transactions recorded for model m. My estimator of the average change in the rate of price depreciation upon warranty expiration is,

$$\sum_{m=1}^{M} n_m(\widehat{\beta}_{3m} + \widehat{\beta}_{4m}12000)$$

I use the nonparametric bootstrap method to calculate standard errors of the resulting estimates.

For vehicles under warranty, I find that prices depreciated an average of $1316 for one-year and 12,000 miles of use. The standard error of this estimate is $70. The average rate of price depreciation increased by only $7 upon warranty expiration. The standard error of this estimate is $16. This pattern holds for American as well as Asian-made vehicles. For American-made vehicles under warranty, the rate of price depreciation was $1829 with a standard error of $96. When the warranty expires, the rate of price depreciation increased by $77 with a standard error of $273. For Asian-made vehicle under warranty, the rate of price depreciation was $1098 with a standard error of $73. Upon warranty expiration, this rate increased by $15 with a standard error of $15.

These results show that the rate of price depreciation remained unchanged upon warranty expiration, a result that indicates that meaningful amounts of adverse selection and efficient sorting are absent in the franchised dealer segment of the used car market.

Buyer trade-ins to dealers

Trade-ins from buyers back to dealers represents another segment of the market in which adverse selection may occur. While dealers are undoubtedly more proficient at recognizing vehicle quality than the average car buyer - they buy and sell dozens of vehicles a week - they typically offer trade-in amounts before mechanics evaluate the conditions of the vehicles.[14] For price, I examine the actual cash value the dealer books into inventory. This amount often varies

[14] Paul Tremblay of East Rock Auto Repair and Romana Primus of Whaling City Ford provided information on this process.

from the payment the owner receives, and represents the dealer's appraisal of the vehicle's actual worth after it has been inspected but before any in-house repairs are made.

The regression model from the previous subsection is used here. The estimate for average price depreciation for one-year and 12,000 miles of use was \$1700 with a standard error of \$229. This rate decreased by \$7 with a standard error of \$63 when the warranty expired. The same patterns hold for American and Asian-made vehicles. American-made vehicles under warranty depreciated by \$1785 per year with a standard error of \$394. That rate decreased by \$37 with a standard error of \$81 upon warranty expiration. For Asian-made vehicles under warranty, prices depreciated by \$1408 per year with a standard error of \$435. When the warranty expired, the rate of price depreciation increased by \$58 with a standard deviation of \$92. These results show that the rate of price depreciation changed little upon warranty expiration, and indicates that meaningful amounts of adverse selection and efficient sorting are absent in the trade-in segment as well.

6 Robustness Analysis

Temporary versus permanent unobserved quality

My identification strategy assumes that adverse selection is absent when the manufacturer warranty is in effect, and is based on the fact that defects discovered during the warranty period are repaired for free. However, if buyers view current problems as indicators of future problems they may have to pay for, some adverse selection may occur during the warranty period. This would weaken the identification strategy of this paper.

For this assumption to hold, low quality must be a transitory characteristic such that once these defects are repaired, the probability that additional defects arise is no greater than that for other cars of that make, model, age, and mileage. This assumption would fail is low unobserved quality is a more permanent trait such that one failure increases the probability that more failures are likely.[15]

[15] Modern day manufacturing techniques have reduced the prevalence of defects due to the manufacturing process. Womack, Jones, and Roos (1991) describe how lean production techniques, pioneered by Toyota, have come to replace mass production, long associated with the big three US car manufacturers. They discuss how mass production is more likely to produce cars with difficult-to-correct inherent defects. Even as recently as 1987, the Framingham, Massachusetts General Motors plant had 135 defective parts for every 100 cars coming off the

I test whether low unobserved quality is a permanent or transitory trait by examining whether repairs are concentrated in a few lemons or are distributed more evenly across vehicles. If repairs are evenly distributed across cars, then current problems say little about the probability of future problems. If repairs are concentrated, this suggests that low unobserved quality is a more permanent trait.

Using CES data on the frequency of repairs for each car part conditional on make, model, age, and mileage, I predict the distribution over the number of repairs per car that would occur if repairs across parts are independent events. This hypothetical distribution assumes that one defect on a particular vehicle has no effect on the probability of another defect on that same vehicle, after controlling for the car's descriptive characteristics. I denote the hypothetical distribution over the number of repairs per vehicle under independence as $\hat{p}(n)_{\perp}$. Then I compare this to the empirical distribution over the number of repairs per vehicle, $p(n)$. If repairs are concentrated in a subset of lemons, the empirical distribution should have more cars with zero repairs and multiple repairs, and less cars with one repair, than under the distribution predicted under independence.

Table 1 shows that these two distributions are close: Problems across parts are close to independent events. There are only 4 percentage points more vehicles with 0 repairs, 0.1 percentage points more vehicles with 2 repairs, and 1.5 percentage points more vehicles with 3 or more repairs in the empirical distribution compared to the independent distribution. While a small percentage of cars seem to require more repairs than chance alone would predict, which raises the possibility that low unobserved condition may persist occasionally, at least some of these extra repairs are likely explained by heterogeneity in repair preferences on the part of customers and mechanics.

Maintenance and repair frequencies of traded vehicles

I provide further evidence that adverse selection is a minor factor in the used car market by showing that maintenance and repair expenditures are similar for traded and continuously-held

assembly line. By contrast, a comparable Toyota plant in 1987 with lean production techniques had 45 defective parts for every 100 cars coming off the assembly line. For the last two decades, however, most new cars sold in the US have been made with lean production techniques. This change has altered the nature of car reliability, shifting the cause of car failures away from the manufacturing process and towards wear and tear and owner maintenance.

vehicles with the same observed characteristics. These results show that the joint effect of adverse selection and efficient sorting on the distribution of the conditions of vehicles that trade is close to zero.

Using CES vehicle repair data, I estimate the effect of purchasing a vehicle in the previous 3 months from a dealership, and from an individual seller through private channels, on an indicator variable for whether repairs are conducted that month. I use a probit specification, and include as control variables, the log of the repair frequency for that vehicle model, model year, and age from Consumer Reports data, the difference between the vehicle's mileage and the average mileage for vehicles with that model, model-year, and age, and the interaction of these two control variables. I limit my sample to vehicles for which the manufacturer warranty has expired to focus on vehicles whose owners bear the costs of repairs. Table 2 contains coefficient estimates from this regression, expressed as marginal effects.

The average frequency of repairs in any given month is 2.7%. The estimate of the effect of having purchased the vehicle in the previous 3 months indicates that this frequency is only 0.16 percentage points higher in the three months following purchase from an independent or franchised dealership. As such, the frequency of repairs is only 6% larger following purchase, an increase that is statistically indistinguishable from zero.

I also estimate the same regression model, but for each repair category separately. These categories are transmission, engine, cooling, electrical, air conditioning, suspension, brakes, and exhaust. Figure 5 presents these coefficient estimates graphically. In the top panel, the darker bar is the monthly repair frequency for the indicated part for continuously-held cars, while the lighter bar is the estimated frequency for cars that were purchased from a dealership within the 3 previous months. The bottom panel contains the estimated differences in repair frequencies between these purchased and continuously-held vehicles.

Examining changes in repair probabilities for specific repair categories reveals that traded vehicles experience more transmission repairs. A dealer-purchased vehicles is 97% more likely to have transmission repairs in the three months after purchase relative to a comparable continuously-held vehicle. This effect is statistically distinguishable from zero (p=0.01). The effect of trade on the repair frequencies for the 7 other repair categories are indistinguishable from zero. The frequencies of repairs for dealer-purchased vehicles relative to comparable continuously-held vehicles are 1.33 for engine (p=0.19), 0.90 for cooling (p=0.72), 0.94 for electrical (p=0.80), 0.80

for air conditioning (p=0.60), 0.85 for suspension (p=0.38), 0.78 for brakes (p=0.15), and 1.24 for exhaust (p=0.51).

These results raise the possibility that adverse selection over the unobserved condition of transmissions may occur, but do not actually identify it separately from the effects of efficient sorting. Given the infrequency with which dealers actually advertise when vehicles have defects, however, adverse selection is likely the cause. The average frequency of transmission repairs increased from 0.3% to 0.6% per month for vehicles that were recently purchased. Although transmission repairs are typically the most expensive system to fix, these repairs likely occur too infrequently to affect market outcomes in a meaningful way: Even if each transmission repair costs $1000, a 0.3 percentage point increase in the probability of repairs only decreases car value by $30 in expectation.

Examining differences in routine maintenance before disposal provides a measure of how owners treat their vehicles before selling them. I find that vehicles that were sold privately had 22% fewer oil changes in the three months prior to sale (p=0.00), and 16% fewer oil changes in the three months prior to trade-in (p=0.08), relative to comparable continuously-held vehicles.[16] Even though this difference is modest - were oil changes to be conducted every 3 months for continuously-held cars, they would be done every 3.6 months for soon-to-be traded cars - it does suggest that vehicle condition may be neglected prior to disposal.[17]

Finally, I examine the repair frequencies of vehicles that were purchased through private channels within the prior three months. These transactions provide an important comparison to dealer traded vehicles. The one-time nature of these transactions limits the effectiveness of reputational factors, such as repeat visits and referrals, in reducing the effects of adverse

[16] Because the CES record vehicle mileage in only one of the five quarters the respondent is interviewed, I am unable to control for vehicle usage in the 3 months prior to sale, except to control for the vehicle's overall mileage.

[17] The frequency of repairs before transaction for traded vehicles is actually modestly higher for equivalent continuously-held vehicle. In the three months prior to sale, traded vehicles received 20% more repairs than equivalent continuously-held vehicles. This difference, however, is statistically indistinguishable from zero (p=0.42). In the three months prior to trading-in a vehicle, 33% fewer repairs were made compared to the equivalent continuously-held vehicle. Again, this difference is statistically indistinguishable from zero (p=0.23). Though these tests do not have sufficient power to detect moderate differences in repair frequencies before disposal, they rule out large differences. An alternative explanation for equivalent repair frequencies prior to disposal could be that traded cars had lower unobserved quality, and required more repairs, than continuously-held cars, but that additional necessary repairs were neglected. This, in fact, is what would occur with adverse selection.

selection. A positive estimate of the joint effect of adverse selection and efficient sorting for private trades, then, would suggest that when sellers do not receive incentives to pursue honest reputations, the effects of adverse selection become important.

I find that privately purchased vehicle had 64% more repairs in the 3 months following transaction relative to comparable continuously-held vehicles (p=0.00). The estimate of this effect is reported in table 2. In a separate regression with a similar specification, I find that cars received 22% fewer oil changes in the 3 months prior to sale (p=0.00). The higher concentration of low condition vehicles in the distribution of privately-traded vehicles raises the possibility that adverse selection is important to this market segment. However, this large joint effect may also be due to an efficient sorting outcome in which sellers of defective vehicles advertised them as such, and sold them to buyers who welcomed these fixer-upper opportunities.

Conclusions

In this paper, I circumvent the challenge of identifying the effect of adverse selection separately from the effect of efficient sorting of vehicles based on their conditions by proposing to estimate them jointly. My estimate of this joint effect, which is based on dealer used car trades to customers, and also customer trade-ins back to dealers, shows that neither adverse selection nor efficient sorting is important in this segment of the used car market. Using Consumer Expenditure Survey data, I provide additional support for this conclusion by showing that vehicles that were recently purchased from a dealership received about the same number of repairs as comparable continuously-held vehicles.

These findings are consistent with Gilligan's (2004) result that trade volumes and rates of price depreciation are directly related, and hence adverse selection is of secondary importance, for transactions of most models of aircraft. The results in Gilligan do suggest, however, that adverse selection affects market outcomes in some cases. Unobserved condition may assume a central role for these airplane transactions because airplane owners face more serious risks from poor condition than car owners do.

The absence of government regulatory structures in California's used car market suggests that market-based mechanisms alone may have mitigated inefficiencies that would have resulted from adverse selection. One such mechanism is vehicle assessment services. These include independent vehicle appraisals provided by auto repair shops, and background vehicle reports supplied by Carfax. Another mechanism is reputation-based incentives that arise from buyers' abilities to withhold repeat visits and referrals from dealers that sell defective vehicles. Buyers can effectively employ these reputation-based incentives - so called because they provide sellers with incentives to pursue good reputation - because they learn about vehicle condition shortly after purchase.

The results I have present for the used car market, and, to a lesser extent, the results presented in Gilligan, stand in contrast to my findings for the auto repair market. In Schneider (2005), I find that the deadweight losses that result from asymmetric information between mechanics and car owners may be large enough to represent half of all repair costs. Since mechanics usually address vehicles' most visible defects regardless of whether they were overtreated, or whether less visible defects were neglected, buyers are limited in their abilities to evaluate mechanics' types. This prevents buyers from effectively employing reputation-based incentives, such as repeat visits and referrals, to encourage good service.

In this paper, I have also presented evidence that is consistent with the existence of adverse selection in the private segment of the used car market. I find that vehicles that traded through private channels received less maintenance before transaction and more repairs after transaction relative to comparable continuously-held vehicles. The inability of buyers to employ reputation-based incentives in this market segment, because of the one-time nature of these transactions, may explain this result as well. The results in Bond (1982, 1984), that traded trucks have inferior condition relative to comparable continuously-held trucks only for trucks older than 10 years, suggests a similar pattern: A significantly higher proportion of older trucks that traded undoubtedly sold through private channels compared to younger trucks that traded. Hubbard (1998) studies the vehicles inspection market and also provides evidence that sellers take actions to acquire good reputations. Hubbard (2002) shows that vehicle owners are highly responsive to these actions, which reveals that clear incentives for sellers to pursue good reputations exist in this inspection market.

Taken as a whole, these differences in the extent to which information asymmetry introduces inefficiencies across these markets suggest that buyers' abilities to provide reputation-based incentives to sellers are important to limiting the harmful effects of information asymmetry.

References

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

[2] Bond, Eric (1982) "A Direct Test of the 'Lemons' Model: The Market for Used Pickup Trucks." American Economic Review, 72, pp. 836-40.

[3] Bond, Eric (1984) "Test of the Lemons Model: Reply." American Economic Review, 74, pp. 801-04.

[4] Emons, Winand and George Sheldon (2003) "The Market for Used Cars: A New Test of the Lemons Model." Manuscript, University of Bern, Switzerland.

[5] Engers, Maxim, Monica Hartmann, and Steven Stern (2003) "Are Lemons Really Hot Potatoes?" Working Paper, University of Virginia.

[6] Genesove, David (1993) "Adverse Selection in the Wholesale Used Car Market." Journal of Political Economy, 101, pp. 644-65.

[7] Gilligan, Thomas (2004) "Adverse Selection and Trade in Used Durable Goods: Evidence from the Market for Business Aircraft." Journal of Political Economy, 112, pp. 1157-80.

[8] Hendel, Igal and Alessandro Lizzeri (1999) "Adverse Selection in Durable Goods Markets." American Economic Review, 89, pp. 1097-1115.

[9] Hubbard, Thomas. (1998) "An Empirical Examination of Moral Hazard in the Vehicle Inspection Market." The Rand Journal of Economics, 29, pp. 406-26.

[10] Hubbard, Thomas (2002) "How Do Consumers Motivate Experts? Reputational Incentives in an Auto Repair Market." Journal of Law and Economics, pp. 437-68.

[11] Lacko, James (1986) "Product Quality and Information in the Used Car Market." Federal Trade Commission Staff Reports.

[12] Kim, Jae-Cheol (1985) "The Market for 'Lemons' Reconsidered: A Model of the Used Car Market with Asymmetric Information." American Economic Review, 75, pp. 836-43.

[13] Porter, Robert and Peter Sattler (1999) “Patterns of Trade in the Market for Used Durables: Theory and Evidence.” Manuscript, Northwestern University Economics Department.

[14] Pratt, Michael and George Hoffer (1984) “Test of the Lemons Model: Comment.” American Economic Review, 74, pp. 798-800.

[15] Schneider, Henry (2005) “A Field Experiment to Measure Agency Problems in Auto Repair.” Working Paper, Yale University.

[16] Wilson, Charles (1980) “The Nature of Equilibrium in Markets with Adverse Selection.” Bell Journal of Economics, 11, pp. 108-30.

[17] Womack, Jones, and Roos (1991) *The Machine That Changed the World: The Story of Lean Production.* Harper Perennial.

Table 1: Distribution Over the Number of Repairs Per Vehicle

n	$\hat{p}(n)_{\perp}$	$p(n)$
0	67.3%	71.3%
1	25.7%	20.0%
2	5.9%	6.0%
3+	1.2%	2.7%

Note: n signifies the number of repairs per vehicle. The middle column contains the predicted distribution over the number of repairs per vehicle assuming repairs across parts are independent events. The last column contains the empirical distribution over the actual number of repairs per vehicle. These frequencies are calculated from Consumer Expenditure Survey data.

Table 2: Repair Frequencies of Traded Versus Continuously-Held Vehicles

Recent Dealer Purchase	0.0016
	(0.0028)
Recent Private Purchase	0.0158
	(0.0070)
Mileage Diff.	2.7E-07
	(3.4E-08)
Log Repair Freq.	0.0182
	(0.0007)
Mileage Diff.*Repair Freq.	-1.6E-06
	(3.6E-07)
N	98,091
Pseudo R-squared	0.032
Predicted Prob. at X-bar	0.023

Note: The dependent variable is an indicator for repairs in that month. The explanatory variables are indicators for having purchased the vehicle in the previous 3 months from a dealer (Recent Dealer Purchase), having purchased the vehicle in the previous 3 month through private channels (Recent Private Purchase), the difference between the vehicle's mileage and the average mileage for all vehicles of that model, model year, and age (Mileage Diff.), the log of repair frequency for that model, model year, and age (Log Repair Freq.), and the interaction of the repair frequency and mileage difference (Mileage Diff.*Repair Freq.). The model assumes a Probit specification, and is estimated using monthly repair data from the Consumer Expenditure Survey. Coefficient estimates are expressed as marginal effects. Standard errors are expressed in parentheses.

Figure 1: Average Repair Expenditures

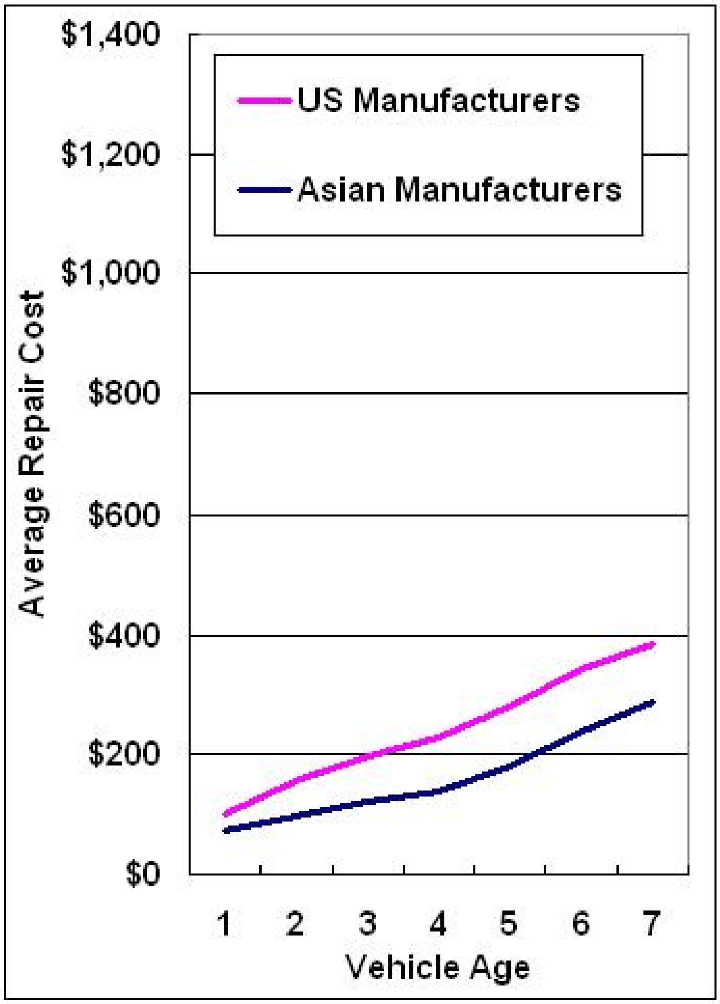

Note: The figure above shows repair expenditures on Asian and domestically-produced vehicles by vehicle age, based on data from the Consumer Expenditure Survey.

Figure 2: 90th Percentile of Conditional Repair Expenditures

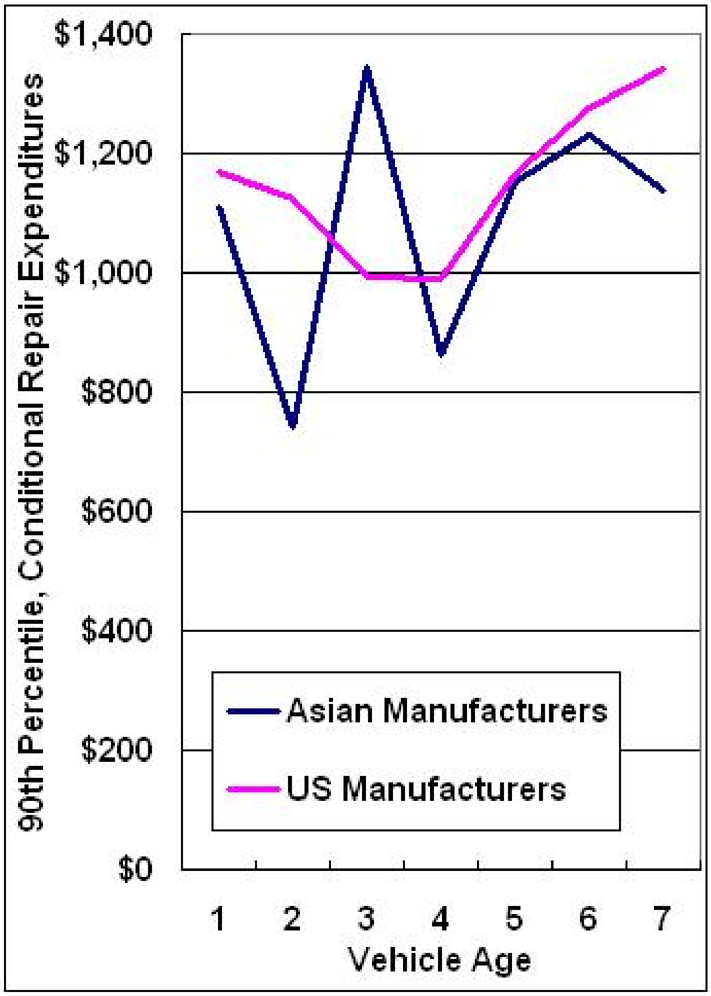

Note: The figure above shows the 90th percentile of repair expenditures on Asian and domestically-produced vehicles, conditional on repairs being conducted, by vehicle age. This figure is based on data from the Consumer Expenditure Survey.

Figure 3: Price Depreciation Curves For Hondas

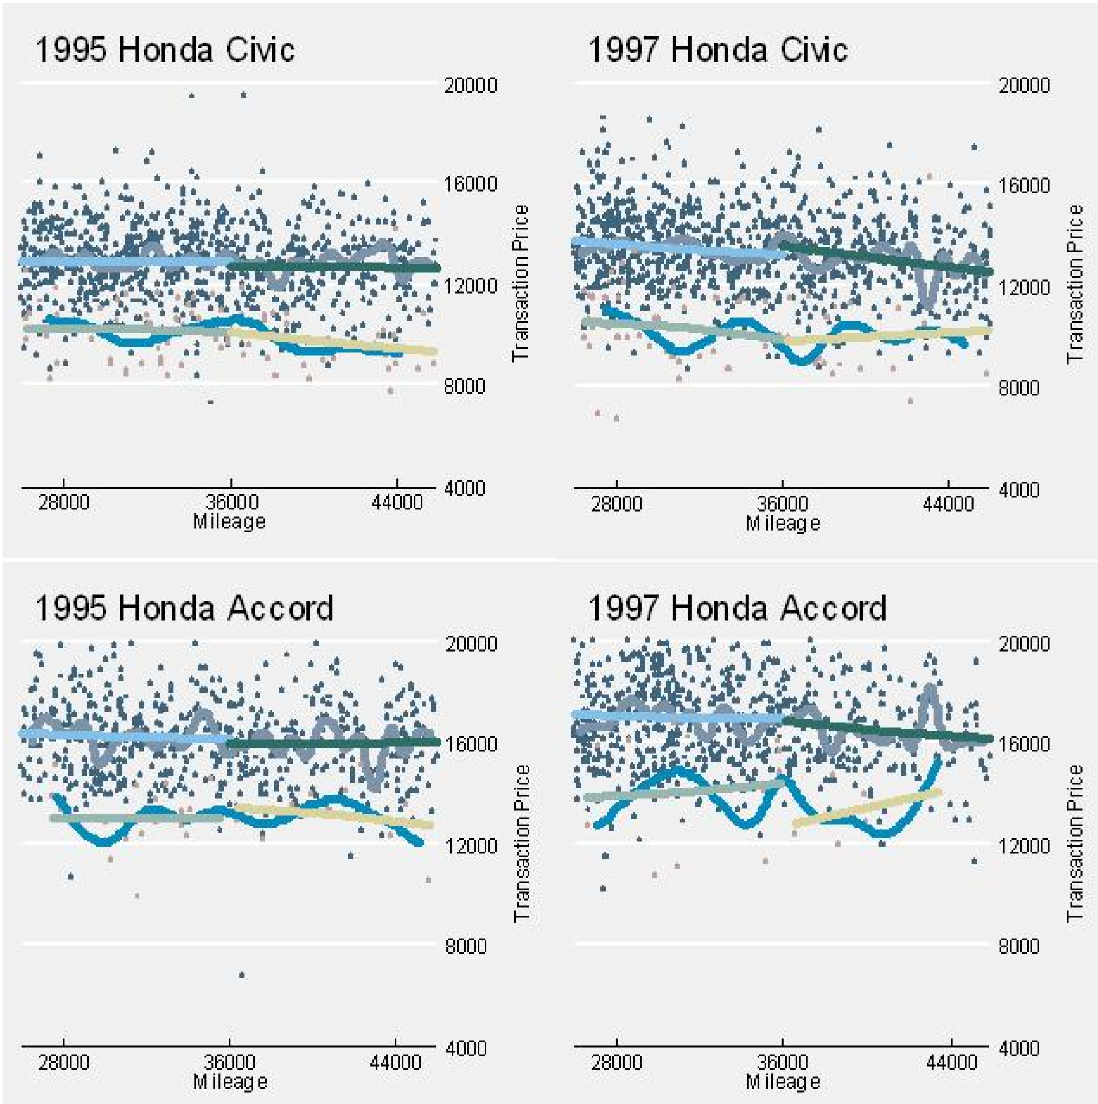

Note: The figures above show price depreciation curves for two sets of transactions for used Hondas Civics and Accords: The top transactions are franchised dealers trades to customers; the bottom transactions are customer trade-ins back to dealers. These vehicles contain prices for cars that are 2.0 to 2.9 years old, and have warranties that expire after 36,000 miles or 3 years. The two line segments for each set of transactions are linear fits of the price line before and after the warranty expires. I also allow the intercept to shift at warranty expiration. The wiggly curve is a median cubic spline fit of transaction prices.

Figure 4: Price Depreciation Curves For Fords

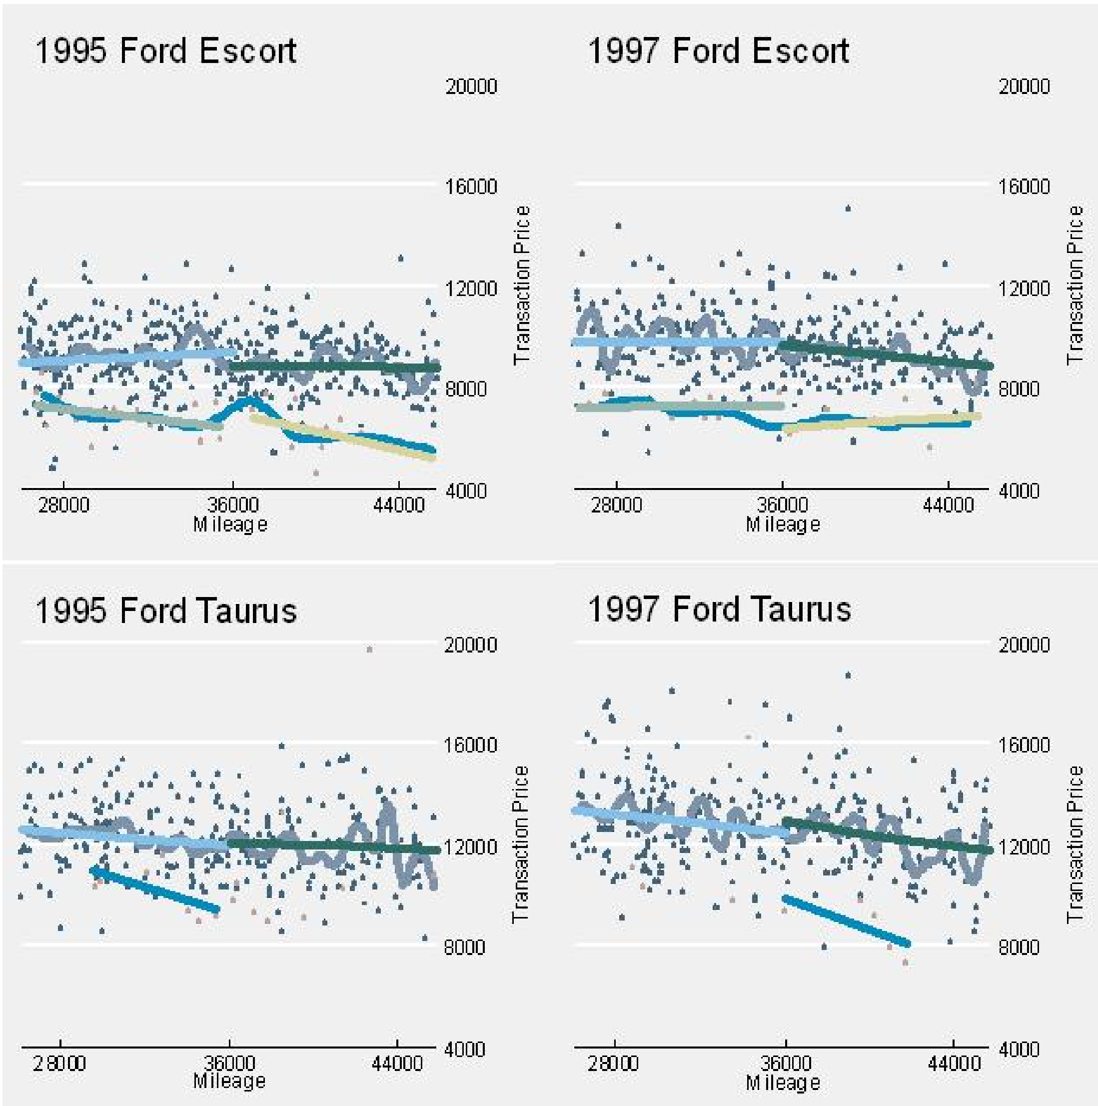

Note: The figures above show price depreciation curves for two sets of transactions for used Ford Escorts and Taurus's: The top transactions are franchised dealers trades to customers; the bottom transactions are customer trade-ins back to dealers. These vehicles contain prices for cars that are 2.0 to 2.9 years old, and have warranties that expire after 36,000 miles or 3 years. The two line segments for each set of transactions are linear fits of the price line before and after the warranty expires. I also allow the intercept to shift at warranty expiration. The wiggly curve is a median cubic spline fit of transaction prices.

Figure 5: Repair Frequencies for Traded Versus Continuously-Held Vehicles

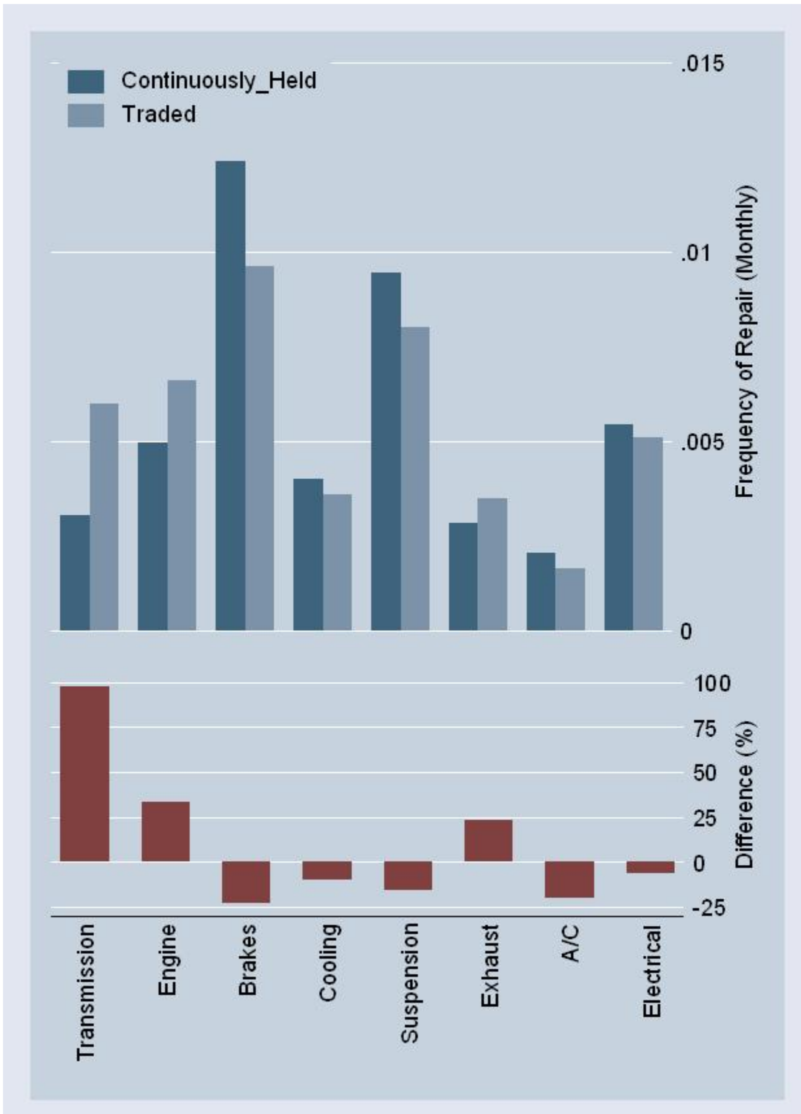

Note: In the top panel, the darker bars are estimates of the monthly repair frequencies for continuously-held cars, while the lighter bar are these repair frequencies for cars that were purchased within the 3 previous months. The bottom panel contains differences between these two estimates. Model, model year, and vehicle age have been controlled for. Transmission is the only category for which this difference is statistically distinguishable from zero.