Bias and Loss Aversion in the Market for Autos

November 2006
PRELIMINARY

Sharon Oster — Yale School of Management
Fiona Scott Morton — Yale School of Management and NBER

We thank Nick Barberis, Glenn Ellison, and Botond Koszegi for helpful comments. Kyle Hood provided research assistance. Address for correspondence: Yale School of Management, Box 208200, New Haven, CT 06520-8200. Email: Fiona.scottmorton@yale.edu, Sharon.oster@yale.edu.

1. Introduction

The fundamental theory of loss aversion, as presented in the work of Kahneman and Tversky (1979, 1984) posits that a consumer's utility depends on a reference point, and that as a result, losses from that reference point will be felt more strongly than gains. Empirical evidence of this effect has been found in the lab in numerous experiments beginning with the classic work of Kahneman, Knetsch and Thaler (1990), in which initially giving someone a mug increases his or her reservation price for that mug over someone given cash. Inducing loss aversion seems to be remarkably easy in the laboratory setting. For example, Kahneman, Knetsch and Thaler (1990) find that traveling in an elevator for a few minutes with a mug causes participants to value the mug more highly.

From the point of view of someone working on behavioral industrial organization, anomalies like loss aversion become most interesting when they redistribute rents in the economy relative to a pure neoclassical model. For this redistribution to be important, the biases must be significant and survive both consumer selection and learning that occurs in real markets, as well as profit-seeking behavior on the part of other, potentially non-biased, agents in the market. To our knowledge, loss aversion has not been studied outside the laboratory setting, other than possibly in work by List (2003) in the baseball card trading market. For our purposes, one of the interesting conclusions taken from List's trading card work (stressed by the University of Chicago website) is that loss aversion is not likely in real world trades involving professional or quasi-professional traders. In List's work, consumers learn from experience and their subsequently smaller bias has a minimal effect on rent distribution. However, in many other markets, particularly for durable goods like freezers, cars, houses, and sofas, consumers do not typically engage in dozens of transactions in a single year and one might well expect biases to survive. This paper examines an important market with low levels of trading to help determine the extent to which biases are manifest in markets and not just in laboratory settings.

In previous work we established that the price of a magazine subscription relative to its newsstand price reflects time-inconsistency on the part of consumers (Oster and Scott Morton (2005)). In that paper we took advantage of heterogeneity across magazines and the fact that publishers had available two pricing instruments, the spot price and the long-term contract (subscription) price, to identify the time inconsistency. Our findings in that paper indicate both that consumers are taking their behavioral anomalies to the market and that managers are choosing measurably different policies in response to that behavior. In this paper we similarly exploit difference in both product and individuals to examine whether there is any evidence that firms take loss aversion into account as they choose the form in which they give discounts off a list price.

In laboratory studies of loss aversion, a principal result is that fewer trades occur than one might expect. If consumers overvalue the loss from a held-good relative to the gains from a new good, barter opportunities are reduced. Thus, Knetsch (1989) finds that only 10% of his sample trade mugs and chocolate, despite an initial random assignment, rather than the predicted 50%. Demonstrating under-trading in marketplaces is more difficult because of the need to construct an appropriate counterfactual. In this paper, we have chosen instead to focus—as we did in the magazine market—on the way in which the bias affects the optimal pricing structure chosen by sellers. In particular, we will be focusing on pricing structures that are revenue-neutral to firms, but not utility neutral to consumers.

The setting we are using in this paper is the market for new cars and used trade-in cars. There are a number of features of this market that make it attractive from our perspective. First, within years and models, cars are homogeneous in characteristics (though perhaps not in quality) with commonly-known used values. As importantly from our perspective, dealers typically have some discretion in the way in which they extract value from consumers. While most car dealers post a list price for a new car, most buyers end up bargaining for a price below list. In some cases, the buyer is simultaneously negotiating for a price for a used car trade-in and the condition for executing the trade in is the simultaneous purchase of a new car. It is in this subset of car purchases that we are interested. At root, the dealer only cares about the profits he can obtain from the two linked transactions. For a new car listed at $25,000 and a trade in with a value of $5000, a dealer is indifferent between an offer of $24000 for the new car and paying $5000 for the trade-in, or an offer of $25000 for the new car while paying $6000 for the trade-in. Buyers who are loss averse, however, may be more price elastic over the trade-in than the new car. This might allow the dealer to raise the price he offers for the trade-in and the price he asks for the new car such that he is indifferent, but the buyer feels better off. With these linked transactions, deals may occur despite the fact of loss aversion given that there are essentially two margins on which the negotiation can occur. Note, although we focus on the used-car transaction, we will refer to the seller of the used car as the "buyer" throughout the paper.

The proposition that a used car might be subject to loss aversion bias has been suggested in the literature before. Indeed, multiple datasets have shown that the amounts car dealers offer for used trade-ins exceeds the value of those trade-ins on average, (e.g. Zettelmeyer, Scott Morton, and Silva-Risso (2001) and Goldberg (1996)). Again, this is possible only because those same dealers can extract off -setting rents from the tied part of the transaction, the haggling over the new car price. Using the subset of the data in ZSS (2001) that comprises our dataset, we find an average overpayment for a trade-in of $996. While this is certainly suggestive of some kind of bias on the part of a transacting party, it does not point to a single explanation.[1]

There is also experimental evidence that consumers prefer a subsidized trade-in even if it causes an equivalent increase in the price for the new car. Purohit (1995) tests precisely this proposition using a sample of MBA students. Having given them an inexpensive pen and announced a market value for it (supposedly derived from another class), he offered the students the chance to buy a more expensive pen using their recently acquired pen as partial payment. Half the students were told the trade-in value of the their pen was a dollar above market value, while the other half were told the trade-in value of their pen was a dollar below market value. The students who were given the subsidy had a higher mean *additional* willingness to pay for the expensive pen ($3.50 versus $2.50). "This suggests that subjects who are overpaid on their trade-ins (when they are sellers) experience a gain that then leads them to be more generous in determining their willingness to pay when they are buyers."

Heterogeneity in our setting will arise in two ways. First, consumers may vary in whether or not they have loss aversion. We rely on previously identified demographic variation in the presence

[1] Note, however, that in six years of presenting car papers, the second author has been asked many times why the trade-in overallowance has a positive mean, and has never heard a convincing explanation.

or strength of loss aversion to identify its existence in the market for used car trade-ins. We further use existing laboratory studies to help identify which types of cars are more or less likely to loss aversion in their owners. A second source of heterogeneity comes from the information set of consumers. To the extent that a consumer is perfectly informed about the dealer's reservation price for a new car, there is no room for a dealer to strategically shift values between trade-ins and new cars. When the consumer is imperfectly informed, the car dealer can strategically exploit any loss aversion by where he places the consumer's bargaining gains. Again, we will be able to use some demographic data to characterize the information set of consumers. Differences across consumers in expected information and loss aversion will then lead to predictions about the distribution of bargaining gains as between trade-ins and new car prices.

One objection that might arise in employing the used car market as a way to test for loss aversion is the importance of asymmetric information, and therefore adverse selection, in this market. For example, drivers may know their car has been exceptionally well treated and is in fact of higher value than might appear. In these cases, the driver's view of the value of the trade-in may be affected both by loss aversion *and also* by superior information. However, in our model we exploit exactly the tendency of buyers to not only have private information that their car is above average, but also to think it is above average when it is not. Thus the dealer will reduce any adverse selection he faces by engaging in the subsidization strategy we document. Furthermore, the dealer is a professional repeat-player with sources of market information and skilled mechanics readily available to him, so we expect him to do as well as possible conditional on all observables of the car and its owner. Lastly, this may not be as large a force as most economists expect; for example, the evidence in Schneider (2006) indicates that there is no significant adverse selection in used cars sold by franchised dealers.

The key to our analysis is the market linking between the two transactions, the trade-in and the new car purchase. Nevertheless, while these transactions are linked from the dealer perspective , there are differentiated from the view of the consumer. In particular, in our modeling, we will assume that consumers are subject to narrow framing as well as loss aversion. That is, individuals, as opposed to dealers, examine each of their transactions, the used car and the new car, in isolation. This is necessary for loss aversion to differentially affect one transaction.

Barberis and Huang (2006) discuss the role of the mix of loss aversion and narrow framing in the context of the equity puzzle. Laboratory evidence of such narrow framing is provided by Janiszweski and Cunha (2004) in looking at the way in which discounts are taken in bundled goods. They find that it matters to customers which of two products in a bundle is assigned a discount of the same dollar magnitude. Consumers prefer a bundle in which the less valued product is discounted. In our example, the bundle consists of a new car and a deal on the used car. Since the new car is the more expensive item, their theory would predict that the consumer will value a subsidy to the used car more than an equivalently higher price for the new car. This is what occurs in practice in our data.

To preview our results, we find shifting of bargaining gains to the trade-in in consumers who are typically thought to be both less well informed and more loss averse: poorer consumers (keep in mind everyone in our dataset is buying a new car), less well educated consumers, consumers from poorer neighborhoods, and minorities. These results are consistent with the strategic use of the loss aversion of these groups by dealers. We also find that market experience (proxied by age of the buyer) reduces the shifting of surplus towards the trade-in, while length of ownership (proxied by age of the trade-in) increases that shifting, both results again consistent with a strategic use by a dealer of the loss aversion of a consumer.

What are the welfare consequences to a consumer of having loss aversion in this market? Our exercise may appear inconsequential because the movement of rents between new and used car does not impact the consumer in a financial way. It is not clear which party gains from the shifting process. On the one hand, because the dealer can create consumer surplus by moving rents from the new car to the used trade-in, it gives him another instrument to extract rents overall from the consumer. On the other hand, the fact that the consumer truly believes her used car is worth more than it is may be a useful tool in a bargaining encounter. As noted above, we find that the same consumers who pay higher prices are those consumers with a high cross-subsidy (the socio-economically disadvantaged). However, the trade-in bias appears to be helpful to consumers; those consumers with a higher cross-subsidy pay less for the package of new car plus trade-in.

2. A Simple Model

Following Koszegi and Rabin (2003), we express the consumer's utility as a sum of consumption utility and a function of the distance from the reference point. We must therefore choose appropriate reference points for our application. Consider a consumer who is negotiating over a new car purchase and wants to trade-in a used car at the same time. As mentioned above, our first assumption is that these decisions exhibit narrow framing on the part of the consumer; the consumer considers each one separately. However, the dealer, an experienced trader, adds them together and evaluates the joint opportunity cost of the two cars and the net price. Throughout the analysis the dealer is not assumed to exhibit loss aversion on the new cars on his or her lot. This assumption is consistent with the work on List on professional baseball card traders (2003) and work by Novemsky and Kahneman (2005).

Next we model the information and expectations the consumer has with respect to the new car as well as the old car. We begin with the new car. The consumer is exposed to advertising of the new car which includes the Manufacturer's Suggested Retail Price (MSRP). In the US retail auto industry, manufacturers publicize MSRP in various media outlets so that consumers could quite plausibly become aware of MSRP more easily than other prices. In practice, manufacturers attempt to set MSRP above market price. Therefore, we assume that all potential purchasers of the car are informed about MSRP. Additionally, the dealer is assumed to have a discount, d, he will give on the list price. Consumers differ in how much they know about d. Highly-informed consumers (using the Internet or some other information source) are able to learn d, and they know it is equal to d^H. Transacting at this discount fully extracts the dealer's surplus. The members of this group have a reference price which is R_n = MSRP-d^H and we refer to them as informed consumers. Uninformed consumers believe that d is equal to d^L where $d^L < d^H$. Uninformed consumers therefore expect to get a discount off of MSRP, but before beginning to bargain they expect a smaller discount than the dealer is in fact willing to offer. The reference price of uninformed consumers is R_n = MSRP-d^L. Both types of consumers expect to purchase a new car at their reference price.

Now we turn to the trade-in vehicle. Both types of consumers overvalue their used cars and take as their reference point their upwardly biased opinion of the car's value. There are a number of reasons why consumers have a reference point above market price for their used car. First, there is ample evidence demonstrating that consumers view their own behavior (driving, repairs) as above average. This leads to a value for the vehicle which is higher than its true condition warrants. The optimistic bias for the used car is likely to be stronger than a similar optimistic bias that could work in the opposite direction on the new car reference price. For example, optimistic biases tend to be stronger for more ambiguous things.[2] A new car is a very specific homogeneous object, whereas the value and condition of a used car is less well-defined. Secondly, bias is stronger for things which are more ego-involved, such as the car the consumer has driven for years, relative to a brand-new car that is not associated with any person.[3]

Additionally, people may overestimate the car's value by ascribing their own tastes in features of the car (scratches, trim level, etc) as broadly representative of others' tastes. We also believe that drivers become fond of their cars, which increases their estimate of their car's value. Lastly, even a consumer who has searched in some kind of trade publication for the market price of the car may have the same bias. Typically in these services the user has to enter the "condition" of the car: excellent, good, fair, poor. Any optimistic or biased estimate of the condition of the car will lead to a predicted value of the car that is too high. The assumption of bias in the value of the used car creates a large area of potential losses for a consumer because her reference point (R_u) is above the market price of the car (P_u^*).

Consumers with loss aversion will experience any unexpected discount on the new car, any price below R^n, as a gain. On the other hand, if the dealer charges them more than R^n for the new car, they will experience that difference as a loss. The same holds for the used car; if a consumer transacted at market price, her losses would be $R_u - P_u^*$. For consumers with loss aversion, a dollar of difference between R_u and P_u^* is weighted more heavily (1+ a) than a dollar of gains (1). (As is standard in the literature, a>0.)

[2] Cite for this.
[3] Correct cite here.

Let us consider what the equilibrium prices look like for the uninformed consumers. The consumer expects to buy the new car at $MSRP - d^L$. The consumer also expects to sell her old car at $R_u > P_u^*$. Now the dealer faces the following problem. Any offer for the trade-in which is below reference price is counted by the consumer as a loss and weighted more heavily in their utility function. Because the uninformed consumer has a reference point above true market price for the new car, the dealer can create utility for this consumer by raising the price of the used car and the new car by the same amount. Call this transfer Δp. The dealer can create utility for the consumer by increasing Δp until the price of the used car hits the consumer's reference price. Thus the difference between d^H and d^L in equilibrium will be exactly this size: $d^H - d^L = R_u - P_u^*$. Because the consumer is unaware of d^H, she does not feel a loss due to not receiving it. Instead, Δp applied to the offer for the used car eliminates losses.

This strategy on the part of the dealer and the uninformed consumer creates a rational expectations equilibrium as follows. Consumers do in fact experience high prices for their own trade-in transactions, so their reference price formation and their expectation that they will sell to the dealer is rational. Likewise, they see a small discount on the new car and so their expectation of a small discount is rational. We make the further assumption that they obtain information from their friends' transaction and that their friends are all also uninformed. Therefore, they see high trade-in prices, a high probability of trade-in and a low new car discount in their sample of other transactions. The dealer is happy to engage in this strategy because he is no worse off financially, and the elimination of losses increases the chance that the consumer transacts at his dealership. The consumer transacts because she feels better off due to the cross-subsidy. The higher price for the trade-in prevents her from feeling a loss, while the higher new car price merely reduces her gains.

We note that there is a logical reason for the dealer to stop raising the two prices (stop cross-subsidizing) when the dealer has uncertainty about a consumer's used reference point. The model has presented the reference point as known, but in an empirical application we would expect a distribution of used reference points. As Δp increases, the used price moves up the used reference price distribution. For reasonable shapes of that distribution (thin tails, normal), at higher prices a given increase in Δp removes losses for fewer and fewer consumers. Most of them are not experiencing losses any more and are in the gains region where a dollar is only worth a dollar. The reference price on the new car puts a bound on how big Δp can be. Provided the dealer does not exceed the consumer's reference price for the new car, no consumer moves from gain to loss status. However, if the dealer continues to increase Δp, at some point the dealer hits the consumer's reference price and at that point the marginal cost of raising price becomes positive as the consumer now will be in the loss region on the new car. Thus marginal cost of raising the offer for the used car is increasing in Δp while the benefit is decreasing, so it is clear the dealer will find it optimal to choose a finite value of Δp.

Now take the case of the fully-informed consumer: Here the consumer knows d^H and is able to fully extract the dealer's surplus in the price of the new car. For a consumer with full information, loss aversion on the used car poses a barrier to transacting. The consumer will demand d^H, so that the dealer will have no surplus to load onto the used car to match the consumer's reference price there. Our prediction here is that the consumer is indifferent over buying the new car at the dealer; she will pay the correct market price regardless of the dealer with whom she transacts. However, she will not be able to get any dealer to pay her reference price for the used car because they have no surplus to carry over from the new car. She will feel a loss if she sells this car to the dealer, and may therefore prefer to sell it privately or give it to a family member or to charity. Her estimate of the value of the car is not confirmed in her market transaction, so in this element of the transaction she does not have rational expectations. With heterogeneity in transactions costs, we may see some consumers selling their used car to the dealer, but in this case we would expect to see a transaction price that is a market price, not a subsidized price. In some ways, this case is the classic example found in the experimental literature where loss averse subjects fail to transact. For these consumers, their perfect information on the "right" new car price prevents the dealer from finding a solution to the no-trade result.

This strategy on the part of the dealer and the informed consumer is internally consistent in several aspects also. Informed consumers see their own transactions and those of their (informed) friends. They obtain large discounts on new cars and thus confirm their price expectations about their new car. They also see the dealer offer low prices for the trade-in; however, they continue to believe that their car has a higher value than the dealer's offer for it. These two facts lead informed consumers to form a low expected probability of selling their used car to the dealer. This low probability of sale is confirmed when the dealer will not cross-subsidize the car and the consumer does not sell. (The fact of expecting to keep the car only exacerbates the "no sale" result as the consumer then feels more unexpected loss if she does sell the car.) The full information of the consumer in the new market prevents the dealer from moving any of the surplus on the new car to the consumer in the used market.

Next we consider the situation where the consumer does not suffer from loss aversion. Such a consumer gains nothing from cross-subsidization. In principle there could be any distribution of the discount between the new price and the used car. However, we assume that when a consumer is indifferent to the location of surplus, the dealer charges market prices. (We will modify this in the empirical work.) Additionally, we allow a consumer who does not suffer from loss aversion to have an estimate of the price of her trade-in that is too high.

If the consumer is informed, she will get the high discount on the new car and receive market price on her trade-in. Because she is not loss averse, when the dealer tells her that her car is only worth p_u^*, she does not feel the loss any differently than she feels a gain; therefore it is not attractive to compensate her more than a dollar for a dollar's loss on the used car. The dealer does not subsidize and the consumer will trade in her used car anyway.

An uniformed consumer without loss aversion will pay a high price for her new car and, by the same mechanism described above, get market price for her trade-in. She is indifferent about trading-in, like the informed unbiased consumer. However, the dealer earns substantial rent from this consumer's new car purchase, and therefore has an incentive to sweeten the trade-in offer to ensure she purchases the new car and does not go a competitor. Therefore, we expect this type of consumer to trade in her used vehicle.

A summary of the patterns generated by the model are contained in Table 1, below. Notice that only consumers who are loss averse and uniformed have a trade-in subsidy. We cannot calculate a trade-in subsidy for the type that does not trade: informed but loss averse consumers who will not be in the linked markets. The least intuitive type seems to us to be the rational but uniformed consumer, who overpays for her new car but immediately accepts the market value of her trade-in. Our prior is that there are relatively few consumers in this cell of the matrix.

Table 1: Predictions of the model		
	Loss aversion	No loss aversion
Uninformed	New car price: high Trade-in price: high $\Delta p > 0$ Trades	New car price: high Trade-in price: low $\Delta p = 0$ Trades
Informed	New car price: low Trade-in price: low $\Delta p = 0$ Does not trade	New car price: low Trade-in price: low $\Delta p = 0$ Trades at the margin

A graphical representation of the situation of loss averse uninformed consumers

Our description of the dealer strategies in the two linked markets when faced with loss averse versus unbiased customers can be well captured by a simple diagram that helps us in the empirical approach. In the data that we will use, we measure the gain to the consumer from the linked transactions package; call this V^T. Recall that consumers must bargain over the components of V^T, namely V^u and V^n.

Let $V^T{}_i$ =the total gain achieved by consumer i in buying a new car and trading in an old car

Let V_i^u = trade in price of consumer i's car less the actual value of that car

Let V_i^n = list car price to consumer i less the new car transaction price to consumer i

As discussed above, the dealer is indifferent between V^u and V^n. In the graph below, we represent this position as a series of equal isocost or isovalue lines, labelled $C_1 \ldots Cn$, or equivalently $V_1 \ldots Vn$, with a slope of negative one. Cost to the dealer and value to the consumer both increase in the northeast direction.

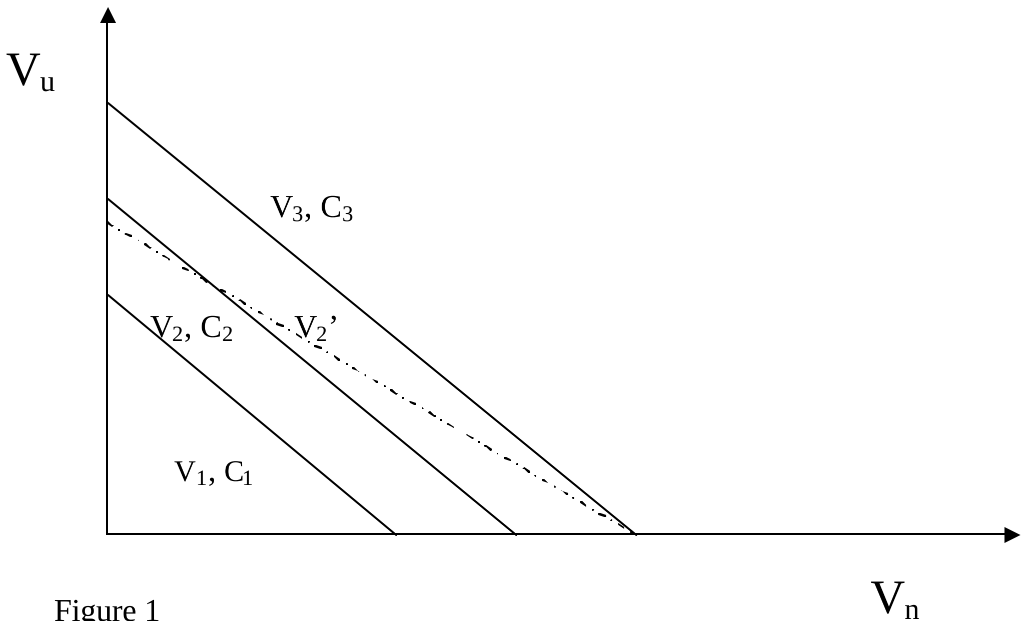

Figure 1

The dealer's problem is to achieve the deal at the lowest cost, while the consumer is interested in pushing for the highest V. With unbiased consumers and dealers, we will not be able to predict where on any given isocost / isovalue line we will end up. All we can predict is that as consumer bargaining power rises, we move to a higher V (C). For a consumer with loss aversion and particular reference points – discussed above – the value line is no longer coincident with the isocost line. Now we expect the V lines to be flatter than the C lines, as in V_2 in our example above. Using our earlier notation, we can represent loss aversion most simply as:

$$V^T_i = \alpha\, V_i^u + V_i^n \qquad (4)$$

where $\alpha>1$.

It follows simply that the optimal strategy for a dealer faced with a consumer with loss aversion will be to choose a Δp that puts more value into the trade-in, since value there has more weight. That is, in Figure 1, the equilibria will all lie closer to the y-axis than they otherwise would.

We do not have data on which consumers are informed or which exhibit loss aversion. However, we can generate some testable hypotheses without this information.

The propensity to trade-in a used car is correlated with the contract price of the new car.
The used car subsidy is correlated with the contract price of the new car.

Additionally, we will use existing evidence from field and laboratory experiments as to which consumers and which transactions should display loss aversion.

Uninformed consumers with loss aversion should have a higher cross-subsidy (and take a greater portion of their gains) in the trade-in compared to rational, informed consumers.

Transactions with trade-ins that have characteristics that generate biased expectations should exhibit a higher cross-subsidy of the trade-in.

Lastly, we would like to present some evidence on consumer welfare. If the dealer simply moves money from one bucket into another, there is no effect on relative rents. One hypothesis is that the dealer can keep some of the rents he transfers. In particular, biased valuations and imperfect information may be correlated with general lack of bargaining sophistication. A second hypothesis is that having a goal higher than the market value of the car helps a consumer in bargaining with a dealer. In the empirical work we will ask whether, for example, we see a correlation in either direction between the profit margin on the total package and Δp for consumers who trade in their used cars.

Consumers with a higher cross-subsidy may gain or lose on the total package.

3. Data and Estimation Strategy

The dataset we use comprises about eight hundred thousand new car purchases at a sample of new car dealerships across the US in 1999. The observations were collected by a market research firm, hereafter referred to as MRF. About 300,000 of these purchases were accompanied by a trade-in transaction and it is this part of the data set we will focus on.

The dataset contains a great deal of detailed information about both the new car and the trade-in. For the new car, we have data on both the transaction price and the average sales price for that same make and model at a very precise level (including, for example, the engine type, number of doors, transmission and trim level). These data will allow us to construct a variable measuring how "good a deal" the buyer got on his or her new car purchase. Our data also report both the transaction price of the trade-in and also the amount at which the dealer's internal accounting system values the car.

The timing of the different steps of the transaction is important. Typically the consumer who wants a new car brings her used car to the dealership. The dealer's mechanics look at the exterior of the car, see its age, mileage, options, and perhaps raise the hood for a visual examination. Then the dealer offers the consumer a trade-in value for the car. This is the "transaction price" we measure for the trade-in. After the transaction is successfully consummated, mechanics at the dealership examine the car carefully. At this point, its value is entered into the dealer's accounting system; this is the "value" of the traded car we use in our analysis. We assume this value is correct and there is no further unobserved quality. In fact, it is in the interest of the dealership to correctly measure the car's value, as often a salesman's commission depends on the net profit of the transaction. Thus, we are able to measure the transaction price relative to the dealer valuation to discover how well the consumer did on the sale of the trade-in.[4]

We first construct several measures of the implied subsidies of the different transactions. The first is *Normalized Used Subsidy* (*NUS*). This is the contracted trade-in dollar amount less the booked

[4] Clearly, a consumer will do better by our measure if the defects on her car are hidden from casual inspection but revealed in the full inspection. We will return to this issue later in the paper.

value of the trade-in divided by the booked value of the trade-in. This provides a percentage measure of the amount of over-payment for the trade-in. We also created *Used Subsidy* (*US*), which is just the dollar difference between the contracted and booked amounts. We found that there was considerable measurement error in the dataset, so we eliminate from the sample those observations where *NUS* is greater than 1.69 or less than -.14. This removed the top and bottom 5% of the data.

Next we created a mean price for each car (recall that a car is a very specific good). The mean price includes transactions with and without trade-ins. In order to be taking a mean of apples and apples, we remove the trade-in subsidy when calculating the price that goes in to this expression. We then define *New Relative Price* (*NRP*) as the contract price of the new car, less the mean price, all divided by the mean price. Again, this creates a measure of the price of the car relative to what others paid. We remove the top and bottom 1% of the sample, as it appears there is also some measurement error in the new price.

We want a measure of *Dp* in order to analyze which transactions have a high cross-subsidy. We proceed as follows. First we take the 10^{th} percentile of the new car price distribution for each car (again, this includes transactions with and without trade-ins and removes the subsidies for comparability). Next, we subtract this 10^{th} percentile figure from the contract price of the new car to get *New OverPayment*, (*NOP*), which tells us how many dollars the consumer left on the table, relative to being approximately the best bargainer in the market. Suppose both *US* and *NOP* were positive. In that case one can see that the cross subsidy would be the smaller of the two. We define *Dp* as the minimum of *US* and *NOP*, or zero if the minimum is negative.[5] The difficulty with this measure of *Dp* is that it is denominated in dollars, and so we might expect it to be larger for higher absolute values of new and used cars. We confirm that the value of the used trade-in and the mean price of the new car a particular consumer sells and buys, respectively, are correlated in our dataset, as one might expect (.5). Thus we want to find some way to normalize *Dp*. We divide *Dp* by the booked value of the trade-in, to create a measure called *normdP*.

[5] We also remove the top 1% of Δp values from the sample.

The variables just described allow us to capture the extent to which a buyer strikes an advantageous bargain in either the new or used car part of the transaction, and measures any correlation between the two. We use two types of data to predict when bias and loss aversion will be strongest for buyers and thus when more of the value will be loaded onto the trade-in piece: car characteristics and buyer characteristics.

Buyer Characteristics

One of the early results of Knetsch and Sinden (1984) in an experimental setting is that when individuals make choices on behalf of other people, they do not display loss aversion. In the case of car deals, it is clear that there are times in which someone in a family other than the principal driver actually negotiates the deal. Unfortunately in our data set, we do not know whether the car traded in was driven by the buyer or someone else in the family. We assume that frequently the two will be the same – in single-person households, for example, and some substantial fraction of the time in a typical married household. However, casual observation of household behavior would suggest that when the buyer is female, it will disproportionately be the case that she is trading in her own car; i.e. wives are less likely to negotiate on behalf of their husbands but the converse is often true. Thus, when the buyer is male, he could be trading in a car he usually drove, or a car his wife usually drove. This suggests, all else equal, that loss aversion will be stronger for the female buyer since the probability is higher that the trade-in is her own. Therefore we include an indicator for the gender of the buyer.[6]

Papers by List (2003, 2004) show that market experience is important in reducing loss aversion. The more often a consumer has traded in the past, the more able he or she is to identify gains from trade and trade accordingly. The best measure we have for market experience is age. In the US, where car-ownership is almost universal, age is likely to be highly correlated with the number of cars a person has owned, and perhaps transacted over, in the past. The transaction dataset provides the age of the buyer. We include age as a spline, as we have no pre-existing view on the functional form it should take. The groups are 25 and below, 26-34, 35-45, 46-55, and 56-65. The omitted group is 66 and above.

[6] Our dataset measures gender by analyzing the first name of the buyer.

There is surprisingly little experimental work on the demographic characteristics that minimize or accentuate behavioral biases. Recent work by Chen et al (2006) on loss aversion in Capuchin monkeys suggests that biases may be at least to some extent built into the brain rather than socially learned. However, there is a small literature on the relationship of demographics to behavioral biases.

List (2004) studies the impact of other demographic characteristics on willingness to trade (lack of endowment effect). He finds no significant effect of demographics: no gender effect, no age effect, and no effect of education.

We also draw on work that suggests that impulsive instincts in general may be partly overcome by cognitive reasoning, suggesting that phenomena like loss aversion might be more modest in individuals with higher cognitive skills (Loewenstein and O'Donoghue (2004)). Recent work by Benjamin and Shapiro (2006) finds laboratory evidence that both small stakes risk aversion and short run discounting are less common among students with higher standardized test scores. One might similarly predict that greater cognitive skills might moderate loss aversion.

The dataset does not, of course, have data on the cognitive abilities of the buyers. We do, however, have data on the census block group of the buyer which allows us to use the census to extract neighborhood data for each buyer. Variables include percent of the population with a college degree, percent who are high school dropouts, percentage blue collar, professional, executive, and averages such as household size, house value, and income. Achievement in the US has been shown to depend both on native ability of the person in question and, of course, on the resources provided by the person's family (Behrman, et al. 1989). Thus we do not assert that average neighborhood income, for example, is an accurate measure of the SAT scores of all of the residents of the neighborhood. However, despite the family component, income will be correlated with cognitive ability due to the financial returns to ability in the economy. Likewise, we are helped by the substantial social and geographic mobility in the US, which causes people to sort into neighborhoods with residents of similar characteristics (Bayer et al (2004)).

The ethnicity frequencies large enough to analyze are Hispanic and Asian. We also know the African-American percentage in the census block. We have no prediction concerning race and ethnicity and the strength of loss aversion. However, in the US, race and ethnicity are correlated with other demographic variables such as education and income, so we expect to see coefficients on race and ethnicity variables that mirror those patterns.

Characteristics of the Trade-in

Strahilevitz and Loewenstein (1998) find that the longer individuals are endowed with an object, the higher their valuation of that object, suggesting that long ownership might accentuate the biases we discuss. Thus, we suggest that the length of time the seller has owned the trade-in will be an important factor in whether or not the person has a biased valuation. Additionally, the older a car gets, the more variance there will be in condition within the model, and therefore the more leeway the consumer will have to exercise her bias and still have statistical support. We do not how long consumers have owned their cars, unfortunately, but we do know the age of the trade-in. At least some of our buyers will have bought their trade-in as a new car. We include the age of the trade-in as a spline as we do not want to impose a linear functional form on the data. Zero to three years is the first category because this is the period when the warrantee holds. A person who always wants a car under warrantee will trade-in at this point. Three to seven years old is the next group, followed by eight to ten. The omitted category is over ten years old. Our expectation is that valuation bias, and thus the amount of value captured by the trade-in, will increase with the age of the car.

The age of the trade-in may also measure a second effect which has been established in the experimental literature by Chapman (1998). This paper shows that the more similar the traded objects are (e.g. two kinds of pen rather than mug and chocolate), the lower the loss aversion. In our setting a person is trading in an old car for money, but may perceive she is trading in an old car for a new car (and some negative money). If this is the case, then perhaps consumers trading in newer used cars will show less loss aversion compared to those trading in older cars that are less good substitutes for the new car. Novemsky and Kahneman (2005) refer to this effect and use traded-in cars as their example of similar goods. Clearly a newer used car will be a closer substitute to a brand new car, in terms of the benefits it can deliver, than an older used car, and so may generate less loss aversion. All these age effects operate in the same direction, so we will not be able to distinguish them in our results.

An off-setting effect mentioned by Purohit (1995) concerns the subsidy as a proportion of market price. As the trade-in gets older, it is worth less, and a subsidy of a particular dollar size becomes an ever larger proportion of its market price, while remaining unchanged as a proportion of the new car price. Thus, if the consumer is convinced by proportion, rather than absolute dollar amounts (in contrast to our model), as the car ages, less subsidy might be as effective in generating a transaction compared to a case where the trade-in was newer.

Just as we observe that behavioral biases can be mitigated by higher level reasoning, so one might expect emotional connections to accentuate these effects. Lerner, Small, and Loewenstein (2004) find that some emotional states like disgust can eliminate loss aversion. We hypothesize that emotional connections are likely to be positive and strongest for an individual's first car. While we do not have information on whether or not a trade in is a first car, we do know the age of the car buyer. Thus, we judge that a car eight years old or more, traded in by a driver under 34 years of age is likely to be trading in his or her first car (*firstcar*). We would expect bias to be stronger for these younger drivers.

Likewise, bias might differ across car subsegment because of the different purposes for which different subsegments are used and therefore the different emotions they arouse. A fun, sporty car might give rise to a different bias than a boring minivan. Okada (2001) finds evidence that frequent pleasant use of a durable good leads the owner to reduce her mental book value of the durable, relative to a good with negative associations. An object with a lower mental book value will be willingly traded at a lower price. Under this theory, the effect of the pleasurable emotion will be to reduce optimistic bias and reduce the price requested for the trade-in (and increase trading in used cars). We assign the traditional subsegments used by the industry to three categories: *fun*, *luxury*, and other.

Summary statistics for the variables of most interest are contained in Table 1.

4. Results

First we present basic evidence in Table 2 that consumers are more likely to trade-in a vehicle when the contract price on the new car is high. Further, a higher price on the new car is correlated with a larger subsidy on the trade-in, as we speculated. This is true both if one examines unconditional correlations, or runs regressions conditional on demographics and regional indicators. Note that the sample examining whether the new car contract price is correlated with the probability of trading-in a used car is the full sample, whereas the correlation between *NUS* and *NRP* uses only transactions with a trade-in.

Next we examine what demographic types have a large cross-subsidy. Looking at Table 3, we see that African-Americans, Hispanics, young people, low income people, low college education, low house value, low executive, and high blue collar are more likely to have a high cross-subsidy for their trade-in. These results indicate that consumers with a high trade-in subsidy exhibit a classic list of characteristics associated with low SES.

The estimated coefficient on female is positive and significant. This may indicate that women have a larger bias for their used cars and/or have more loss aversion, or it may indicate that women are more likely to be transacting over a car they actually drove. Suppose both genders have a biased valuation for a car they drove, but less bias for a car they did not drive. We would estimate a stronger effect for women because when they trade, they trade their own car, whereas men often trade someone else's car.

We further find strong evidence that market experience reduces the effect of loss aversion. The coefficients on the age spline are monotonically decreasing: the youngest age group does best on the trade-in, 14% more than the omitted group, followed by a steady decline (9%, 5%, 2%, 0%) until the 66 and over group.

Table 4 introduces features of the car that the literature suggests should matter to either the level of loss aversion or the amount of bias. Age of the trade-in vehicle is the most important. Buyers do better negotiating for cars they have owned for more than three years, but are not ancient. Cars older than 10 years old are the omitted group and by construction have a zero effect on the cross-subsidy. The age effect is consistent both bias increasing with the length of ownership and also with the old and new cars being worse substitutes as the trade-in ages. The effect varies across specifications somewhat; in column 1 the 0-3 age category has an estimated coefficient of .021, while the older categories 4-7 and 8-10 have estimated coefficients of .033.

In column 2, our measure of whether the car is a 'first car' has a negative and significant coefficient, which is the opposite of our theoretical prediction that positive emotion enhances loss aversion. It is consistent with Okada (2001). Recall his theory is that pleasurable use lowers the mental accounting value of the durable good, making an owner more willing to trade. We also add the indicators *fun* and *luxury* which refer to the sub-segment of the traded-in car. Sporty categories form the first variable and the union of all the luxury categories forms the second. Buyers extract a higher cross-subsidy on a fun car, while luxury cars are no different from other types. The Okada (2001) theory is again supported by this result.

Column 2 continues to display a strong age pattern (.19, .10, .05, .025, 0) as we found in column 1. Likewise, the age of the trade-in has a positive effect on the cross-subsidy (.020, .030, .037) until the eleven and older category.

We now return to the issue of unobservable quality. Recall that in our setting, the dealer learns about all unobservable quality upon carrying out a full inspection. The final booked value of the car in the denominator contains all information about the value and condition of the car. A buyer with a car that looks good on the outside, but has a ruined transmission, might be offered an average price for the car, but then the dealer would revise down the booked value upon inspection. Then this consumer's *NUS* would be relatively high due to the new information. Our empirical findings are consistent with this story, in the sense that the highest values of *NUS* are earned by the socio-economically most disadvantaged buyers. Thus it would not be surprising if these cars were discovered to have below average condition after trade-in offers were made.

Of course, if this were an empirical regularity, we would expect the salesman to correct for it in the average level of his offers. He is able to observe the socio-demographic status of the buyer with some accuracy due to her presence at the dealership and the fact that car salesmen are selected to have this kind of observational skill. He can adjust his offers downwards if he suspects low condition that will be discovered *ex post*. Secondly, the salesman is typically compensated by a percentage of the net margin on the deal. The margin on the deal is calculated using the booked value of the trade-in. Thus if the trade-in is worth less than the salesperson thought, his financial reward falls. So it is not in the interest of the salesman to make persistent *mistakes* about the value of the car – though of course he may choose to offer more than the trade-in is worth if it creates a deal where there would not otherwise be one, as in our model above.

Note also that any feature of the car or buyer that is observable to the dealer when he makes the trade-in offer cannot be information learned from the inspection that lowers the booked value of the car. The trade-in's age, for example, and the age and gender of the buyer are fully known at the time of the trade-in offer.

We briefly engage in some robustness checks by running the same regression on different sub-samples of the data. A key feature of our setting is the flexibility of both transaction prices, which allows bias and loss aversion to have an impact. This flexibility will not exist in negotiations where the two parties first agree on the new car price before discussing the trade-in. With this timing, the dealer commits to a price at which he will sell the new car and cannot raise that price. Only later does the consumer physically arrive at the dealership with her trade-in, and at that point the dealer can examine the trade-in and make an offer. (Note that dealers will not make binding trade-in offers without physical inspection of the used vehicle for obvious reasons.) Buyers who have used the Internet to get a price quote would be an example of a group where the bargaining over the trade-in might be independent of the new car price because that has been fixed in a previous negotiation. In our sample, we know who used Autobytel.com, an Internet service that puts buyers in touch with dealers who make new car offers over the phone or by email. As expected, we find that on average, *normdP* is 4% smaller than it is for other buyers. (dP is $238 less.)

In our results thus far we have found that the groups least likely to use the Internet (socio-economically disadvantaged) are the ones who get the most for their trade-ins. Our concern is that perhaps every consumer has the same bias and loss aversion but the Internet shoppers do not show it because the dealer can not trade off between the two prices. In Table 5, column 1, we restrict the sample to buyers who used Autobytel.com. The results will be weaker and less significant due to the much reduced sample size (265,000 down to 12,000). However, we can see the same patterns persist: African-American online consumers continue to have a higher cross-subsidy, as do young people and people with less expensive houses. The effect of trade-in age, however, reverses its pattern: young trade-ins have a greater subsidy than old ones. College, executive, blue collar, and Hispanic are no longer significant, and female is also no longer significant. However, many of these variables display estimated coefficients similar to those generated by the full sample, though their standard errors are higher.[7] We conclude that the demographic results hold up in the Internet sample, but not those for the trade-in bias. This finding needs further investigation.

We also limit the sample to those who did not obtain financing at the dealership. This group is likely to be more financially sophisticated than average: either they have cash to pay for the car or they are able to obtain an auto loan at their bank or credit union.[8] Indeed, we find that their cross-subsidy is lower by 6% (*normdP*) or $333 in absolute dollars on average. Nonetheless, this sophisticated group generates similar results: lower income, lower house values, more blue collar, more African-American, more female, and younger consumers have a higher cross-subsidy. Perhaps because of the smaller sample size, we lose significance of college, executive, and Hispanic. Older trade-ins have a higher cross-subsidy. 74% of transactions with a trade-in obtain financing from the dealer, so our sample size for this specification is about 70,000. We conclude that more financially sophisticated people still exhibit the patterns we expect from bias and loss aversion.

---

[7] *Female* is mis-measured in our dataset, so it is not surprising it suffers attenuation bias in a smaller sample. MRF takes the first name on the purchase contract and uses it as the buyer's first name, which creates a problem in the case of joint ownership. Thus if the buyers are Jane and John Doe, the buyer is coded as female. If the buyers are John and Jane Doe, the buyer is coded male.

[8] Ayres, cite

Finally, we measure the correlation between *normdP* and total surplus overall on the combined package (*Total Price*). We do this using the whole sample, in which case we control for having a trade-in, and we also use only the trade-in sub-sample. These results are reported in Table 6. We also repeat the specification using a measure of surplus on the package that controls for model, month, region, and so forth, as we expect these factors to affect the price of the car and we do not know if they are correlated with a tendency to have a subsidy. This dependent variables is called *Adjusted Price*. Finally, we run the complete price specification from earlier work (ZSS (2001)) and show that *normdP* has a negative impact on total price paid for the package above and beyond the demographic factors we have already identified.

In short, we find that those buyers who have a higher cross-subsidy for their trade-in also have a negative error term – they pay less than others – in the overall price of the package. The low SES demographics that ZSS identified as being associated with higher prices continue to predict higher prices as well as higher cross-subsidies. However, conditional on demographics, a higher cross-subsidy leads to lower net price on the package. Our interpretation of this finding is that a biased, uninformed, but stubborn consumer, who thinks her trade-in is worth more than it really is, gains a bargaining advantage over the dealer relative to the situation with no trade-in.

5. Conclusion

We find that, in bargaining over a trade-in vehicle, a consumer does better if she has owned the car for a longer time (proxied by trade-in age) and if she is young and therefore without much market experience buying and selling cars. These results are consistent with loss aversion affecting trade in the market for used cars. Car dealers will shift money between new and used car prices in order to create utility for the consumer. Consumers who appear to have strong loss aversion and therefore desire this shift are lower income, blue collar, minority, non-college educated, and who live in poorer neighborhoods. Women may exhibit stronger loss aversion but in our data we cannot separate the differences between genders and the “driver” effect.

Our setting has advantages for studying the effects of bias and loss aversion because of the ease with which the salesman can move money from one product to the other. However, that ease arises from the underlying fact that the two sources of money are identical for the dealer and for the consumer's bank account, though not perhaps for the consumer's utility. Given the potential lack of a financial effect in our setting, we are interested in determining if there are welfare consequences from loss aversion and biased values for used cars. When we perform empirical tests of the relationship between the net package price and the amount of cross-subsidy, we find that a higher cross-subsidy redistributes rents from the dealer to the consumer. We interpret this finding to mean that the optimistic value the consumer places on her trade-in serves as a valuable bargaining tool against the dealer. Furthermore, the biases we study have consequences for welfare and profits.

Second, we infer from the fact that the dealer adjusts the price of the trade-in that he gains from doing so – he reduces the probability of losing the combined deal and induces consumers to trade. If there is no such margin to adjust, some owners of used cars would not upgrade their used car to a new car. Our results provide predictions for which subsegments of the used car market will be disproportionately active in the absence of the ability to subsidize the trade-in: cars owned by older people, younger cars, cars driven by men, and cars owned by wealthier and better educated people.

The market for cars is not one in which most consumers trade frequently. For example, a baseball card collector probably trades many more cards in a five year span than a typical new-car consumer trades cars. This leads us to reconcile the difference between our findings and List (2004) by pointing to differences between consumers' market experience in the two contexts. Of course, we also find that market experience is important in reducing bias, but we note that our estimated coefficient falls slowly across 45 years of market experience. So learning appears to occur fairly slowly in the car market. Our findings demonstrate that loss aversion remains a significant feature of one important real-world market, and suggests to us that it may be a feature of other markets where the frequency of trading is not high.

References

Barberis, Nicholas and Ming Huang (2006) "Loss Aversion / Narrow Framing Approach to the Equity Premium Puzzle" draft Yale School of Management

Bayer, Patrick, Robert McMillan and Kim S. Reuben (2004), "What Drives Racial Segregation? New Evidence Using Census Microdata," Journal of Urban Economics, 56(3): 514-35.

Behrman, Jere R., Robert A. Pollak and Paul Taubman (1989), "Family Resources, Family Size, and Access to Financing for College Education," Journal of Political Economy, 97(2): 398-419.

Benjamin and Shapiro (2006) "Title?" draft, University of Chicago.

Chapman, Gretchen (1998) "Similarity and Reluctance to Trade" *Journal of Behavioral Decision-making*: 11:1 (December) 47-58.

Chen, M. Keith, Vankat Lakshimarayanan, and Laurie Santos (2006) "How Basic are Behavioral Biases? Evidence from Capuchin Monkey Trading Behavior" forthcoming *Journal of Political Economy.*

Janiszweski, Chris and Marcus Cunha (2004) "The Influence of Price Discount Framing on the Evaluation of a Product Bundle" *Journal of Consumer Research*, 30 (March) 534-546.

Khaneman, Daniel, Jack Knetsch, and Richard Thaler (1990), "Experimental Tests of the Endowment Effect and the Coase Theorem," *Journal of Political Economy*, 98 (December) 1325-48.

Khaneman, Daniel and Amos Tversky (1979) "Prospect Theory: An Analysis of Decision Under Risk" *Econometrica* 47 (March): 263-91

--------- and -------- (1984) "Choices, Values, and Frames," *American Psychologist*, 39 (April) 341-50.

Knetsch, Jack (1989) "The Endowment Effect and Evidence of Noreversible Indifference Curves" American Economic Review, 79 (December), 1277-84.

Knetsch, Jack and J.A. Sinden (1984) "Willingness to Pay and Compensation Demanded: Experimental Evidence of an Unexpected disparity in Measures of Value" *Quarterly Journal of Economics*: 99 (3): 507-521.

Lerner, Jennifer, Deborah Small, and George Loewenstein (2004) "Heart Strings and Purse Strings: The Carry-Over of Effects of Emotions on Economic Transactions" *Psychological Science:*15:5:337-341

List, John (2003) "Does Market Experience Eliminate Market Anomalies?" *Quarterly Journal of Economics* (Feb) 41-71.

List, John (2004) “Neoclassical Theory versus Prospect Theory: Evidence from the Marketplace” *Econometrica*: 72:2:615-25.

Loewenstein and O’Donoghue (2004) “Animal Spirits: Affective and Deliberative Processes in Economic Behavior” draft, Carnegie Mellon University, Cornell University.

Novemsky, Nathan and Daniel Kahneman (2005) “The Boundaries of Loss Aversion,” *Journal of Marketing Research*, Vol 42 (May), 119-128.

Okada, Erica (2001) “Trade-ins, Mental Accounting, and Product Replacement Decisions” *Journal of Consumer Research*: 27: 433-446.

Oster, Sharon and Fiona Scott Morton (2005) “Behavioral Biases Meet the Market: the Case of Magazine Subscription Prices” *Berkeley Electronic Journals in Economic Analysis & Policy Advances*:5:1

Purohit, Devavrat (1995) “Playing the Role of Buyer and Seller: The Mental Accounting of Trade-ins” *Marketing Letters*:6:2:101-110.

Strahilevitz and Loewenstein (1998) “The Effect of Ownership History on the Valuation of Objects” *Journal of Consumer Research*, 25: 276-289

Scott Morton, Fiona, Florian Zettelmeyer, and Jorge Silva-Risso (2001) “Internet Car Retailing” *The Journal of Industrial Economics*:49:4

Zettelmeyer, Florian, Fiona Scott Morton, and Jorge Silva-Risso (2003) “Consumer Information and Discrimination: Does the Internet Affect the Pricing of New Cars to Women and Minorities?” *Quantitative Marketing and Economics*:1:1:65-92, 2003.

Table 1. Summary Statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
TRADEACV	256365	7628.275	6541.288	0	187000
US	256364	876.8989	1480.503	-6000	19732
NUS	256139	.1718532	.3088756	-.1398	1.689474
cprice	693336	23790.52	8121.045	7084	88518
NRP	693336	.0155203	.0772464	-.1789823	.298987
dP	256393	868.5071	1351.825	0	9028
normdP	256139	.169144	.2961153	0	1.689474
lnprice	693336	10.01045	.3214808	8.690474	11.41559
tradeage	245804	5.530569	3.337363	0	18
female	693336	.3578856	.4793786	0	1
custage	693336	43.87526	14.07768	16	100
pctblk	693336	5.821779	14.25296	0	100
pcthisp	693336	8.123472	10.19755	0	55.32995
pctasia	693336	4.913506	7.917888	0	100
college	693336	31.33289	17.78549	0	100
income	693336	57115.5	25059.62	10403	150000
lesshs	693336	12.27969	10.42061	0	100
exec	693336	17.54538	8.065803	0	100
blue	693336	25.95208	14.90583	0	100
hsval	693336	165813.3	100029.3	7500	500000

Table 2. Correlations				
Unconditional correlations				
			Correlation (p-value)	Sample
NUS		NRP	.384 (.000)	trade-ins only
Trade-in		NRP	.238 (.000)	all
Conditional regressions				
Dependent variable	*Controls*	*Indep. var.*	*Coefficient (s.e.)*	*N*
Trade-in (probit)	constant	NRP	3.85 (0017)	744,571
Trade-in (probit)	demographics, region	NRP	4.02 (.023)	688518
NUS	constant	NRP	1.33 (006)	266,800
NUS	demographics, region	NRP	1.33 (.007)	256139

Table 3: Demographic determinants of normdP		
Variable	(1) tobit	(2) tobit vans only
Female	.018* (.002)	
Pct Black	.0014* (.00008)	
Pct Hispanic	.0011* (.0002)	
Pct Asian	-.00016 (.0002)	
Age 25	.141* (.0052)	
Age 34	.088* (.0043)	
Age 45	.046* (.0040)	
Age 55	.0192* (.0043)	
Age 65	.0021 (.0048)	
Income	-2.0e-6* (2.6e-7)	
Income^2	7.9e-12* (1.6e-12)	
Pct College	-.00042* (.00017)	
Pct less HS	-.00032 (.00021)	
Med house value	-5.5e-7* (2.7e-8)	
Pct executive	-.0011* (.00026)	
Pct blue collar	.00095* (.00017)	
Pseudo r-squared	.044	
Observations	245110	

* indicates significant at the 5% level or better. Controls included but not reported: region dummies, end of month, end of year, weekend, professional, technical, percent home ownership

Table 4: Car and Demographic determinants of *normdP*		
Variable	(1) tobit	(2) tobit
Female	.017* (.002)	.015* (.002)
Pct Black	.0014* (.00009)	.0014* (.00009)
Pct Hispanic	.0011* (.0002)	.0010* (.0002)
Pct Asian	-.0001 (.0002)	-.0002 (.0002)
Age 25	.141* (.005)	.195* (.009)
Age 34	.087* (.004)	.104* (.005)
Age 45	.046* (.004)	.051* (.004)
Age 55	.019* (.004)	.025* (.004)
Age 65	.002 (.005)	.006 (.005)
Income	-2.1e-06* (2.6e-07)	-2.1e-06* (2.6e-07)
Income^2	7.9e-12* (1.6e-12)	7.8e-12* (1.6e-12)
Pct College	-.0004* (.0002)	-.0004* (.0002)
Pct less HS	-.0003 (.0002)	-.0003 (.0002)
Med house value	-5.5e-07* (2.7e-08)	-5.5e-07* (2.7e-08)
Pct executive	-.0011* (.0003)	-.0011* (.0003)
Pct blue collar	.0009* (.0002)	.0010* (.0002)
First car	---	-.0444* (.0068)
TradeAge03	.0209 (.004)	.0199* (.004)
TradeAge47	.0338 (.004)	.030* .004)
TradeAge810	.0322 (.004)	.037* (.004)
Fun	----	-.052* (.003)
Luxury	----	-.0018

		(.0028)
Pseudo r-squared	.045	.045
Observations	256139	256139

* indicates significant at the 5% level of better. Controls included but not reported: region dummies, end of month, end of year, weekend, professional, technical, percent home ownership

Table 5: Different samples: determinants of delta P		
Variable	(1) Internet only	(2) no dealer financing
Female	.011 (.012)	.0105* (.0045)
Pct Black	.0015* (.0006)	.0008* (.0002)
Pct Hispanic	-.0007 (.0010)	-.00015 (.0004)
Pct Asian	-.0002 (.0010)	.00012 (.00045)
Age 25	.098* (.033)	.067* (.011)
Age 35	.025 (.028)	-.019* (.008)
Age 45	-.0086 (.028)	-.056* (.006)
Age 55	-.033 (.029)	-.065* (.007)
Age 65	-.040 (.033)	-.036* (.007)
Income	-2.5e-06 (1.3e-06)	-2.2e-06* (4.7e-07)
Income^2	1.9e-11 (7.7e-12)	7.7e-12* (2.8e-12)
Pct College	-.0008 (.0008)	.00012 (.00032)
Pct less HS	-.0017 (.0012)	-.00046 (.00043)
Med house value	-7.1e-7* (1.2e-7)	-5.2e-07* (5.1e-08)
Pct executive	-.0001 (.0013)	-.00089 (.00049)
Pct blue collar	.0005 (.0009)	.0012* (.0003)
TradeAge03	.093* (.020)	-.092* (.007)
TradeAge47	.071* (.019)	-.033* (.007)
TradeAge810	.013 (.022)	.033* (.008)
r-squared		
Observations	12,277	69,932

* indicates significant at the 5% level of better. Controls included but not reported: region dummies, end of month, end of year, weekend, professional, technical, percent home ownership

Table 6: Total profit earned by dealer				
Unconditional correlations				
	Total price	normdP	-.19	Trade-ins only
	Adjusted price	normdP	-.18	Trade-ins only
	Total price	normdP	-.09	Full sample
	Adjusted price	normdP	-.08	Full sample
Simple regressions: full sample				
	Total price	Adjusted price		
NormdP	-.198* (.002)	-.184* (.002)		
Trade-in indicator	.047* (.0009)	.049* (.0009)		
Observations	693,082	647,272		
R-squared	.013	.012		
Regressions with demographic controls and car fixed effects: full sample				
Variable	Total price	Adjusted price		
normdP	-.011* (.0003)	-.142* (.002)		
Trade-in indicator	.0026* (.0001)	.058* (.0009)		
Observations	642805	642805		
R-squared	.977	.104		

* indicates significant at the 5% level or better. Lowest panel includes all demographics and car fixed effects, month indicators, region indicators, and transaction characteristics.