CONSUMER BELIEFS AND BUYER AND SELLER BEHAVIOR IN THE VEHICLE INSPECTION MARKET

Thomas N. Hubbard

Working Paper **6245**

NBER WORKING PAPER SERIES

CONSUMER BELIEFS AND BUYER AND SELLER BEHAVIOR IN THE VEHICLE INSPECTION MARKET

Thomas N. Hubbard

Working Paper 6245
http://www.nber.org/papers/w6245

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 1997

I would like to thank Tim Bresnahan, Andrew Dick, Kevin Murdock, Roger Noll, Robert Porter, Jean-Laurent Rosenthal, Andrea Shepard, Scott Stern, Eric Talley, Darrell Williams, Frank Wolak, seminar participants at AEA, NBER, UCLA, the University of Illinois, and the Federal Trade Commission, and two anonymous referees for helpful comments. Financial support from the Lynde and Harry Bradley Foundation, and an Alfred P. Sloan Doctoral Dissertation Fellowship are gratefully acknowledged. This paper is part of NBER's research program in Industrial Organization. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

Consumer Beliefs and Buyer and Seller Behavior
in the Vehicle Inspection Market
Thomas N. Hubbard
NBER Working Paper No. 6245
October 1997
JEL Nos. L14, L15, D83
Industrial Organization

ABSTRACT

Moral hazard exists in "diagnosis-cure" markets because sellers have an incentive to shade their reports of buyers' condition to increase the short-run demand for the treatments they supply. The California vehicle emission inspection market offers a rare opportunity to examine how incentives operate in such markets. This paper investigates why sellers help vehicles pass inspections, focusing on multiperiod mechanisms such as those in reputation models. I show that the demand individual firms face is sensitive to inspection outcomes. Consumers are 30% more likely to return to a firm at which they previously passed than one at which they previously failed. If, over the long run, an independent garage fails one additional vehicle per month, this decreases demand by 5.6 inspections per month on the average. This figure is lower for service stations and new car dealers. Consumers' behavior is consistent with a learning model in which they have diffuse initial priors regarding the probability they fail at individual firms, and Bayesian update using two to three inspection outcomes at each firm.

Thomas N. Hubbard
Department of Economics
University of California, Los Angeles
Los Angeles, CA 90095
and NBER
hubbard@econ.ucla.edu

1. Introduction

Transactions involving services are not simple exchanges. Production takes place after buyers and sellers agree to the terms of trade. When buyers can neither perfectly observe nor costlessly verify quality, moral hazard problems arise. Sellers can take actions which affect the size and allocation of the gains from trade. This situation appears in transactions mediated by markets as well as those that take place "within firms." In "diagnosis-cure" markets such as health care, automotive repair, and legal services, sellers supply information and related services. Doctors, mechanics, and lawyers have an incentive to shade their reports of buyers' condition to increase demand for the services they supply.

The California vehicle emission inspection market offers a rare opportunity to investigate how incentives operate in a "diagnosis-cure" market, due to the availability of detailed transaction-level panel data. In previous work, Hubbard (1997) finds that inspectors generally help vehicles pass inspections, even though they could increase short-run demand for emission-related repairs by helping them fail. Inspectors' actions cut the fraction of vehicles which fail from about 40% to about 20%.

This paper investigates the incentives which encourage inspectors and their firms to help vehicles pass. The focus is on multiperiod mechanisms such as those in reputation models.[1] There are two research goals. One is simply to provide evidence regarding why inspectors tend to help vehicles pass. I estimate a model of consumer choice, concentrating on two empirical relationships. One is how the outcome of their vehicle's previous inspection affects the probability consumers choose the same firm in the future. The other is how firms' failure rates — the fraction of inspections they perform which are failures — are related to individual consumers' choice. Empirical estimates of these relationships allow one to compute the elasticity of firms' demand with respect to inspection outcomes. If demand is sensitive to inspection outcomes, this is evidence that buyers provide sellers incentives via multiperiod mechanisms.

The other goal is to better understand what underlies consumers' behavior, and thus the

[1]Single-period contingent contracts are uncommon in this market. Part of the reason for this is that they are discouraged by the environmental regulators who oversee the market. Adverse selection may also undermine such arrangements. Consumers may have private information about their vehicle's condition.

mechanism through which incentives arise in this market. There several distinct theoretical models in which multiperiod mechanisms align incentives.[2] I examine the empirical results in light of one of these models. In this working model, relationships between individual consumers' choice and inspection outcomes arise solely because of learning: consumers believe inspection outcomes provide information about firm characteristics they do not directly observe. Under these conditions, I show that estimates of the two central relationships allow one to compute a measure which indicates how informed consumers are. Assuming that consumers have initial priors regarding the probability they fail at individual firms which are characterized by a Beta distribution, and use Bayes' rule to update from these priors, I show that the two relationships also characterize consumers' initial priors. Applying these results, the empirical estimates of the consumer choice model characterize consumers' beliefs and informedness. Interpreting the results with this model provides further intuition for what may drive buyers' (and thus sellers') behavior in this market.

I find that consumers' choice of firms is strongly related to both the outcome of their previous inspection and firms' failure rates. If a firm passes, rather than fails, a consumer's vehicle, this increases the probability the consumer chooses the firm the next time his vehicle needs an inspection by about 30% on the average. From the relationship between choice and failure rates, in the long run, failing one additional vehicle per month would decrease demand by 5.6 inspections per month on the average across the independent garages in my sample, 2.4 on the average across the service stations, and 1.7 across the new car dealers. Multiperiod mechanisms encourage sellers to supply passing inspections. Although consumers cannot completely observe inspectors' actions, they provide sellers strong incentives to help them pass because demand is quite sensitive to inspection outcomes. Assuming that the data reflect a "Nash-in-prices" equilibrium, failing one additional vehicle per month implies a loss of \$46.71 in inspection profits per month. These demand shifts do not appear to be due to consumers' learning about their vehicle. If they solely reflect consumers'

---

[2]There are two classes of "reputational" models in the literature. In one class, agents have complete information about each others' attributes. Equilibrium behavior is maintained by participants' commitment to punish parties when they observe something that signals a deviation from equilibrium strategies. (Klein and Leffler (1980), Green and Porter (1984), Abreu, Pearce, and Stachetti (1986)). The data cannot reject this class of models; the empirical results may reflect a very complicated trigger equilibrium. In another set of models, agents have incomplete information about each others' attributes ("types"). Agents' decisions reflect their desire to signal they are certain types. (Fama (1980) and Holmstrom (1982) explore this in the context of managerial shirking.) The model used to interpret the data falls in this second class.

learning about unobserved firm characteristics, they are consistent with a model in which all consumers know the outcomes of two to three inspections at each firm in the market, and have diffuse initial priors regarding the probability they fail at individual firms. Buyers' and sellers' behavior in this market may reflect that consumers are skeptical regarding the extent individual firms are prone to help them pass and apply information from a very small proportion of transactions in the market in forming their conjectures.

An outline of the rest of the paper follows. Section 2 describes the relevant features of the inspection market, discusses the different models in which multiperiod mechanisms can supply firms quality incentives, and derives the learning model I use to interpret the results. In section 3, I describe the data. In section 4, I construct the empirical framework and develop the econometric model used in estimation. Section 5 contains the estimation results and analysis. Section 6 concludes.

2. Analytic Framework

The Market for California Vehicle Inspections

In most parts of California, drivers must obtain an emission certificate each time they change their vehicle's registration and biennially upon registration renewal.[3] In general consumers can only obtain a certificate once their vehicle passes an emission inspection.[4] Inspections and any associated emission-related repairs have little or no private value. Consumers prefer passing inspections to failing them because "passes" relieve them of a regulatory requirement that is costly to fulfill.

Private firms such as independent garages, service stations, and new car dealers supply emission inspections. Inspections have two parts: an "emission test" in which inspectors measure the composition of vehicles' exhaust, and an "underhood test" in which they check the physical condition of emission control equipment. Vehicles pass inspections when they pass both parts. Inspectors employed by these firms conduct inspections and complete emission-related repairs. These individuals have discretion in how to conduct inspections, and if the vehicle fails, which

---

[3]New vehicles and those older than the 1966 model year are exempt, as are those with diesel engines. For more detail about the California vehicle inspection market, see Hubbard (1996, 1997).

[4]Vehicles can obtain a waiver if the cost of the repairs required to satisfy applicable standards exceeds a model-year-specific amount. During 1992, the period of my data set, this ranged from $50 to $350.

repairs to recommend. They can affect inspection results in several ways. They can influence tailpipe emission readings by warming vehicles up. They can influence the outcome of the underhood test by simply being more or less lenient in applying the relevant technological standards. Actions that affect the probability vehicles fail or the cost of repairs given a failure affect consumers' cost of registering their vehicle. Moral hazard exists when consumers can neither directly observe nor costlessly verify the effect of these actions.

Regulators oversee the inspection market. They prefer that inspection outcomes be determined by vehicles' actual emission condition, not by actions taken by inspectors that make vehicles' emission condition seem different than it actually is.[5] They attempt to limit how inspectors affect inspection results in two ways. First, as much as possible, they control the inspection procedure with software routines embedded in inspectors' emission analyzers.[6] Second, they conduct covert audits. In these, undercover state officials bring a vehicle designed to fail an inspection to an inspection supplier. If it passes without preinspection repairs, the inspector and the firm are given citations.

Below is a stylized description of the decisions buyers and sellers make and the incentives that affect them. This highlights why buyers' beliefs regarding how inspectors will behave may differ across firms and vary with previous inspection outcomes. This motivates the specification of the empirical model and interpretation of results.

The timing of events follows. Consumers first form beliefs about the cost of obtaining a passing inspection at different firms. They then choose a firm. An inspector at the firm they select then chooses how to conduct the inspection. Because emissions are stochastic, nature then moves; this determines the inspection outcome.[7] The next period then begins. Consumers next choose among firms when they next need an inspection. If the outcome was a "fail," this is soon after the

---

[5]Regulators are not averse to all of the ways in which inspectors can exercise discretion. For example, they permit (and even encourage) inspectors to warm up vehicles before inspecting them to the extent that it makes vehicles' condition during the inspection more representative of their condition in normal operation.

[6]For example, the machines can determine whether the probe that measures tailpipe emissions is in a vehicle's tailpipe.

[7]Throughout this paper, I assume that the support of the distribution of outcomes is the same for all actions.

initial inspection, often after consumers purchase repairs. If it was a "pass," it is the next time they need to change or renew the vehicle's registration. This paper examines consumers' choice of firms for their first inspection within an "inspection cycle" — not their choice of where to obtain repairs or reinspections.

Firms choose their organizational characteristics, the goods and services they supply, and prices toward maximizing profits. Organizational characteristics include hierarchies and compensation schemes. I take these to be exogenous with respect to individual inspectors' and consumers' decisions. Inspectors choose how to conduct inspections to maximize their utility, which is a function of income and effort. Firms' characteristics imply incentive structures which affect how inspectors behave. At most firms, part of inspectors' and mechanics' compensation is based on piece rates. Inspectors have an incentive to help vehicles fail because their firms have local market power in supplying emission-related repairs. If they believe demand is sensitive to inspection results, they face a trade-off between helping vehicles fail and helping them pass.

Consumers choose among firms to maximize expected utility.[8] For many, this is approximately equivalent to minimizing the cost of obtaining a passing inspection. Some, however, may have preferences for particular firms — for example, their vehicle's new car dealer — unrelated to cost. The cost of obtaining a passing inspection includes the inspection price and time and travel costs. It also includes all costs associated with failing an inspection. I will refer to these as "repair costs," although they include the price of reinspections and time costs as well as repair prices. These costs equal zero when vehicles pass, and are positive when they fail. Consumers are uncertain about repair costs because they cannot perfectly determine their vehicles' emission condition (or forecast what it will be during the inspection), and may not be able to perfectly anticipate how inspectors will exercise their discretion. Given inspectors' actions, expected repair costs may be higher for older vehicles, at firms which do not offer free reinspections, and for consumers who place a relatively high value on their time. Consumers' expectation of inspectors' actions depends critically on their beliefs and information about firms' characteristics, and may be related to previous inspection outcomes.

---

[8]I will assume consumers maximize current period expected utility. Consumers may value the information they receive about firms while transacting with them, but the expected value of this information is the same across firms.

supply. The other is inspection outcomes. Inspection outcomes are not public information. The inspection outcomes an individual "observes" include both those from his own experiences and those about which he hears from others.

Let $g(n, N, \theta)$ represent an individual's conjecture a firm at which he has observed N inspections, of which n are failures. θ is a vector of parameters. Assume $g(n, N, \theta)$ is linear in n:

$$g(n,\ N,\ \theta)\ =\ w_0(N,\theta)\ +\ w_1(N,\theta)\ n \tag{1}$$

One class of functions which satisfy this take the form:

$$g(.)\ =\ [1 - w(N,\theta)]\,k_0(\theta)\ +\ w(N,\theta)\frac{n}{N} \tag{2}$$

In this case, the conjecture is a weighted average of $k_0(\theta)$ and the fraction of inspections the individual observes at the firm which are failures. $k_0(\theta)$ is the individual's conjecture, given only public information. If $w(N,\theta)=1$, $k_0(\theta)$ is irrelevant and g(.) takes the form of a classical maximum likelihood estimate. Alternatively, if one sets $k_0(\theta) = \alpha/(\alpha+\beta)$ and $w(N,\theta) = N/(\alpha+\beta+N)$, one obtains:

$$g(.)\ =\ \frac{\alpha + n}{\alpha + \beta + N}\ =\ \frac{\alpha}{\alpha + \beta + N}\ +\ \frac{1}{\alpha + \beta + N}\ n \tag{3}$$

α/β defines $k_0(\theta)$. $(\alpha + \beta)$ determines the weight the individual places on $k_0(\theta)$, given $N>0$.

Mathematically, this expression is equal to the expectation of a random variable distributed $B(\alpha + n, \beta + N - n)$, where B denotes the Beta distribution and α and β are parameters. One well-known result from conjugate distribution theory is that $B(\alpha + n, \beta + N - n)$ is the posterior distribution one obtains by starting with initial priors $B(\alpha, \beta)$ regarding the unknown mean of a binomial distribution, and Bayesian updating using N random draws from the distribution, n of which are ones.[9] This special case is of particular interest because it shows (sufficient) conditions in which conjectures which are linear in n arise from a model in which consumers Bayesian update. In the present context, the conditions are that an individual's initial priors regarding the probability

---

[9]Degroot (1970). The Beta distribution allows for a wide range of unimodal and J-shaped densities which have a finite support. The uniform is a special case. See the Appendix for more about its properties.

Suppose consumers have complete information about firms' characteristics and how they affect inspectors' behavior. Expected repair costs may be related to previous inspection outcomes because consumers anticipate that inspectors behave differently according to whether they previously passed or failed. This would be the case in trigger equilibria. What may help maintain such an equilibrium is that inspectors may not observe certain consumer characteristics which affect their preferences among firms — such as where they live or work. Inspectors may draw inferences about these from how individual consumers respond to previous transaction outcomes and discriminate accordingly. In such a model, some consumers may not use simple loyalty-boycott strategies, but inspectors discriminate against those who (optimally) return after failing. The empirical framework and data can be used to test whether all consumers and firms are using certain simple strategies. However, one cannot reject all models in which demand shifts occur because consumers believe firms discriminate according to previous outcomes. The results may reflect a very complicated trigger equilibrium.

Suppose instead that consumers have incomplete information about firms' characteristics. Then expected repair costs may be related to previous inspection outcomes because consumers use them to infer which firms have unobserved characteristics that make them "consumer-friendly." The following subsection develops a model in which relationships between transaction outcomes and expected repair costs reflect only changes and differences in consumers' beliefs about firms' unobserved attributes. Under such circumstances, a learning model alone forms the structure through which transaction outcomes affect expected repair costs. Consumers do not believe inspectors discriminate. One can then interpret consumers' observed behavior in terms of their initial beliefs about firms' unobserved characteristics and how much information they use in updating these beliefs. While I discuss this model in the context of the inspection market, it is applicable for analyzing consumer behavior in other markets where consumers are uncertain about product quality, and in which researchers have transaction-level panel data.

A Learning Model

Suppose each consumer forms a point estimate (a "conjecture") regarding the probability of failing an inspection at each firm. Assume conjectures reflect two sets of information. One of these includes variables all consumers directly observe such as inspection prices and the services firms of failing an inspection at a firm have a B(α, β) distribution, and that the information — inspection outcomes at that firm — he uses in updating these priors takes the form of binomial random variables drawn from a common distribution. Note that "initial priors" here represent individuals' beliefs, given public information. Assuming that g(.) takes the "Beta-Bayes" form, w_1 is a function of the number of inspection outcomes used by the individual in forming his conjecture and parameters which characterize his initial prior regarding the probability of failing at the firm.

One quantity of interest is how much an individual's conjecture changes with a single outcome. For an individual consumer using the above learning rule, the difference in his conjectures is:

$$g(n+1,N,\theta) - g(n,N,\theta) = w_1(N,\theta) \tag{4}$$

Another quantity of interest is how much an individual's conjecture about a firm differs with the firm's failure rate across all individuals. One can use g(.) to calculate this as well. Let F_j be a firm's failure rate across all individuals. Let $E_n(g(n(F_j,N), N, \theta,))$ be a consumer's expected conjecture about a firm with actual failure rate F_j, and at which he has N observations. Assume that these observations are random draws from the set of all inspections at the firm. The expectation is taken over n, the number of failures.

$$h(N,\theta) = \frac{dE_n(g|N,\theta,F_j)}{dF_j} = Nw_1(N,\theta) \tag{5}$$

If one obtains these two quantities for an individual, one obtains N from their ratio. One can determine how many transaction outcomes the individual has observed at that firm if g(.) is linear in n and the inspection outcomes the individual observes are random draws from all transactions at the firm. Assuming that individuals observe random draws from a common distribution implies that the vehicles whose inspection outcomes they observe are not systematically different than the population of vehicles the firm inspects. It also implies that consumers believe that inspectors and firms neither behave differently toward individual consumers from inspection to inspection nor discriminate across individuals.

Note that g(.) is not a function of inspection results at firms other than the one at hand. This rules out situations where consumers use inspection outcomes to infer unobserved characteristics of their vehicle which persist over time and affect the probability they fail. It is possible to specify a g(.) that would accommodate this, but it would be difficult to obtain an expression for h(.) when g(.) is a function of inspection outcomes at the firm at hand and at other firms. One reason is that one would have to take an expectation over outcomes at other firms, and it is unclear how to do so without knowing at which firms each of these outcomes occurred.[10] Another is that conjectures at one firm would generally be a function of conjectures at other firms. The empirical framework and data provide a means of testing whether g(.) reflects consumers' learning about their vehicle. I describe this in section 4.

Hypothesis Tests

Expected repair costs are equal to the product of R_{ij}, the individual's expected cost of repairs given that he fails an inspection at firm j, and $g(n_{ij}, N_{ij}, \theta_{ij})$.

$$\begin{aligned} E(R_{ij}) &= R_{ij} g(n_{ij}, N_{ij}, \theta_{ij}) \\ &= R_{ij}(w_0(N_{ij}, \theta_{ij}) + w_1(N_{ij}, \theta_{ij}) n_{ij}) \end{aligned} \qquad \textbf{(6)}$$

For ease of exposition, assume that consumers are not uncertain about R_{ij}: unobserved firm characteristics only affect the probability they fail, not the cost of repairs given they fail.[11]

Suppose w_1=0 for all individuals at all firms. This would hold if, for example, all consumers believe that they are fully informed about each firm in the market: either because their initial priors have zero variance or they have enough information through transaction outcomes in the market so that the outcome of a single transaction outcome means little. Then $g(n_{ij}+1, N_{ij}, \theta_{ij})-g(n_{ij}, N_{ij}, \theta_{ij}) = 0$ and $h(N_{ij}, \theta_{ij}) = 0$ for each individual and at each firm: individuals' choice will be unrelated to n_{ij}

[10]For example, h(.) may be a function of failure rates at other firms, but which other firms?

[11] Assuming that consumers are only uncertain about failure probabilities has little qualitative impact because repairs and inspections are effectively bundled in equilibrium — a very small proportion of consumers change firms after failing their initial inspection. One could construct a learning model based on conjectures about failure probabilities and repair costs given failure (or, equivalently, expected repair costs), but the main analytical points would not change and the exposition would become somewhat more complicated.

and F_j. Finding that either or both of these quantities does not equal zero implies that at least some consumers believe that they are not fully informed about each firm in the market.

Suppose N=0 for all consumers at all firms in the market: no consumer has at hand any information about specific transaction outcomes at any firm in the market. Then $h(N_{ij}, \theta_{ij}) = 0$: conjectures will be unrelated to failure rates. If one finds that conjectures are in fact related to failure rates, then one can reject the hypothesis that consumers have no information about any firm in the market. At least some consumers are informed via inspection outcomes. Therefore, finding that neither $g(n_{ij}+1, N_{ij}, \theta_{ij})$-$g(n_{ij}, N_{ij}, \theta_{ij}) = 0$ nor $h(N_j, \theta_j) = 0$ implies that some consumers a) believe that they are incompletely informed about firms' characteristics, b) believe that transaction outcomes are informative, and c) have accumulated some information about firms via transaction outcomes. If one can empirically estimate these quantities, one can test the null hypotheses that consumers (or subsets of consumers) are completely informed or completely uninformed.

Characterizing Consumers' Informedness and Initial Priors

The hypothesis tests allow one to examine very extreme hypotheses about the extent to which consumers are informed. One might expect these hypotheses to be rejected. An alternative hypothesis is that consumers are to some extent, but not completely, informed. One would like to take estimates of $g(n_{ij}+1, N_{ij}, \theta_{ij})$-$g(n_{ij}, N_{ij}, \theta_{ij})$ and $h(N_j, \theta_j)$ and solve for N_j and θ_j. These quantities would indicate the extent to which each consumer is informed via inspection outcomes and, under the Beta-Bayes assumptions above, characterize each individual's initial priors. The individual estimates could then be used to characterize the distribution of informedness and initial priors across consumers in the market. It is easy to specify an empirical model of individuals' choice of firms, conditional on n_{ij} and N_{ij}, (their "history" at firm j). However, individuals' histories — which may include the outcome of other individuals' transactions — are at best partially observed. One could account for this by treating an individual's history as an unobserved variable, parameterizing the distribution of information in the market, and integrating out over the set of possible histories. But such a model would carry a huge computational burden because the number of possible histories is generally very large.[12]

This paper takes a simpler approach toward gleaning information about consumers' informedness and initial priors from the data. Assume $N_{ij} = N$ and $\theta_{ij} = \theta$. This requires that the amount of information individuals have about firms is the same across individuals and firms, and under the Beta-Bayes assumptions, individuals' initial priors are such that $(\alpha + \beta)$ is constant across individuals and firms.[13] Under these symmetry assumptions, one can solve for $(\alpha + \beta)$ and N directly from the estimates of $g(n+1, N, \theta)-g(n, N, \theta)$ and $h(N, \theta)$. Estimates of $(\alpha + \beta)$ and N obtained through this method provide metrics which characterize consumers' initial priors and informedness. The qualitative conclusions drawn from these metrics may extend to situations where the symmetry assumptions are relaxed.

3. Data

The data are similar to those used in Hubbard (1997). They include 7519 observations of vehicles which received their initial inspections in Fresno, California between late August and mid-November, 1992. This is the set of all individuals who obtained their initial inspections during this time at one of twenty-nine firms in the north part of the city.[14] This cluster of firms is located in a dense, commercially-zoned corridor which is approximately 3 miles by 1 mile. Most of the firms are on North Blackstone Road, an extremely busy multilane road. The region's boundaries are chosen so that all firms have a competitor within one-half mile that is also within the region, and no firm has a competitor within one-half mile that is outside of the region. I examine demand at a cluster of firms rather than the entire city to make the empirical work more tractable. The results of the demand model estimated below are conditional on consumers' selecting to purchase an inspection from a firm in the cluster.

---

[12]This is true even if one assumes that N is known and constant across individuals unless N is very small and there are few firms in the choice set: if an individual has observed N inspections at each of K firms in the market, there are $(N+1)^K$ possible histories.

[13]Note that different individuals can have different information about firms (but not different amounts). Individuals' initial priors can vary across firms as long as $(\alpha + \beta)$ is constant. For example, their expected probability of failure can vary with variables they observe directly such as price or other firm characteristics. Finally, note that initial priors need not be consistent in the game-theoretic sense.

[14]"Initial" means that they are the vehicles' first inspections within the period. The cluster of firms I examine comprised about 30% of the inspection suppliers in the city. These firms supplied about 30% of the inspections.

Each observation includes firm and vehicle characteristics, and inspection results. There is no information about consumer characteristics other than the characteristics of their vehicle and where they purchase inspections. I obtained the inspection price at each firm in a telephone survey. I calculate failure rates over the entire August-November 1992 period for each firm. While this perfectly measures the true failure rate during this period (I have all observations at the firms in my sample), it is an imperfect measure if consumers use information from transactions outside this period. I do not have data from immediately before August 1992. If there is sampling error, the empirical model presented below is poorly specified. Fortunately, even if consumers' information is based on periods longer than the time from which my sample is drawn, there are many observations at most of the firms. If inspection policies at firms are constant over time, it is reasonable to assume that the failure rate between August and November 1992 is very close to that defined over longer periods from which consumers may observe transactions.

I acquire information about consumers' previous transactions by using inspection data from the entire state of California between August and November 1990, when many of the vehicles in my sample were receiving their previous inspections. Using vehicle identification numbers, I am able to match about one-third of the 1992 observations to 1990 observations.[15] Because inspections are required when vehicles change owners, and "change of ownership" inspections shift vehicles' inspection cycles so that their next scheduled inspection is two years after the ownership change, very few of the vehicles inspected during both August-November 1990 and August-November 1992 were owned by different individuals at these times.[16] This helps in two ways. First, it allows me to interpret "same vehicle" as "same consumer" or "same household" when I am able to match 1990 and 1992 observations. Similarly, it allows me to interpret cases where vehicles were previously inspected outside Fresno county — new-to-market vehicles — as new-to-market consumers. If a vehicle inspected outside of Fresno county during August-November 1990 were sold to an individual

[15] I conjecture that, of the 1992 observations I was not able to match, 35-40% are because they were receiving off-cycle "change-of-ownership" inspections, 20-25% are new vehicles, and 5-10% are vehicles that were previously registered in another state. The remaining non-matches are vehicles for which the VIN was misentered by the inspector, and those receiving biennial inspections during August-November 1992 whose previous inspections happened to miss the August-November 1990 window.

[16] The exceptions to this are when changes of ownership happened to occur very close to the same time vehicles would have been otherwise due for inspections.

living in the county between 1990 and 1992, the vehicle would have been inspected at that time, then not for another two years. I generally would not observe these vehicles being inspected during August-November 1992.

Table 1 provides a first look at the data. Of the 7519 vehicles observed in 1992, I was able to identify 1990 inspections for 2704, or 36%. 263 of these 2704 were at firms outside of Fresno county. Of the 2441 which were observed in 1990 in Fresno county, 391, or 16%, failed the 1990 inspection. Of those that passed, 38.8% chose the same firm in 1992; of those that failed, 25.3% did. Of the 1286 which were observed in 1990 at a firm within the 29-firm cluster, 13.8% failed the 1990 inspection. Of those that passed, 71.7% chose the same firm in 1992; of those that failed, 55.9% did. These proportions are higher for the "old to cluster" than the "old to market" subsamples because, by definition, "new to cluster" consumers did not choose the same firm in 1990 and 1992.[17] These raw numbers indicate that consumers are substantially more likely to return to firms at which they previously passed than those at which they previously failed. However, they are by no means certain to return after passing, nor are they certain not to return after failing. This is evidence against the simplest equilibria in which homogeneous consumers discipline firms by following simple loyalty-boycott rules or strategies. Consumers are about equally likely to return conditional on failing either part of the test, but they are more likely to return if they failed either part than both.

Inspection outcomes are correlated across time. Of the 2704 vehicles observed in both years, 41.8% of those which failed in 1990 failed in 1992, but only 16.0% of those which passed in 1990 did. Part of this is due to differences in the vehicles' age and make. Table 2 reports results from four simple logits. The dependent variable equals one if the vehicle failed its 1992 inspection, and zero otherwise. "Fail in 1990" equals one if the vehicle failed its 1990 inspection, and zero otherwise. The second and third columns add a full set of vehicle age and make dummies. In each specification, the fail in 1990 dummy is positive and significant. The probability deltas at the bottom of the table report differences in the estimated probability a vehicle failed its 1992 inspection when "fail in 1990" equals one and zero, holding the other independent variables at their sample means.

[17]They are also conditional on choosing a firm within the cluster in August-November 1992. The sample does not include vehicles inspected at a firm within the cluster in 1990, but elsewhere in 1992. 71.7% and 55.9% thus overstate the proportions that chose the same firm across all individuals who obtained inspections at these firms during August-November 1990.

Including the age and make dummies cuts this figure by more than half, but it is still 12-14 percentage points. This provides evidence that vehicle characteristics other than age and make influence inspection results in a way that persists from year to year. If consumers do not directly observe these characteristics, they may use inspection outcomes toward drawing inferences about vehicle condition. The fourth column includes a full set of firm dummies, a dummy variable which equals one if the vehicle was inspected at the same firm as in 1990 and zero otherwise, and an interaction between "same firm" and "fail in 1990." The results suggest that inspectors helped consumers who had purchased inspections at their firm in 1990 pass more than those who had obtained them elsewhere, regardless of whether the vehicle passed in 1990. Firms did not "punish" consumers who returned after failing.

Table 3 contains the inspection price, number of observations, share of observations, failure rate, and "station type" for each firm in my sample. Prices range from $19.76 to $65.00. The average inspection price over firms is $39.32; the average price over inspections is ten dollars lower, because more inspections take place at lower-price firms. Over half of the observations are at only three of the twenty-nine firms in my sample. Failure rates range from 2.8% to 33.3%. Thirteen of the firms are new car dealers, eight are independent garages, seven are service stations, and one is a tune up shop.[18] The tune up shop has by far the largest market share, completing more than 25% of the inspections of these firms. Failure rates are positively correlated with market share. This is probably due to the fact that most of those with low failure rates are new car dealers. These firms tend to have the highest inspection prices and labor rates. Furthermore, low failure rates may not indicate that inspectors at these firms generally help vehicles pass, because the vehicles they inspect tend to be newer and lower-emitting.

4. Empirical Framework

Specification of Demand

Assume that consumers choose among firms to maximize utility in each period. Let V_{ij} be consumer i's indirect utility from choosing firm j. Divide indirect utility into cost- and non-cost-related components:

[18]The high percentage of new car dealers is due to the fact that the cluster of firms includes an "auto row."

$$V_{ij} = \alpha C_{ij} + \mu_{ij} \quad \textbf{(7)}$$

C_{ij} is consumer i's expected cost of obtaining a passing inspection, given that he or she chooses firm j for the vehicle's initial inspection. μ_{ij} captures consumer i's idiosyncratic taste for the quality of service firm j provides. I specify μ_{ij} as:

$$\mu_{ij} = \gamma(OD)_{ij} + \eta_{ij} \quad \textbf{(8)}$$

where OD_{ij} ("own dealer") equals one if firm j is a new car dealer that sells consumer i's brand of vehicle and zero otherwise. I permit η_{ij} to be correlated among firms within station types; this accounts for the possibility that consumers may have non-cost-related tastes for the service at new car dealers, independent garages, etc.

I specify C_{ij} as:

$$C_{ij} = p_j + E(R_{ij}) + \varepsilon_{ij} = p_j + R_{ij}g(n_{ij}(F_j),N_{ij},\theta_{ij}) + \varepsilon_{ij} \quad \textbf{(9)}$$

The cost of obtaining a passing inspection at firm j is equal to the price of the initial inspection, "expected repair costs," and time and transportation costs.

I specify expected repair costs, $E(R_{ij})$, as a reduced form. In the basic specification, it is:

$$E(R_{ij}) = f(X_{vi},X_{ci},X_{oi}) + \beta_1 D_{ij}^1 + \beta_2 D_{ij}^2 + \beta_3 F_j + \beta_4 W_{ij} + \beta_5 Z_j + \beta_6 (Z_j X_{vi}) + \xi_j + \zeta_{ij} \quad \textbf{(10)}$$

where:

— $f(X_{vi}, X_{ci}, X_{oi})$ is an arbitrary function of vehicle and consumer characteristics, and the vehicle's previous inspection outcomes,

— D_{ij}^1 is a dummy that equals one if consumer i was observed to obtain a previous inspection at station j, and zero otherwise,

— D_{ij}^2 is a dummy that equals one if $D_{ij}^1=1$ and the consumer passed the previous inspection, and zero otherwise,

— F_j is firm j's failure rate across all consumers,

— W_{ij} is a dummy that equals one if a warranty applies for emission-related work for vehicle i at firm j, and zero otherwise,[19]

— Z_j is a vector of station type dummies,

— ξ_j is firm characteristics observable to the consumer but not the econometrician, and

— ζ_{ij} is an error term.

The error term includes sampling error and specification error.

Inspection outcomes affect expected repair costs at firms in general through the individual-specific term $f(X_{vi}, X_{ci}, X_{oi})$. This term is not empirically identified. β_2 indicates how much more single inspection outcomes affect expected repair costs at the firm where the inspection takes place. β_3 indicates how much expected repair costs differ with differences in firms' failure rates across all consumers. Along with the basic specification, I estimate specifications in which these vary with the age of consumers' vehicles and type of firm, and in which β_2 varies with firms' failure rates.

Under these assumptions, consumer i chooses firm j iff:

$$V_{ij} > V_{is} \ \forall \ s \neq j \Leftrightarrow \alpha(p_j + E(R_{ij}) + \varepsilon_{ij}) + \mu_{ij} > \alpha(p_s + E(R_{is}) + \varepsilon_{is}) + \mu_{is} \ \forall \ s \neq j \qquad \textbf{(11)}$$

Assuming that $\alpha(\xi_j + \zeta_{ij} + \varepsilon_{ij}) + \eta_{ij}$ is independent of the other right hand side variables, and has a generalized extreme value distribution, I then can estimate the model's parameters with a nested logit. This provides estimates of individual consumers' demand at each firm. One can then aggregate across consumers and obtain estimates of each firm's demand, how much it changes with individual transaction outcomes, and its elasticity with respect to failure rates.

Endogeneity Issues

Assuming that ξ_j is independent of the other explanatory variables brings up a familiar econometric issue.[20] In this model, ξ_j includes objective factors: for example, whether firms charge for reinspections, their labor rates, whether they can complete repairs "on the spot," whether

[19]Federal law requires vehicle manufacturers to provide 5-year, 50,000-mile warranties that cover emission-related repairs. I assume that these warranties only apply at a vehicle's "own dealer."

[20]See Berry (1994), Goldberg (1995).

consumers can leave their vehicles to be repaired, and so on.[21] It does not include subjective factors such as "consumer friendliness" which arise from firm characteristics consumers do not directly observe.[22] Because firms choose these characteristics simultaneously with inspection prices and their other characteristics, ξ_j may be correlated with included explanatory variables: most prominently prices and failure rates. If this is the case, then when estimating the model described below, the coefficients on these variables would pick up the effect of these unobserved characteristics as well. If in equilibrium firms with unobserved characteristics that consumers value more charge higher prices and have higher failure rates, then the price and failure rate coefficients will be biased toward zero. Alternatively, suppose that failure rates are lower at firms which offer free reinspections. Then β_3 would reflect both the effect of "free reinspection" on consumers' expected cost of repairs (both by itself and through its effect on inspectors' conduct) and the effect of the inferences consumers draw from observations of transactions not in the data.

To investigate this, I estimate the parameters of the model using a two-step procedure suggested by Berry (1994). The maintained assumption is that ξ_j, j=1,...,29, is mean independent of the observed (to the econometrician) characteristics of all firms in the market. In this procedure one first estimates product-level "mean utility levels" using individual data, then regresses predicted "mean utility levels" on firm characteristics using instrumental variables. I use the station type and number of inspectors of each firm's closest geographic competitor as instruments in the section stage. In what follows, I show that the estimates obtained using this procedure do not provide strong evidence that price and failure rate are econometrically endogenous, but this may be due to a lack of good instruments and small sample size. The sample size in the second stage is equal to the number of firms (in this case, only 29).

For the most part, I will discuss the estimates obtained under the assumption that ξ_j is independent of all included explanatory variables. This assumption is satisfied if consumers believe that, of the attributes they directly observe, only those in Z_j and W_{ij} influence expected repair costs.

---

[21]Some of these are picked up in the station type dummies — for example, the new car dealer dummy may pick up the effect of the higher average labor rates at new car dealers. These controls do not pick up differences among firms within station types.

[22]This is unlike other consumer choice models, in which reputations are considered product characteristics consumers directly observe.

If this holds, then the interpretations outlined above apply. If not, then the estimates demonstrate empirical phenomena that are still of interest with respect to sellers' incentives. They show whether, conditional on the prices and characteristics observed in the data, consumers believe that obtaining a passing inspection is more expensive at firms whose failure rate across all consumers is high. Part of the explanation would be that consumers know that firm characteristics they observe (but which are not included in the model) are associated with better conduct and lower costs. If this is the case, relationships between failure rates and consumers' choice still imply that consumers are presenting sellers incentives to treat them well. The market provides sellers incentives to adopt organizational characteristics that both are observed by consumers and are associated with low failure rates.

Implications of the Estimates in a Learning Model

Suppose that in the underlying theoretical model g(.) is linear in n. Suppose for now that $R_{ij}=R$, a positive number which is constant across consumers and firms. One implication of this assumption is that learning about one's vehicle affects expected repair costs equally across firms. Then β_2 is the empirical analog to $R[g(n_{ij}+1, N_{ij}, \theta_{ij})-g(n_{ij}, N_{ij}, \theta_{ij})] = Rw_1$; β_3 is the empirical analog to $Rh(N_{ij}, \theta_{ij}) = RNw_1$. H_0: $\beta_2 = 0$ corresponds to H_0: $w_1 = 0$ for all consumers at all firms. H_0: $\beta_3 = 0$ corresponds to H_0: $Nw_1 = 0$ for all consumers at all firms. Rejecting the joint hypothesis H_0: $\beta_2 = 0$, $\beta_3 = 0$ implies the alternative that some consumers believe that they are incompletely informed, and apply some information from transaction outcomes in forming their conjectures.

The ratio β_3/β_2 is the empirical analog to $h(N_{ij}, \theta_{ij})/(g(n_{ij}+1, N_{ij}, \theta_{ij})-g(n_{ij}, N_{ij}, \theta_{ij}))$. Under the symmetry assumptions outlined earlier this equals N, the number of inspections each individual has observed at each firm. If one assumes a value for R, one can also solve for $(\alpha + \beta)$ under the Beta-Bayes assumptions: this equals $(R-\beta_3)/\beta_2$. These characterize consumers' informedness and beliefs assuming that the observed behavior reflects only consumers' learning about firms.

One can capture differences in R_{ij} which correspond to vehicle or firm characteristics observed in the data in specifications with interaction effects. If learning about vehicles does not change consumers' preferences among firms, this simply allows β_2 and β_3 to differ across consumers or firms but does not change how one interprets them. Differences in R_{ij} across firms take on greater qualitative importance if consumers prefer certain types of firms — for example, those at which repairs are inexpensive — more if they believe that their vehicle is likely to fail. If learning about vehicle characteristics does not affect expected repair costs equally across firms, β_2 may reflect changes in preferences due to learning about vehicles as well as those due to updating about incumbent firms' unobserved attributes.

I investigate this by examining whether consumers' preferences with respect to firms they did not choose in 1990 differ depending on whether they passed or failed their previous inspection. Suppose consumers switch firms after failing because they prefer certain types of firms more when they believe their vehicle is likely to fail than when they believe it is likely to pass. Then consumers whose vehicles failed in 1990 should systematically choose different firms than those whose vehicles passed, conditional on the identity of their incumbent firm. Particularly strong evidence for the hypothesis that switching arises from consumers' learning about their vehicles would be if consumers who receive failures and switch firms tend to then choose the same firms as consumers with older vehicles.

5. Results

Estimates of the basic specification are in Table 4. The left side reports nested logits. The first column uses the entire sample; the second uses only vehicles which were observed in both 1990 and 1992. The right side of the table reports estimates from the two step method. In each specification, the omitted station type is "independent garage."

Looking at the nested logits, the price coefficient is negative and significant. From the coefficient on the previous inspection dummy, consumers are more likely to choose firms at which they were observed two years before than their competitors, even if they failed their initial inspection at that time. The coefficient on the previous inspection*pass interaction indicates that they are much more likely to return if they previously passed. From the failure rate coefficient, conditional on their prices, their previous inspection outcome, station type, etc., consumers are more likely to choose firms with low failure rates than those with high ones. They are more likely to choose independent garages and service stations, and less likely to choose new car dealers, the older their vehicle. The estimates in column (1) indicate that conditional on the other included variables, consumers are more likely to choose a firm if it is their vehicle's own new car dealer, particularly if a Federally-mandated emission warranty applies. The coefficient on the "warranty applies" dummy turns insignificant and is imprecisely estimated in column (2) because relatively few vehicles receiving inspections in both

1990 and 1992 were still on warranty during 1992. Applying a Wald test, the inclusive value parameters are jointly different than one at any conventional significance level; one can thus reject independence of irrelevant alternatives.

The numbers at the bottom of the table are derived from the coefficient estimates; these normalize the coefficients on previous inspection*pass and failure rate by that on price. These correspond to β_2 and β_3: how much expected repair costs change with individual inspection outcomes and differ with differences in firms' failure rate across all consumers. From the point estimates in column (2), failing, rather than passing, a previous inspection is associated with the same difference in the probability a consumer chooses a firm as a \$22.42 change in price. Likewise, having a ten percentage point higher failure rate is associated with the same difference in market share as having a \$4.67 higher price. These figures are lower when using the estimates from the balanced panel, due mainly to the higher (in absolute value) price coefficient.

Examining the two-stage estimates on the right side of the table, the coefficients on variables which vary across individuals are about the same as those in the simple logits. Because of the small sample size and paucity of good instruments, standard errors for the parameters estimated in the second stage are very high. The price and failure rate coefficients are smaller in absolute value than those estimated above and are not significantly different from zero, but are still of the correct sign. The point estimate for the failure rate/price ratio is similar to that in the nested logits. The point estimate for the previous inspection*pass/price ratio is higher than that in the nested logits. In both cases, however, the standard errors are extremely high.[23] Using a Hausman test, one does not reject the null of exogeneity at any conventional significance level. The low quality of instruments means, however, that this statistical test has little power to reject.

Table 5 contains results from four specifications which are estimated using simple nested logits. The first two investigate whether consumer behavior differs with variables which are correlated with the probability vehicles fail. The first contains a full set of interactions with vehicle age. The only statistically significant interactions are with price, the new car dealer and tune up shop

[23]The standard errors reported for the previous inspection*pass/price ratio probably understate the true standard errors. The previous inspection*pass and price coefficients come from two different stages. The formula used to compute the standard errors on this ratio treats the point estimate for previous inspection*pass as a constant rather than an estimated value.

dummies, and the warranty applies dummy. Consumers with older vehicles are more price sensitive than those with newer ones. This is probably due to the fact that such individuals are, on the average, less wealthy. The interactions on the previous inspection and previous inspection*pass dummies and failure rate are not statistically significantly different from zero. From these estimates, neither the relationship between previous inspection outcome and choice of firms nor that between failure rate and choice differs with vehicle age. One reason may be that although newer vehicles are less likely to fail than older ones, they are more expensive to repair when they do.

The second specification includes interactions with a dummy variable which equals one if the vehicle failed its 1990 inspection, and zero otherwise. If consumers who previously failed have different preferences than those who previously passed, switching may reflect learning about their vehicles. Looking at the estimates on the "fail in 1990" interactions, none are statistically significant. Using a likelihood ratio test, one cannot reject the null that these parameter estimates are jointly equal to zero at any conventional significance level. Consumers may use inspection outcomes to update their beliefs about their vehicles' underlying condition, but there is no evidence that this changes consumers' preferences among non-incumbent firms.[24] This lends support to the hypothesis that the reason consumers are more likely to switch firms after failing than passing has to do with changes in conjectures with respect to their incumbent firm.

The third specification includes interactions between failure rate and several variables, including a "new to market" dummy variable which equals one if the vehicle's 1990 inspection was outside of Fresno county.[25] From the failure rate*previous inspection coefficient, consumers are more likely to choose a firm at which they previously failed, the lower its failure rate. The interaction between failure rate and previous inspection*pass is positive and large, but not statistically significant. This provides weak evidence that consumers are more outcome-sensitive at firms with high failure rates. From the coefficient on the new car dealer interaction, small

---

[24]I have also run specifications in which I model consumers' choice of firms as a function of a firm-specific constant, previous inspection and previous inspection*pass dummies, and interactions between the firm constants and a "fail in 1990" dummy. The interactions test whether consumers who failed in 1990 are more likely to choose particular firms than those who passed. One cannot reject the null that these interactions are jointly equal to zero.

[25]The omitted interaction is (independent garage or tune up shop)*failure rate. Because there is only one tune up shop, one cannot separately identify a tune up shop*failure rate interaction.

differences among new car dealers' failure rates have the same relationship with consumers' choice as larger differences among independent garages. This may reflect that repair costs given a failure tend to be higher at new car dealers than at other firms. Alternatively, consumers might expect a low proportion of vehicles to fail at new car dealers and be more sensitive to hearing of a failing inspection. Adding the failure rate and failure rate*service station coefficients, consumers are more likely to choose service stations with high failure rates than those with low ones. The new to market interaction is small and not statistically significant. One interpretation is that consumers who moved to Fresno in the previous two years and those who have lived in the county for at least two years are equally informed. This suggests that to the extent information spreads across consumers, it does so relatively quickly.

The fourth specification includes variables which interact previous firm with "failed emission test," "failed underhood test," and "failed both." The omitted interaction is "passed inspection." Parameter estimates are negative and significant for both the emission and underhood interactions, but they are not statistically different from each other. The "failed both" interaction is statistically zero. Failing both parts of the inspection affects consumers' choice more than failing either the emission or the underhood test. Its effect is statistically the same as the effect of failing the emission test plus that of failing the underhood test.

Implications for Firms' Incentives

For each consumer observed at one of the 29 firms in my sample in 1990, one can calculate the estimated probability they choose the same firm in 1992, conditional on passing and failing. The difference in the probabilities — the probability "deltas" — shows the extent to which firms' demand from existing individual customers changes with respect to single transactions. One can also calculate own price and failure rate elasticities of demand for each firm. Although they are not the focus of this paper, own price elasticities are of interest because they provide a check on the model. If they appear unreasonable, this may indicate that endogeneity is a problem or that the demand system is otherwise poorly specified. If one assumes that the data reflect a single-period Nash equilibrium in prices, one can calculate the mark-ups and marginal costs implied by the own price elasticity estimates and use these to check the model as well.

The parameter estimates from the second column in Table 4 imply that, averaged across consumers observed at one of the 29 firms in my sample in 1990, the probability that a consumer chooses the same firm in 1992 is .70 if they passed in 1990 and .54 if they failed. The average difference is .16; the 25th and 75th percentile values are .12 and .20. These probabilities are conditional on choosing one of the firms in my cluster in 1992; the unconditional probabilities and differences are somewhat lower. These figures imply that on the average, if a firm passes, rather than fails, a consumer's vehicle this increases the probability the consumer chooses the firm two years hence by about 30%.

Table 6 contains mean estimated elasticities, mark-ups, and marginal costs derived from the full sample estimates in the first column in Table 4. Failure rate elasticities are relatively high at independent garages: on the average, they are approximately -1. Having a 24% failure rate rather than a 16% failure rate, all else equal, would reduce demand by one-half. Elasticities are lower at the other station types. One can convert these elasticities to a different measure which answers the question: if the firm changed its organizational characteristics such that one additional vehicle failed per time period, given the distribution of vehicles it inspects how much would demand decrease per period?[26] This is given in the columns labeled dQ/d(outcome). The failure rate elasticities imply that over the long run failing one additional vehicle per month would decrease demand by 5.6 inspections/month on the average across the independent garages in my sample, 2.4 on the average across the service stations, and 1.7 across the new car dealers. Individual firms' demand is quite sensitive to their failure rate.

Multiperiod mechanisms play an important role in aligning buyers' and sellers' incentives. When vehicles fail, this lowers the probability their owners choose the firm in the future, and tends to reduce demand across all consumers by an economically significant amount. Inspectors, who are also mechanics, have an incentive to help vehicles pass because they are generally paid a function of the work they complete. Firms have an incentive to adopt organizational characteristics that encourage inspectors to do so. The way they compensate inspectors and mechanics may be a manifestation of this.

---

[26]Using the formula for own failure rate elasticity of demand and the relationship: $dF_j/d(\text{outcome})=[(Q_{fj}+1)-Q_{fj}]/Q_j$, where F_j is the failure rate at firm j, Q_j is the number of inspections per period, and Q_{fj} is the number of failures per period, one obtains $dQ_j/d(\text{outcome})=\eta_{ff}/F_j$.

Independent garages have the highest own price elasticities and the tune up shop the lowest. The own price elasticities range between -0.46 and -4.23, at observed quantities. Four firms' estimated elasticities are less than one in absolute value, implying that they are operating on the inelastic portion of their residual demand curves. While this is inconsistent with profit-maximizing behavior if one assumes that firms choose prices to maximize single period profits, taking their own and others' characteristics (which include all those affecting inspection conduct) as given, these assumptions may not hold.[27] Setting this aside for a moment and assuming that the data reflect a "Nash-in-prices" short-run equilibrium, these imply that new car dealers tend to have high mark-ups (generally $25-35) and independent garages low ones (generally $6-10). Marginal costs tend to be high at new car dealers and low at service stations: possibly because of differences in labor or opportunity costs. As a whole, these calculations lend support to the model; except possibly for the low price elasticity estimate for the tune up shop, the quantities do not appear to be unreasonable.

The far right column uses the estimated inspection mark-ups to obtain estimates of firms' loss in inspection profits from failing an additional vehicle per month. This is the product of each firm's estimated mark-up and dQ/d(outcome). Under the "Nash-in-prices" assumption, this is by construction the same across firms and equal to $-\beta_3$: $46.71.[28] In a model where firms choose organizational characteristics as well as prices, they would choose such characteristics such that the marginal benefits from being slightly less consumer-friendly were equal to the marginal costs. This figure is appealing in this context, because it seems within the range of firms' average incremental profits from supplying emission repairs to an additional failing vehicle.

<u>Implications for Consumer Beliefs and Informedness</u>

In this section I examine what the parameter estimates imply about consumers' beliefs and informedness under the working assumption that changes and differences in expected repair costs reflect only changes and differences in their beliefs about firms' unobserved characteristics. The ratios at the bottom of Table 4 are both significantly different from zero. Assuming that g(.) is linear in n, one can reject the null hypotheses that all consumers are completely informed about all firms,

---

[27]Firms may not set price myopically, or may choose variables other than prices in the short run.

[28]In a single period Nash equilibrium in prices, $p_j - mc_j = -p_j/\eta_{pj}$. Multiplying the right side by dQ/d(outcome) gives $-(p_j\,\eta_{fj})/(F_j\,\eta_{pj})$. Substituting $\eta_{fj}/\eta_{pj} = -\beta_3 F_j/p_j$ gives the stated result.

and that all consumers are completely uninformed about firms via inspection outcomes. The estimates are consistent with a model in which consumers believe there are unobserved differences across firms and they are to some extent informed about such differences through inspection outcomes.

The ratio of the two estimates provides a metric for the extent to which they are informed. Using the full sample estimates from columns 1 and 2 in Table 4, this ratio is equal to 2.08 and 2.65, respectively. Consumers are behaving as if they have conjectures about firms which are linear in n, believe that the inspection outcomes they observe at individual firms are random draws from a common pool, and use outcomes from two to three inspections at each firm in the market in forming their beliefs. The empirical estimates are consistent with a model in which consumers are incompletely informed about individual firms, and use only a very small fraction of the total number of inspections in the market in forming their conjectures about suppliers.

Assuming that initial priors are distributed $B(\alpha,\beta)$, one can use the estimates to characterize consumers' initial priors. This is reported in Table 7. I calculate estimates for $(\alpha+\beta)$ under various assumptions about R, the expected cost of repairs given they fail.[29] This is reported in the first column of each panel. If one assumes a value for the mean of the initial priors, one can use the estimate of $(\alpha+\beta)$ to compute an estimate of the variance of the initial priors. This is in the second column; I assume the mean to be 0.20, which is approximately equal to the fraction of vehicles which fail initial inspections. To provide better intuition regarding the magnitude of the variance, I calculate the support of a uniform distribution that would have the same variance. This is in the third column. Within the range of R considered here, the estimates are consistent with a model in which consumers have diffuse initial priors. Even assuming that R=$120, the estimates imply initial priors that have about the same variance as a U(0,0.6) distribution. Thus, one interpretation of the empirical results is that consumers believe that there is considerable unobserved heterogeneity in the market and choose among firms based on a moderate amount of information from their and others' experiences. Single inspection outcomes strongly influence their beliefs about individual firms. Firms have strong reputational incentives to supply passing inspections.

---

[29]The state of California estimates that the average price of repairs, given that a vehicle fails, is $50-75. This does not include time costs and other costs associated with failing.

6. Conclusion

Buyers often cannot directly detect shirking in service markets, and cannot always tell how much sellers shade diagnoses. Like "output" in standard moral hazard models of labor supply, diagnoses indicate the extent to which sellers act in buyers' interest. Sellers have an incentive to supply diagnoses in which they appear to be acting in buyers' interest. Examples of these are when repairmen propose simple fixes, doctors prescribe generics instead of more expensive but chemically identical brand name drugs to self-insured patients, brokers advise holding stocks rather than selling them, and attorneys recommend that cases not be pursued. Sellers may have more opportunity to supply diagnoses which appear favorable in the vehicle inspection market than in other markets. Regulations require buyers to purchase inspections regardless of their vehicles' condition.[30] Interpreting diagnoses may be more difficult when the "first-best" diagnosis is to recommend extensive treatment.[31]

Buyers' beliefs about firms are important in "diagnosis-cure" markets. Skepticism about the motivation of individual repairmen, brokers, attorneys, and other suppliers of advice combined with beliefs that there exists unobserved heterogeneity across sellers means that sellers have reputational incentives, particularly when switching costs are low. This is reminiscent of Fama (1980). Incentives are weaker when consumers are naive about sellers' private objectives, believe that sellers are homogeneous, or when switching costs are high. These conditions seem not to characterize most "blue collar" service markets. For example, conventional wisdom about auto repair markets and highly-publicized incidents of fraud tends to make buyers wary about the advice they receive. They may be more common in "white collar" service markets such as health care. Portrayals of doctors as objective suppliers of medical advice and treatment, combined with the informational advantage incumbents have in supplying treatment to their existing patients, may dampen demand-side quality incentives.

---

[30]Indeed, were it not for regulators' attempts to encourage inspectors to fail high-emitting vehicles, very few vehicles would fail inspections.

[31]Demand incentives may encourage sellers to supply verification for their diagnoses in such circumstances. Alternatively, an industry structure may evolve in which some firms supply advice but not extensive treatments, as in Wolinsky (1993).

Appendix

The standardized $B(\alpha, \beta)$ $[\alpha>0, \beta>0]$ distribution encompasses a wide range of densities whose support is (0,1). The uniform U(0,1) is the special case $\alpha=\beta=1$. If $\alpha>1$, $\beta>1$, the density has a single interior mode. If $\alpha<1$, $\beta<1$ it is U-shaped. If $\alpha<1$, $\beta>1$ it is "J-shaped": monotonically decreasing within the unit interval. If $\alpha>1$, $\beta<1$ it is monotonically increasing. This class of distributions does not allow densities with multiple interior modes. The expected value of a random variable drawn from a $B(\alpha, \beta)$ is $\alpha/(\alpha + \beta)$. It is therefore an increasing function of the ratio of the parameters, α/β. The variance is $E(X)[1-E(X)]/(\alpha + \beta + 1)$, where E(X) is the expected value. Holding the ratio of the parameters constant (and thus the expected value), the variance is a decreasing function of $(\alpha + \beta)$. If an individual's initial priors are distributed $B(\alpha, \beta)$, $(\alpha + \beta)$ is a measure of how diffuse they are around the individual's expectation.

References

Abreu, Dilip, David Pearce, and Ennio Stachetti, "Optimal Cartel Equilibria with Imperfect Monitoring," Journal of Economic Theory, 39 (1986), 251-69.

Berry, Steven T., "Estimating Discrete-Choice Models of Product Differentiation," Rand Journal of Economics, 25 (1994), 242-62.

DeGroot, Morris H., Optimal Statistical Decisions, McGraw-Hill, New York, 1970.

Fama, Eugene F., "Agency Problems and the Theory of the Firm," Journal of Political Economy, 88 (1980), 288-307.

Goldberg, Pinelopi Koujianou, "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry," Econometrica, 63 (1995), 891-952.

Green, Edward J. and Robert H. Porter, "Non-cooperative Collusion Under Imperfect Price Information," Econometrica, 54 (1984), 532-54.

Holmstrom, Bengt, "Managerial Incentive Problems: A Dynamic Perspective," in Essays in Economics and Management in Honor of Lars Wahlbeck, Helsinki, Swedish School of Economics, 1982.

Hubbard, Thomas N., "Agency Relationships in the Vehicle Emission Inspection Market: Empirical Analysis and Public Policy Implications," Doctoral Dissertation, Stanford University, 1996.

Hubbard, Thomas N., "An Empirical Examination of Moral Hazard in the Vehicle Inspection Market," Manuscript, University of California, Los Angeles, 1997.

Klein, Benjamin, and Keith B. Leffler, "The Role of Market Forces in Assuring Contractual Performance, Journal of Political Economy, 89 (1981), 615-641.

Wolinsky, Asher, "Competition in a Market for Informed Experts' Services," Rand Journal of Economics, 24 (1993), 380-398.

Table 1
Previous Inspection Results and Stay/Switch Decision

							Percent Choosing Same Firm
Observations	7519						
Previous Inspection Not Observed	4815	64.0%					
Previous Inspection Observed	2704	36.0%					
New to Market			263	9.7%			
Old to Market			2441	90.3%			
Old to Market, Passed Previous					2050	84.0%	38.8%
Old to Market, Failed Previous					391	16.0%	25.3%
Old to Market, Failed Emission Only					204	10.0%	29.9%
Old to Market, Failed Underhood Only					120	5.9%	25.0%
Old to Market, Failed Both					67	3.3%	11.9%
New to Cluster			1418	52.4%			
Old to Cluster			1286	47.6%			
Old to Cluster, Passed Previous					1109	86.2%	71.7%
Old to Cluster, Failed Previous					177	13.8%	55.9%
Old to Cluster, Failed Emission Only					105	9.5%	58.1%
Old to Cluster, Failed Underhood Only					51	4.6%	58.8%
Old to Cluster, Failed Both					21	1.9%	38.1%

Table 2
Inspection Outcome Logits

Dependent Variable: 1992 Inspection Outcome

Variable	1	2	3	4
C	**-1.66**	-1.25	-1.01	-1.03
	(0.06)	(0.80)	(0.77)	(1.01)
Fail in 1990	**1.33**	**0.82**	**0.77**	**0.85**
	(0.11)	**(0.12)**	**(0.13)**	**(0.15)**
Same Firm				**-0.46**
				(0.14)
Same Firm*Fail in 1990				-0.27
				(0.30)
Vehicle Age Dummies?	No	Yes	Yes	Yes
Vehicle Make Dummies?	No	No	Yes	Yes
Firm Dummies?	No	No	No	Yes
Probability "Delta"				
Fail in 1990	0.258	0.135	0.122	0.131
Same Firm				-0.055

N=2704, standard errors in parentheses.
Bold indicates significance of t-test H0: b=0 of size 0.05.

Table 3
Firm Prices, Quantities, Failure Rates, Station Types

Firm	Price	N	Share	Failure Rate	Station Type
1	$19.95	1908	25.38%	25.2%	Tune Up Shop
2	$21.95	1247	16.58%	16.2%	Independent Garage
3	$19.95	995	13.23%	16.5%	Service Station
4	$19.95	300	3.99%	19.1%	Independent Garage
5	$19.95	276	3.67%	5.4%	New Car Dealer
6	$59.95	208	2.77%	5.8%	New Car Dealer
7	$65.00	205	2.73%	3.9%	New Car Dealer
8	$31.00	205	2.73%	15.1%	Service Station
9	$60.00	167	2.22%	3.4%	New Car Dealer
10	$65.00	163	2.17%	4.9%	New Car Dealer
11	$23.95	163	2.17%	31.3%	Independent Garage
12	$50.00	159	2.11%	7.5%	New Car Dealer
13	$29.95	151	2.01%	5.3%	Independent Garage
14	$63.00	143	1.90%	8.4%	New Car Dealer
15	$30.00	143	1.90%	11.9%	New Car Dealer
16	$25.00	130	1.73%	17.7%	Service Station
17	$31.00	128	1.70%	17.2%	Service Station
18	$30.00	119	1.58%	14.3%	Independent Garage
19	$60.00	96	1.28%	6.3%	New Car Dealer
20	$24.00	90	1.20%	6.7%	Service Station
21	$29.95	82	1.09%	13.4%	Service Station
22	$24.95	75	1.00%	24.0%	Service Station
23	$55.00	73	0.97%	2.8%	New Car Dealer
24	$65.00	70	0.93%	4.3%	New Car Dealer
25	$59.00	65	0.86%	4.6%	New Car Dealer
26	$19.76	58	0.77%	24.1%	Independent Garage
27	$25.00	45	0.60%	15.9%	Independent Garage
28	$62.00	37	0.49%	8.1%	New Car Dealer
29	$49.95	18	0.24%	33.3%	Independent Garage

Number of Observation	7519
Mean Inspection Price (Firms)	$39.32
Mean Inspection Price (Inspections	$29.27
Corr(Price, Share)	-0.36
Corr(Failure Rate, Share)	0.26

Table 4
Consumer Choice Equation Estimates

Dependent Variable:	Firm	Number of Firms:	29
-LogL at Alpha=0, Beta=0:	25319		

	(1)	(2)	Stage 1	Stage 2: IV	Stage 2: OLS
Number of Observations	7519	2704	2704	29	29
-LogL	18791.4	5524.8	4941.5		
R-Squared				0.551	0.505
Price	**-0.035**	**-0.048**		-0.014	**-0.028**
	(0.001)	**(0.003)**		(0.039)	**(0.008)**
Previous Inspection	**1.900**	**2.047**	**1.998**		
	(0.146)	**(0.151)**	**(0.149)**		
Prev. Inspection*Pass	**0.788**	**0.772**	**0.668**		
	(0.152)	**(0.156)**	**(0.158)**		
Failure Rate	**-1.643**	**-2.044**		-0.624	-1.124
	(0.108)	**(0.237)**		(5.768)	(1.635)
New Car Dealer	**-0.270**	**-0.789**		-0.224	0.078
	(0.102)	**(0.187)**		(0.664)	(0.392)
Tune Up Shop	**0.449**	**0.799**		**1.550**	**1.473**
	(0.064)	**(0.130)**		**(0.562)**	**(0.514)**
Service Station	**-0.773**	**-0.579**		-0.191	-0.226
	(0.082)	**(0.137)**		(0.375)	(0.254)
New Car Dealer*Vehicle Age	**-0.170**	**-0.087**	**-0.086**		
	(0.008)	**(0.016)**	**(0.015)**		
Tune Up Shop*Vehicle Age	**-0.030**	**-0.035**	**-0.036**		
	(0.006)	**(0.011)**	**(0.011)**		
Service Station*Vehicle Age	**-0.019**	-0.009	-0.011		
	(0.006)	(0.010)	(0.010)		
Warranty Applies	**1.004**	0.250	0.163		
	(0.091)	(0.230)	(0.190)		
Own Dealer	**1.794**	**2.106**	**1.612**		
	(0.093)	**(0.156)**	**(0.170)**		
Inc. Value -- NCD	**0.813**	**0.800**	**0.568**		
	(0.039)	**(0.060)**	**(0.053)**		
Inc. Value -- Service Statior	**0.564**	**0.602**	**0.605**		
	(0.032)	**(0.037)**	**(0.039)**		
Inc. Value -- Garage	**0.221**	**0.352**	**0.498**		
	(0.014)	**(0.032)**	**(0.042)**		
Constant				0.596	**1.070**
				(2.122)	**(2.122)**
Previous Inspection*Pass/Price	**$22.42**	**$16.11**		$47.81	**$23.86**
	($4.40)	**($3.35)**		($133.07)	**($7.39)**
Failure Rate/Price	**$46.71**	**$42.66**		$44.66	$40.76
	($2.24)	**($4.42)**		($320.36)	($60.92)

Standard errors in parentheses, bold indicates significance of t-test H0: b=0 of size 0.05.
Omitted station type is independent garage.
Stage 1 includes 29 firm dummies (not shown).
Stage 2 instruments are station type and number of inspectors of firms' closest geographic competitor.

Table 5
Consumer Choice Equation Estimates
Vehicle Age, Previous Result, Failure Rate, Component Outcome Interactions

Dependent Variable:	Firm	Number of Firms:	29
Number of Observations:	2704	-LogL at Alpha=0, Beta=0:	9105

-LogL	5507.7		5506.6		5475.0		5522.3	
	Parameter Estimate	Standard Error	Parameter Estimate	Standard Error	Parameter Estimate	Standard Error	Parameter Estimate	Standard Error
Price	**-0.037**	**0.005**	**-0.037**	**0.005**	**-0.059**	**0.004**	**-0.048**	**0.003**
Previous Inspection	**1.881**	**0.388**	**1.904**	**0.338**	**3.531**	**0.447**	**2.817**	**0.098**
Prev. Inspection*Pass	0.427	0.416	0.399	0.417	-0.021	0.463		
Failure Rate	**-2.049**	**0.481**	**-2.034**	**0.483**	**-2.832**	**0.304**	**-2.044**	**0.236**
New Car Dealer	**-0.785**	**0.202**	**-0.778**	**0.204**	-0.147	0.238	**-0.791**	**0.187**
Tune Up Shop	**0.907**	**0.139**	**0.899**	**0.139**	**0.884**	**0.132**	**0.791**	**0.130**
Service Station	**-0.519**	**0.141**	**-0.522**	**0.142**	**-1.352**	**0.184**	**-0.587**	**0.138**
Warranty Applies	**-2.964**	**1.067**	**-2.977**	**1.068**	0.200	0.258	0.249	0.230
Own Dealer	**2.074**	**0.288**	**2.085**	**0.291**	**2.513**	**0.183**	**2.107**	**0.155**
Vehicle Age*Price	**-0.0011**	**0.0005**	**-0.0010**	**0.0005**				
Vehicle Age*Prev. Ins.	0.016	0.027	0.015	0.026				
Vehicle Age*Prev. Pass	0.051	0.032	0.053	0.032				
Vehicle Age*Failure Rate	-0.003	0.043	-0.011	0.045				
Vehicle Age*NCD	**-0.078**	**0.018**	**-0.080**	**0.019**	**-0.090**	**0.016**	**-0.087**	**0.016**
Vehicle Age*Tune Up	**-0.041**	**0.011**	**-0.038**	**0.012**	**-0.036**	**0.011**	**-0.035**	**0.011**
Vehicle Age*Serv. Sta.	-0.014	0.011	-0.014	0.011	-0.010	0.010	-0.099	0.010
Vehicle Age*Warranty Applies	**0.736**	**0.235**	**0.740**	**0.235**				
Vehicle Age*Own Dealer	-0.002	0.033	-0.004	0.035				
Fail in 1990*Price			0.005	0.006				
Fail in 1990*Failure Rate			0.289	0.603				
Fail in 1990*NCD			0.071	0.243				
Fail in 1990*Tune Up			-0.107	0.164				
Fail in 1990*Serv. Sta.			-0.010	0.156				
Fail in 1990*Own Dealer			0.052	0.325				
Failure Rate*Prev. Ins.					**-7.351**	**2.008**		
Failure Rate*Prev. Ins.*Pass					3.784	2.179		
Failure Rate*NCD					**-14.303**	**2.608**		
Failure Rate*Service Station					**4.804**	**0.685**		
Failure Rate*New to Market					-0.174	0.549		
Prev. Inspection*Fail Emission							**-0.539**	**0.201**
Prev. Inspection*Fail Underhood							**-0.921**	**0.262**
Prev. Inspection*Fail Both							0.037	0.483
Inc. Value -- NCD	**0.772**	**0.060**	**0.773**	**0.060**	**0.939**	**0.072**	**0.800**	**0.060**
Inc. Value -- Service Station	**0.622**	**0.039**	**0.623**	**0.039**	**0.615**	**0.039**	**0.602**	**0.037**
Inc. Value -- Garage	**0.380**	**0.035**	**0.380**	**0.035**	**0.365**	**0.033**	**0.349**	**0.032**

Standard errors in parentheses, bold indicates significance of t-test H0: b=0 of size 0.05.
Omitted station type is independent garage.

Table 6
Mean Estimated Elasticities, Markups, Marginal Costs
By Station Type

	Frate Elasticity	dQj/ d(outcome)	Own Price Elasticity	Markup	MC	d(Ins. Profits)/ d(outcome)
Service Station	-0.38	2.39	-1.37	$19.67	$6.88	$46.71
	(0.14)	(0.22)	(0.54)	($1.94)	($5.31)	
Independent Garage	-1.06	5.64	-3.14	$8.95	$18.61	$46.71
	(0.41)	(1.21)	(0.64)	($3.21)	($6.67)	
New Car Dealer	-0.10	1.67	-2.00	$27.67	$27.24	$46.71
	(0.04)	(0.13)	(0.52)	($2.26)	($13.66)	
Tune Up Shop	-0.28	1.09	-0.46	$43.18	-$23.23	$46.71
	N/A	N/A	N/A	N/A	N/A	

Standard deviations in parentheses.
Final four columns calculated assuming single period Nash equilibrium in prices.

Table 7
Consumers' Informedness and Initial Priors

Specification	Table 4, Column 1			Table 4, Column 2		
Beta2 Estimate		$22.42			$16.11	
Beta3 Estimate		$46.71			$42.66	
N		2.08			2.65	
Expected Costs, Given Failure	Alpha+Beta	Variance of Initial Priors	Support of Uniform	Alpha+Beta	Variance of Initial Priors	Support of Uniform
$60	0.59	0.10	1.10	1.08	0.08	0.96
$80	1.48	0.06	0.88	2.32	0.05	0.76
$100	2.38	0.05	0.75	3.56	0.04	0.65
$120	3.27	0.04	0.67	4.80	0.03	0.58

The mean of the initial priors is assumed to be 0.2.