

How Do Consumers Motivate Experts? Reputational Incentives in an Auto Repair Market

Thomas N. Hubbard*
April 26, 2001

Moral hazard exists in expert service markets because sellers have an incentive to shade their reports of buyers' condition to increase the short-run demand for their services. The California vehicle emission inspection market offers a rare opportunity to examine how reputational incentives work in such a market. I show that consumers are 30% more likely to return to a firm at which they previously passed than one at which they previously failed, and that demand is sensitive to firms' failure rate across all consumers. These and other results suggest that demand incentives are strong in this market because consumers believe that firms differ greatly in their consumer-friendliness and are skeptical even about those they choose. Weak demand incentives in other expert service markets are not a direct consequence of moral hazard, but rather its interaction with switching costs and consumers' beliefs that firms are relatively homogeneous.

*Graduate School of Business, University of Chicago, and National Bureau of Economic Research. Email: thomas.hubbard@gsb.uchicago.edu. I would like to thank Tim Bresnahan, Judy Chevalier, Andrew Dick, Kevin Murdock, Randy Kroszner, Roger Noll, Robert Porter, Jean-Laurent Rosenthal, Andrea Shepard, Scott Stern, Eric Talley, Robert Topel, Darrell Williams, Frank Wolak, many seminar participants and several 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.

1. Introduction

Transactions involving services are not simple exchanges. Production takes place after buyers and sellers agree to the terms of trade. Moral hazard problems arise when buyers can neither perfectly observe nor costlessly verify quality. Sellers can take actions that affect the size and allocation of the gains from trade. In expert service 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 short-run demand for the services they supply.

Theoretical models show that demand-side quality incentives can be weak in expert service markets.¹ Along with highly-publicized incidents of fraud, these models have led academics and policy-makers to suspect market failure and explore how to improve expert service markets' performance, especially in health care but also in blue-collar service markets such as auto repair. But expert service markets need not fail to the degree some theoretical models imply. Reputational concerns may encourage doctors, mechanics, and other experts to act in consumers' interest to some degree, even though consumers can neither perfectly observe nor costlessly verify service quality.

Although there is some sense that reputations are important in expert service markets, how they work and the key institutions that support them are not well understood. This largely reflects empirical difficulties. In game-theoretic models, reputational incentives depend critically on relationships between agents' decisions and their previous experiences. Empirical research that studies how reputations work in light of these models therefore requires individual-level panel data on demand and experiences. Such data are rarely available to researchers. Consequently, applied work on reputations such as that compiled and discussed by Klein (1997) has not been highly empirical.

This paper helps fill this gap by investigating how reputation mechanisms work in one segment of the auto repair market – the market for California vehicle inspections. My previous work (Hubbard (1998)) showed that inspection suppliers tend to help vehicles pass, even though they stand to benefit from repairing those that fail; this paper probes the deeper question of why demand

¹Darby and Karni (1973), Wolinsky (1993), Taylor (1995). Health economists label this "inducement" to emphasize physicians' incentive to overprescribe care; see Gruber and Owings (1996) and its references.

incentives are strong. Evidence on this question can help improve the focus of remedial efforts in other expert service markets.²

The analysis exploits unusually rich individual-level panel data. I estimate a model of consumer choice, concentrating on two empirical relationships. One is how the probability consumers switch firms is affected by the outcome of their vehicle's previous inspection. The other is how consumers' choice relates to a measure of firms' aggregate performance they do not directly observe: the fraction of vehicles that fail inspections (their "failure rates").³ 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 inspectors and firms help consumers pass because of dynamic incentives, not just because consumers are able to bribe inspectors within single-period contingent contracts. I find strong evidence that this is the case. Consumers are 30% more likely to choose a firm at which their vehicle previously passed than failed. Relationships between demand and failure rates indicate that, in the long run, failing one additional vehicle per month would lower a firm's monthly inspection revenues by an average of \$97.69 and inspection profits by \$46.71.

These and other empirical relationships also provide evidence regarding how and why dynamic incentives work in this market in light of different classes of theoretical models of reputation.⁴ The evidence is neither consistent with the hypothesis that firms and consumers are playing simple trigger equilibrium strategies nor that consumer behavior reflects their learning about their vehicles. Applying an incomplete information framework in which consumers learn about firms' unobserved characteristics, consumers are behaving as if they have information from a small

²Remedial efforts in this market seek to dampen rather than strengthen demand incentives, due to the inspection program's environmental objectives. In markets without such objectives, remedial efforts would seek to strengthen demand incentives.

³Consumers do not directly observe individual firms' failure rates in part because regulators have successfully prevented this information from being published. The context is thus different than in other expert service markets such as financial management in which consumers observe good aggregate performance measures. See Chevalier and Ellison (1997) for an investigation of the incentives of mutual funds.

⁴Klein and Leffler (1981), Shapiro (1983), and Green and Porter (1984) are examples of complete information reputation models. Fama (1980), Holmstrom (1982), Kreps and Wilson (1982), Milgrom and Roberts (1982), Tadelis (1999) apply incomplete information frameworks. The latter class of models informs the career concerns literature in labor economics; see Chevalier and Ellison (1999) for a recent empirical study.

number of inspections in the market and have weak priors about individual firms' type, or consumer-friendliness. The results are consistent with the view that consumers in this market believe that firms differ greatly in their consumer-friendliness and are skeptical even about those they choose. Combined with low switching costs, this explains why demand incentives are strong. It also implies that demand incentives may be weak in other expert service markets where switching costs are higher and consumers are less skeptical, such as health care. Measures that lower switching costs or encourage skepticism in these other markets may improve their performance, even if consumers do not have good information about firms' aggregate performance.

An outline of the rest of the paper follows. Section 2 describes the relevant features of the inspection market and discusses my previous findings in light of this paper's research goals. Section 3 develops the analytic framework. In section 4, I describe the data. I also test whether consumers and firms are playing simple complete information strategies, and whether switching reflects consumers' learning about their vehicle. In section 5, I construct the empirical framework and develop the econometric model used in estimation. Section 6 contains the estimation results and analysis. Section 7 concludes.

2. 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. In general consumers can only obtain a certificate once their vehicle passes an emission inspection.⁵ Inspections and any associated emission-related repairs have little or no private value, and there is little evidence that consumers whose vehicles fail purchase repairs that have lasting emission effects.⁶ 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

⁵New vehicles and those older than the 1966 model year are exempt, as are those with diesel engines. 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. For more detail about the California vehicle inspection market, see Hubbard (1996, 1998).

⁶Hubbard (1997b) summarizes research that strongly suggests that vehicles are repaired so that they are "clean for a day." Also, I have tested whether vehicles that failed an inspection in 1990 exhibit different results in 1992 as a function of the firm that failed (and likely repaired) them, and found no evidence of such differences.

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 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 perfectly 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. 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. For example, the machines can determine whether the probe that measures tailpipe emissions is in a vehicle’s tailpipe. 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.

Previous Findings

Two results from Hubbard (1998) motivate this paper and shape its analytic framework.

The first is that inspectors tend to help vehicles pass when consumers would bear the cost of emission-related repairs. Vehicles are generally much less likely to fail inspections at private firms than inspections conducted by state officials outside of the normal inspection process, holding vehicle characteristics constant.⁷ The only exception to this is when emission repairs are covered

⁷State officials conducted about 2,500 inspections per year outside of the normal program. Regulators compared failure rates from these inspections with those at private firms to evaluate the program. These inspections were conducted on roadsides on vehicles chosen at random. Drivers were neither relieved of inspection requirements if their vehicles passed nor penalized if they failed.

by warranties: late-model vehicles inspected at new car dealers are not less likely to fail. Overall, actions taken by inspectors at private firms cut the fraction of vehicles that fail from about 40% to about 20%. This result implies that demand-side incentives are quite strong in this market. This paper investigates the source of these incentives by examining consumer behavior.

The second result is that there exist systematic differences in the extent to which inspectors at different firms help vehicles pass. Holding constant vehicle characteristics, failure probabilities for both parts of the inspection are much higher at “chain stores” such as Pep Boys and Sears than at independent garages or service stations, and increase with the number of inspectors firms employ. In the previous paper, I explain how these and other patterns reflect differences in the extent to which firms’ organizational characteristics expose individual inspectors to market incentives. For example, free rider problems weaken inspectors’ incentive to help vehicles pass at firms with many inspectors.⁸

The existence of cross-firm heterogeneity in this market shapes this paper’s analytic framework. If all cross-firm differences in inspection conduct were associated with organizational characteristics that consumers can directly observe, firms’ incentives to choose consumer-friendly organizational features would be straightforward. Firms would choose their organizational characteristics based on a trade off between their cost of implementing a consumer-friendly organizational structure and the additional business it would bring in. Consumers would choose among firms, knowing in advance how much inspectors would help them pass. Like in hedonic models, cross-firm heterogeneity would persist in equilibrium because of differences in consumers’ willingness to pay for consumer-friendliness. For example, some consumers may choose chain stores because they value convenience, even though they know that chain stores tend to be less consumer-friendly than other firms.

But some cross-firm differences may not be associated with things consumers can directly observe or verify. When firms’ “type” is unobservable, consumers make choices under uncertainty. Incentive mechanisms are more complicated because they hinge on how things consumers can

⁸The data also contain some evidence on repair intensities for vehicles that fail. Relationships between repair intensities and organizational characteristics mirror those between failure probabilities and characteristics, suggesting that incentives affect inspectors’ behavior with respect to inspections and repairs similarly.

potentially observe but that are not necessarily public information – such as inspection outcomes – change their beliefs about firms’ type. Firms’ incentive to be a consumer-friendly type is weak when consumers’ beliefs about firms are insensitive to outcomes, especially if consumers also have little information about firms outside of their own experiences.

3. Analytic Framework

The timing of events follows. Firms choose their organizational structure, lines of business, and prices. Consumers form beliefs about the cost of obtaining a passing inspection at different firms. In forming these beliefs, consumers may be uncertain about both the condition of their vehicle and the degree to which inspectors at individual firms will help them pass. 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.⁹ 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 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.¹⁰

Firms choose their organizational characteristics, the goods and services they supply, and prices toward maximizing profits across all their lines of business. Organizational characteristics include hierarchies and compensation schemes. I treat these as fixed over long horizons, and 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 that 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.

⁹Throughout this paper, I assume that the support of the distribution of outcomes is the same for all actions.

¹⁰Over 90% of vehicles that fail an inspection are reinspected at the same firm, usually on the same day.

Consumers choose among firms to maximize expected utility.¹¹ 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 that do not offer free reinspections, and for consumers who place a relatively high value on their time.

Relationships between consumers’ choice of firms and previous inspection outcomes can arise in both complete and incomplete information reputation models. They can also arise because inspection outcomes change consumers’ preferences across firms through their beliefs about their vehicles.

Complete Information Models

Suppose consumers have complete information about firms’ characteristics and how they affect inspectors’ behavior. Suppose also that inspection outcomes do not affect consumers’ preferences across firms through their beliefs about their vehicle. Consumers may update about their vehicle, but this shifts expected repair costs by the same amount across firms.

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 that 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,

¹¹I 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.

some consumers may not use simple loyalty-boycott strategies, but inspectors discriminate against those who (optimally) return after failing.

The empirical framework and data cannot reject all models in which demand shifts occur because consumers believe firms discriminate according to previous outcomes, because equilibria can be supported by very complicated strategies. However, one can test whether consumers and firms are using certain simple strategies. For example, one can test whether all consumers are using simple loyalty-boycott strategies by examining whether they always return after passing and never return after failing. One would expect to reject this hypothesis: it is likely that some consumers find it optimal to return after failing. One can test a more interesting class of complete information equilibria by examining whether firms discriminate against consumers who return after failing. Finding that this is not the case empirically makes complete information interpretations of the data less plausible, since this pattern of discrimination would underlie most of the complete information equilibria supported by simple supplier strategies.

Incomplete Information about Firms

Suppose instead that consumers do not believe inspectors discriminate, but are unable to observe firms' type directly. Then expected repair costs may be related to previous inspection outcomes because consumers use them to infer firms' type. The magnitude of relationships between consumers' choice of firm and a) their previous inspection outcome, and b) firms' failure rate across all consumers reflect the strength of their priors about firms' type and the degree to which they utilize information from their and others' inspection outcomes in forming their beliefs (their "informedness").

Suppose consumers believe they are effectively completely informed about firms' type. This could be either because they believe to be no unobserved cross-firm heterogeneity or because their informedness via inspection outcomes is very high. Then there should be no relationship between their choice of firms and their previous inspection outcome. One can therefore test the proposition that consumers are completely informed by testing whether the probability they choose a firm at which they were previously inspected is the same, regardless of whether they passed or failed. Finding that consumers are less likely to choose a firm at which they previously failed than passed implies that they are incompletely informed. The more sensitive their choice is to previous

inspection outcomes, the weaker their priors are about firms' type.

Suppose consumers are completely uninformed via inspection outcomes. Then controlling for firm characteristics they directly observe, there should be no relationship between their choice of firms and firms' failure rates. Therefore, if one finds such a relationship, one can reject the null hypothesis that consumers are completely uninformed. The stronger the relationship, the more informed consumers are. Strong relationships suggest that information from inspection outcomes diffuses significantly across consumers in the market.

Relationships between consumers' choice and their previous inspection outcome and firms' failure rates therefore indicate what motivates firms and inspectors to help vehicles pass. If consumers' choice is not related to failure rates but is very sensitive to their previous inspection outcome, then demand-side incentives are entirely due to inspections' outcomes' effect on single consumers' priors. If consumers' choice is not sensitive to their previous inspection outcome but is strongly related to firms' failure rates, incentives instead arise because consumers are well-informed about firms' type. The empirical results thus shed light on the likelihood that the strong demand-side incentives in this market are due to individual consumers' weak priors, well-working informational networks, or both.

Switching and Learning about Vehicle Condition

As noted above, inspection outcomes can affect expected repair costs not just through consumers' beliefs about inspectors' behavior, but also through their beliefs about their vehicle. If failing an inspection changes expected repair costs disproportionately across firms, consumers will switch firms not just because their beliefs about how inspectors behave change, but also to obtain a more appropriate match between their vehicle and firm. This is the main alternative interpretation of switching behavior.

One can test this interpretation in the following way. If consumers switch because of updating about their vehicles' condition, those who switch firms after passing should tend to choose different firms than those who switch firms after failing. In particular, those who switch after failing should move toward the same firms that tend to inspect older (i.e., high-emitting) vehicles. Finding that this is the case supports the hypothesis that switching in part reflects changes in consumers' beliefs about vehicle condition. Finding that it is not suggests instead that changes in consumers'

beliefs about vehicle condition do not induce switching. One can then interpret switching in light of the models outlined above.

4. Data

The data are similar to those used in Hubbard (1998). They include 7519 observations of vehicles that 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.¹² This cluster of firms is located in a dense, commercially-zoned corridor that 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.

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

¹²“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.

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.¹³ 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.¹⁴ 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 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 that 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 that 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,

¹³I 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.

¹⁴The 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.

by definition, “new to cluster” consumers did not choose the same firm in 1990 and 1992.¹⁵ 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 complete information equilibria in which homogeneous consumers discipline firms by following simple loyalty-boycott 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 periods. Of the 2704 vehicles observed in both years, 41.8% of those that failed in 1990 failed in 1992, but only 16.0% of those that passed in 1990 did. Part of this is due to differences in the vehicles’ age and make. Table 2 reports results from five 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. Including the age and make dummies cuts this figure by more than half, but it is still 12-14 percentage points. This suggests 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. It is therefore important to test the hypothesis that switching behavior arises because of changes in consumers’ beliefs about their vehicle.

The fourth column tests whether firms discriminate against consumers who return after failing inspections. This specification includes a full set of firm dummies, a dummy variable that 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 coefficient on same firm*fail in 1990 does

¹⁵They 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.

not indicate that firms “punish” consumers who return after failing. Vehicles that were inspected at the same firm they were inspected in 1990 tended to pass more than those that were inspected elsewhere, regardless of whether the vehicle passed in 1990. If the results reflect a complete information equilibrium, the equilibrium is supported by supplier strategies that do not dictate that they always discriminate against consumers who return after failing.

This result casts doubt on complete information “trigger equilibrium” interpretations in general. One can reconcile the patterns in the fourth column with a more complicated trigger equilibrium in which suppliers do not always discriminate against consumers who return after failing, but consumers can always anticipate when they will. But complete information equilibria become less plausible when they are based on more complicated triggers. Complicated triggers impose stricter informational requirements on firms and consumers. In this case, consumers and firms would not only have to remember the consumer’s previous inspection outcomes, but also when particular outcomes should lead to a “punishment stage” and when they should not. Because of this, I will interpret further results in light of incomplete information reputation models rather than complete information ones.

The fifth column tests whether relationships between inspection outcomes and switching arise for spurious reasons. If mobile consumers tend to drive high-emitting vehicles, one would observe relationships between switching and inspection outcomes. But these relationships would not reflect that inspection outcomes affect individuals’ demand at their incumbent firms. I investigate this by testing whether – conditional on their age, make, and where they are inspected – vehicles driven by new-to-market consumers are more likely to fail inspections than those driven by old-to-market consumers. The premise is that new-to-market consumers are more mobile than old-to-market ones. The coefficient on the new-to-market dummy in the fifth column is negative but not statistically significant. Vehicles driven by individuals new to Fresno are not more likely to fail than those driven by longer-term residents; if anything, they are more likely to pass. This test does not indicate that mobile consumers drive higher-emitting vehicles, and thus does not provide support for the hypothesis that relationships between switching and inspection outcomes reflect unobserved consumer heterogeneity rather than demand shifts.

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.¹⁶ 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. The table also indicates that some firms have high (low) market shares despite relatively high (low) prices and/or failure rates, suggesting that other firm characteristics such as location and whether advance appointments are necessary affect consumers' choice.

To sum up, simple patterns in the data suggest that dynamic incentives motivate suppliers in this market. Individual consumers are more likely to return to firms at which they previously passed than failed. There is no evidence that this reflects unobserved consumer heterogeneity. The data also provide evidence against simple trigger equilibria: firms do not discriminate against consumers who return after failing. The data do indicate that consumers may learn about their vehicle from inspection outcomes: vehicles that failed in the past are more likely to fail in the future, conditional on their make and age. It is therefore possible that consumers switch firms more after failing to obtain a more appropriate match for their vehicle.

5. 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-

¹⁶The high percentage of new car dealers is due to the fact that the cluster of firms includes an "auto row." None of the estimates of relationships between choice and inspection results change when eliminating new car dealers and the people that choose them from the sample.

related components:

$$V_{ij} = \alpha C_{ij} + \mu_{ij} \quad (1)$$

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:

$$\mu_{ij} = \gamma(OD)_{ij} + \eta_{ij} \quad (2)$$

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:

$$C_{ij} = p_j + E(R_{ij}) + \epsilon_{ij} \quad (3)$$

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 base specification, it is:

$$E(R_{ij}) = f(X_{vi}, X_{ci}, X_{oi}) + \beta_1 D_{ij}^1 + \beta_2 D_{ij}^2 + \beta_3 F_j + \beta_4 W_{ij} + \beta_5 Z_j + \beta_6 (Z_j X_{vi}) + \xi_j + \zeta_{ij} \quad (4)$$

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,¹⁷

— 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.

Under these assumptions, consumer i chooses firm j iff:

$$V_{ij} > V_{is} \quad \forall s \neq j \Leftrightarrow \alpha(p_j + E(R_{ij}) + \epsilon_{ij}) + \mu_{ij} > \alpha(p_s + E(R_{is}) + \epsilon_{is}) + \mu_{is} \quad \forall s \neq j \quad (5)$$

Assuming that $\alpha(\xi_j + \zeta_{ij} + \epsilon_{ij}) + \eta_{ij}$ 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 inspection outcomes, and its elasticity with respect to failure rates.

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. In the base model, I do not include interactions between inspection outcomes and firm characteristics. This model therefore embeds the assumption that if consumers use inspection outcomes to learn about their vehicle, this affects $E(R_{ij})$ by the same amount across firms. In other specifications, I include such interactions. The coefficients on interactions between previous inspection outcome and firm characteristics form the basis of a test for whether switching reflects consumers' learning about their vehicle. Finding that the interactions are statistically different than zero implies that consumers who previously failed and switch choose different firms than those who passed and switch. This would suggest that switching reflects consumers' learning about their vehicle, especially if those who fail then tend to choose the same firms consumers with older vehicles do. Finding that they are not different from zero suggests that relationships between consumers' choice and inspection outcomes reflect changes in their beliefs about firms, not their vehicles.

Many other papers, especially in the marketing literature, examine relationships between

¹⁷Federal 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."

individuals' choice of products and their purchase histories. These papers seek to distinguish among various factors that lead to serial correlation in purchases: unobserved consumer characteristics ("heterogeneity"), brand loyalty, habit, and so on. β_1 , the coefficient on D_{ij}^1 , captures the effects of these sources of serial correlation. Distinguishing among the reasons why consumers tend to choose the same firm they did in the past, irrespective of their previous inspection outcome, is not the focus of this paper. I therefore treat D_{ij}^1 as a control rather than a variable of interest.

The coefficients of interest are β_2 and β_3 . β_2 indicates how much more single inspection outcomes affect expected repair costs at the firm where the inspection takes place than at other firms. β_3 indicates how much expected repair costs differ with differences in firms' failure rates across all consumers. Assuming that these coefficients reflect only their beliefs about firms, these reflect the strength of their priors and their informedness. If $\beta_2=0$, inspection outcomes do not change priors at the margin; if not, they do. Higher values of β_2 (in absolute value) indicate weaker priors – consumers are more skeptical or uncertain about individual firms' type. If $\beta_3=0$, consumers are completely uninformed, where informedness is defined as knowledge about inspection outcomes other than their own previous one. Higher values of β_3 indicate better-informed consumers. Strong relationships between choice and both previous inspection outcomes and failure rates indicate that consumers are both informed and have weak priors. If consumers update beliefs in a Bayesian fashion, this would imply that consumers' "initial priors" – their beliefs about firms, given they have no information via inspection outcomes – are diffuse. In Bayesian models where individuals have continuous unimodal initial priors, priors narrow as individuals become more informed. Finding that both relationships are strong therefore suggests that consumers' behavior reflects a high degree of uncertainty and skepticism: they believe to be large underlying differences in auto repair firms' consumer-friendliness, and are unsure about inspectors' incentives even at the firms they choose.

Endogeneity Issues

Assuming that ξ_j is independent of the other explanatory variables brings up a familiar econometric issue.¹⁸ 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

¹⁸See Berry (1994), Goldberg (1995).

consumers can leave their vehicles to be repaired, and so on.¹⁹ It does not include subjective factors such as "consumer friendliness" that arise from firm characteristics consumers do not directly observe.²⁰ 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 positively biased. Alternatively, suppose that failure rates are lower at firms that 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 $\xi_j, j=1, \dots, 29$, 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 second 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

¹⁹Some 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.

²⁰This is unlike other consumer choice models, in which reputations are considered product characteristics consumers directly observe.

that, of the attributes they directly observe, only those in Z_j and W_{ij} influence expected repair costs.²¹ If this holds, then the interpretation of β_3 above applies. If not, it has a slightly different interpretation, but one that is still of interest with respect to sellers' incentives. Positive and significant estimates of β_3 then in part would indicate that consumers know that firm characteristics they directly observe (but that 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.

The next section presents the results from the demand estimates. I first estimate the coefficients of the reduced form utility function and compute individual consumers' and market demand elasticities with respect to inspection outcomes. These elasticities indicate the strength of dynamic incentives. I then test whether switching reflects learning about vehicles by examining whether the demand of consumers who switch after passing is different from those who switch after failing. Last, I examine the implications of the estimates, assuming that relationships between inspection outcomes and choice only reflect consumers' learning about firms' unobserved attributes.

6. Results

Estimates of the base specification are in Table 4. The left side reports nested logits. The first column uses the entire sample; the second uses only vehicles that 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."

The nested logit estimates in the first two columns show this paper's most important empirical patterns: consumers' choice of firms is both sensitive to their previous inspection outcome and to firms' failure rates. The positive and significant coefficient on the previous inspection*pass interaction indicates that they are much more likely to return to the place that inspected their vehicle

²¹If firms are able to signal their type through their appearance, endogeneity is a problem. At one point, I visited every inspection-supplying firm in Fresno, including all of the firms in this sample. I did not observe any obvious variables that are omitted from the specification and consumers would interpret as a signal. One reason for this is that obvious signals would invite regulatory scrutiny. Regulatory scrutiny also explains why very few firms – and none in this sample – offered consumers contingent "pass or don't pay" contracts during the time of this sample.

two years before if they previously passed than failed. The negative and significant failure rate coefficient indicates that, 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. The bottom of the table reports estimates of β_2 and β_3 , which normalize these coefficients by the price coefficient. These estimates indicate that expected repair costs vary substantially both with single inspection outcomes and firms' failure rates. From column (1), 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 somewhat lower, \$16.11 and \$4.27 respectively, when using the estimates from the balanced panel, due mainly to the higher (in absolute value) price coefficient.²²

The rest of the estimates provide further detail with respect to demand patterns. The price coefficient is negative and significant; demand is downward-sloping. From the coefficient on the previous inspection dummy, consumers are more likely to choose firms at which their vehicle was inspected two years before than any other individual firm, even if their vehicle had failed. They are more likely to choose independent garages (the omitted station type) 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.

Examining the two-stage estimates on the right side of the table, the coefficients on variables

²²Using the "old to market" sample raises the question of whether the average consumer in the market behaves differently than that in the "old to market" subsample. I investigated this by estimating the model with interactions between an "old to market" dummy and price, failure rate, and the station type dummies. I omitted the previous inspection and previous inspection*pass variables to put the matched and unmatched observations on equal footing. "Old to market" consumers are slightly more price-sensitive, as suggested by this table, but there is no evidence of differences in failure rate sensitivity.

that vary across individuals are about the same as those in the simple logits. 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 β_3 is similar to that in the nested logits. The point estimate for β_2 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 that are estimated using simple nested logits. The first two investigate whether consumer behavior differs with variables that 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 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 because their emission control systems tend to be electronic rather than mechanical.²⁴

The second specification includes interactions with a dummy variable that 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

²³The 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.

²⁴In specifications not reported here, I also investigated whether consumers with particular types of vehicles (old vehicles, foreign makes, trucks) tended to select firms that treated these vehicles well relative to their average. I found no evidence of sorting along these more detailed dimensions.

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 their preferences among non-incumbent firms.²⁵ The result that consumers are more likely to switch firms after failing than passing does not appear to be due to consumers' updating about their vehicle's condition.

The third specification includes interactions between failure rate and several variables, including a "new to market" dummy variable that equals one if the vehicle's 1990 inspection was outside of Fresno county.²⁶ 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 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.²⁷

²⁵I 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.

²⁶The 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.

²⁷I have also tested whether the price and station type dummy coefficients differ for new to market consumers and find no evidence that they are.

The fourth specification includes variables that 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. The point estimate is greater in absolute value for the latter, providing some evidence that consumers are more sensitive to underhood than emission failures, but the difference is not statistically significant. 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 that answers the question: if the firm changed its organizational characteristics (e.g., its internal incentive structure) such that one additional vehicle failed per time period, given the distribution of vehicles it inspects how much would demand decrease per period?²⁸ This is given in the column labeled $dQ_j/d(\text{outcome})$. 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. On the average across all firms, failing an additional vehicle per month would decrease inspection revenues by \$97.69 per month. 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 is a manifestation of this.

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

²⁸Using 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 $dQ_j/d(\text{outcome}) = \eta_{F_j} / F_j$.

assumptions may not hold.²⁹ Setting this aside for a moment and assuming that the data reflect a Bertrand-Nash short-run equilibrium in which firms choose prices conditional on their product characteristics and demand, 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 are not 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 $dQ/d(\text{outcome})$. Under the "Nash-in-prices" assumption, this is by construction the same across firms and equal to β_3 : \$46.71.³⁰ 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.

Implications for Consumer Beliefs and Informedness

Assuming that β_2 and β_3 reflect only changes and differences in consumers' beliefs about firms, finding that both are significantly different from zero implies that consumers are neither completely informed nor completely uninformed. The magnitude of β_2 indicates that consumers' priors are weak, even with respect to the firms they choose. Outcomes have the same effect on consumers' subsequent choices as a large change in the inspection price. Combined with β_3 , weak priors exist even though consumers are informed to some extent. This suggests that consumers have diffuse initial priors. In a working paper (Hubbard (1997a)), I develop a more formal version of the incomplete information model and show that the results are consistent with a model in which consumers have $U(0,0.6)$ initial priors about the probability they fail at each firm, and have updated

²⁹Firms may not set price myopically, or may choose variables other than prices in the short run.

³⁰In a single period Nash equilibrium in prices, $p_j - mc_j = -p_j/\eta_{pj}$. Multiplying the right side by $dQ/d(\text{outcome})$ gives $-(p_j \eta_{ij})/(F_j \eta_{pj})$. Using formulae for nested logit demand elasticities, $\eta_{ij}/\eta_{pj} = -\beta_3 F_j/p_j$; substituting this into the previous expression gives the stated result.

their priors using two random inspection outcomes at each firm.

Consumers' demand patterns are as if they believe there is considerable heterogeneity in the market and choose among firms after observing a small amount of information about each firm. Furthermore, from the failure rate*new to market interaction in Table 5, old-to-market consumers are not behaving as if they are more informed about firms than new-to-market ones; one possible explanation for this is that old-to-market consumers put little weight on information that is not recent.

Interpreting the results in light of incomplete information models, I reach the following conclusions. Part of the reason strong demand incentives exist in this market is that consumers believe that there are significant differences in firms' consumer-friendliness, and are skeptical even about the firms they actually choose. While consumers are to some extent informed about cross-firm differences from their and others' previous experiences, single inspection outcomes strongly influence their beliefs about individual firms. Firms have strong reputational incentives to supply passing inspections as a consequence.

7. 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. Some characteristics of the inspection market make it more straightforward for sellers to signal that they are acting in their customers' interest than in other markets. Regulations require buyers to purchase inspections regardless of their vehicles' condition.³¹ Interpreting diagnoses may be more difficult when the "first-best" diagnosis is to recommend extensive treatment. But demand incentives would then encourage sellers to

³¹Indeed, were it not for regulators' attempts to encourage inspectors to fail high-emitting vehicles, very few vehicles would fail inspections.

develop such signals – that is, to supply verification – in such circumstances.³² Alternatively, an industry structure may evolve in which some firms supply advice but not extensive treatments, as in Wolinsky (1993).

Buyers' beliefs about firms are important in expert service 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.

³²See Jin (2000) for a study of accreditation incentives in health care.

References

- Berry, Steven T., "Estimating Discrete-Choice Models of Product Differentiation," Rand Journal of Economics, 25 (1994), 242-62.
- Chevalier, Judith, and Glenn Ellison, "Risk Taking by Mutual Funds as a Response to Incentives," Journal of Political Economy, 105 (1997), 1167-1200.
- Chevalier, Judith, and Glenn Ellison, "Career Concerns of Mutual Fund Managers," Quarterly Journal of Economics, 114 (1999), 389-432.
- Darby, Michael R., and Edi Karni, "Free Competition and the Optimal Amount of Fraud," Journal of Law and Economics, 16 (1973), 67-88.
- 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.
- Gruber, Jonathan, and Maria Owings, "Physician Financial Incentives and Caesarean Section Delivery," Rand Journal of Economics, 27 (1996), 99-123.
- 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., "Consumer Beliefs and Buyer and Seller Behavior in the Vehicle Inspection Market," NBER Working Paper #6245, October 1997a.
- Hubbard, Thomas N., "Using Inspection and Maintenance Programs to Regulate Vehicle Emissions," Contemporary Policy Issues, 15 (1997b), 52-62.
- Hubbard, Thomas N., "An Empirical Examination of Moral Hazard in the Vehicle Inspection Market," Rand Journal of Economics, 29 (1998), 406-426.
- Jin, Ginger Z., "Competition and Disclosure Incentives: An Empirical Study of HMOs," University

of Maryland, 2000.

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

Klein, Daniel B., Reputation: Studies in the Voluntary Elicitation of Good Conduct, University of Michigan Press, Ann Arbor, 1997.

Kreps, David and Robert Wilson, "Reputation and Imperfect Information," Journal of Economic Theory, 27 (1982), 253-279.

Milgrom, Paul and John Roberts, "Predation, Reputation, and Entry Deterrence," Journal of Economic Theory, 27 (1982), 280-312.

Shapiro, Carl, "Premiums for High Quality Products as Returns to Reputations," Quarterly Journal of Economics, 98 (1983), 659-679.

Tadelis, Steven, "What's in a Name? Reputation as a Tradeable Asset," American Economic Review, 89 (1999), 548-563.

Taylor, Curtis R., "The Economics of Breakdowns, Checkups, and Cures," Journal of Political Economy, 103, 53-74.

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 <u>Choosing Same Firm</u>
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	5
C	-1.66 (0.60)	-1.25 (0.80)	-1.01 (0.77)	-1.03 (1.01)	-0.94 (1.02)
Fail in 1990	1.33 (0.11)	0.82 (0.12)	0.77 (0.13)	0.85 (0.15)	0.84 (0.15)
Same Firm				-0.46 (0.14)	-0.49 (0.14)
Same Firm*Fail in 1990				-0.27 (0.30)	-0.26 (0.30)
New to Market					-0.21 (0.19)
Vehicle Age Dummies?	No	Yes	Yes	Yes	Yes
Vehicle Make Dummies?	No	No	Yes	Yes	Yes
Firm Dummies?	No	No	No	Yes	Yes

Probability "Delta"

Fail in 1990	0.258	0.135	0.122	0.131	0.130
Same Firm				-0.055	-0.059

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

<u>Firm</u>	<u>Price</u>	<u>N</u>	<u>Share</u>	<u>Failure Rate</u>	<u>Station Type</u>
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: -LogL at Alpha=0, Beta=0:	Firm 25319	Number of Firms:	29			
	(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 (0.001)	-0.048 (0.003)		-0.014 (0.039)	-0.028 (0.008)	
Previous Inspection	1.900 (0.146)	2.047 (0.151)	1.998 (0.149)			
Prev. Inspection*Pass	0.788 (0.152)	0.772 (0.156)	0.668 (0.158)			
Failure Rate	-1.643 (0.108)	-2.044 (0.237)		-0.624 (5.768)	-1.124 (1.635)	
New Car Dealer	-0.270 (0.102)	-0.789 (0.187)		-0.224 (0.664)	0.078 (0.392)	
Tune Up Shop	0.449 (0.064)	0.799 (0.130)		1.550 (0.562)	1.473 (0.514)	
Service Station	-0.773 (0.082)	-0.579 (0.137)		-0.191 (0.375)	-0.226 (0.254)	
New Car Dealer*Vehicle Age	-0.170 (0.008)	-0.087 (0.016)	-0.086 (0.150)			
Tune Up Shop*Vehicle Age	-0.030 (0.006)	-0.035 (0.011)	-0.036 (0.110)			
Service Station*Vehicle Age	-0.019 (0.006)	-0.009 (0.010)	-0.011 (0.100)			
Warranty Applies	1.004 (0.091)	0.250 (0.230)	0.163 (0.190)			
Own Dealer	1.794 (0.093)	2.106 (0.156)	1.612 (0.170)			
Inc. Value -- NCD	0.813 (0.039)	0.800 (0.060)	0.568 (0.530)			
Inc. Value -- Service Station	0.564 (0.032)	0.602 (0.037)	0.605 (0.390)			
Inc. Value -- Garage	0.221 (0.014)	0.352 (0.032)	0.498 (0.420)			
Constant				0.596 (2.122)	1.070 (0.428)	
Previous Inspection*Pass/Price (Beta2)	-\$22.42 (\$4.40)	-\$16.11 (\$3.35)		-\$47.81 (133.07)	-\$24.23 (\$7.39)	
Failure Rate/Price (Beta 3)	\$46.71 (\$2.24)	\$42.66 (\$4.42)		\$44.66 (320.36)	\$40.76 (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

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

Standard deviations in parentheses.

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