THE JOURNAL OF INDUSTRIAL ECONOMICS 0022-1821
Volume XLIX December 2001 No. 4

INTERNET CAR RETAILING*

FIONA SCOTT MORTON,† FLORIAN ZETTELMEYER‡ AND JORGE SILVA-RISSO§

We investigate the effect of Internet car referral services on dealer pricing of automobiles in California. Customers of an online service pay on average 2% less for their car ($450 for the average car). 25% of the savings come from purchasing at low-price dealerships affiliated with the online service. The remaining 75% stem from information provision by the online service, bargaining by the service on behalf of consumers, and cost efficiencies. A consumer receiving the mean online price does better than 65% of offline consumers, conditional on the car being purchased.

I. INTRODUCTION

The Internet is expanding rapidly into every market and many geographic locations. While much attention focuses on so-called 'new economy' businesses, an interesting aspect of the Internet revolution is the change being forced on traditional industries. The resulting threat to intermediaries such as traditional brokerage and music labels was not hard to predict. However, the impact of the Internet on some other industries was less clear. In 1995, for example, the popular press thought the Internet would not be a good sales channel for cars. After all, went the argument, consumers would always want to 'kick the tires' before buying a car. While consumers remain interested in physically inspecting a car, the Internet has nonetheless become an important complement to the car-buying process.

* We thank Melanie Webber and Craig Ima from Autobytel.com for providing us with purchase referral data and Keiko Powers from J.D. Power and Associates for help with the transaction data. We thank Ian Ayres, Meghan Busse, Chris Denove, Jean-Pierre Dube, Jonathan Feinstein, Thomas Hubbard, Ganesh Iyer, David Levine, Barry Nalebuff, Sharon Oster, Andrea Shepard, Scott Stern, Ivo Welch, an anonymous referee, and seminar participants at Harvard, the NBER E-Commerce Conference, U.C. Berkeley, and Yale Universities for helpful comments and suggestions. All errors are the responsibility of the authors. Fiona Scott Morton was partially supported by NSF grant 9810178. Florian Zettelmeyer was partially supported by the UC Berkeley Junior Faculty Research Grant.

† Author's affiliation: School of Management, Yale University, PO Box 208200, New Haven CT 06520-8200, USA.
email: fiona.scottmorton@yale.edu

‡ Haas School of Business, UC Berkeley, Berkeley CA 94720-1900, USA.
email: florian@haas.berkeley.edu

§ J.D. Power and Associates, Power Information Network, 30401 Agoura Road, Agoura Hills, CA 91301, USA.
email: jorge.silva-risso@powerinfonet.com

In 2000, for example, 54% of all new vehicle buyers used the Internet in conjunction with buying a car (J.D. Power and Associates [2000b]).

This paper investigates web-based auto retailing to understand whether the Internet has changed the product market behavior of established firms. In particular, we investigate how much and in what ways Internet car referral services affect dealer pricing of automobiles in California.

Internet car referral services are one of three types of sites related to new car purchases; the other two types are informational sites and sites that offer cars at posted prices. At the core of both informational and referral sites is detailed information about individual cars, including current market conditions and invoice pricing. The Internet has dramatically decreased consumer's cost of obtaining car-related information. In addition to being available at close to zero marginal cost, Internet information is of higher quality (more timely, more detailed, customizable) than information available from offline sources. While informational sites are only indirectly involved in the car purchase, referral sites such as Autobytel.com, Autoweb.com, and Carpoint.com establish contractual relations with dealers and pass on consumers' purchase requests or 'qualified leads' to contract dealers. Sites that offer cars at posted prices (although the actual sale is still performed through an affiliated dealership) are still in their infancy.

In this paper, we capitalize on the fact that Internet purchase referral services do not control the prices offered by the dealer. This means that prices that are offered to consumers continue to be the choice variable of dealers, whether a consumer walked into the showroom or was referred by an Internet service. In addition, the referring Internet services do not systematically learn transaction prices. To date it has not been possible, except through self-reported surveys, to assess what impact Internet referral services have on prices. The first question we answer in this paper is whether consumers who use Internet purchase referral services to buy a car pay less for an equivalent car than consumers who do not. Analyzing purchase referral data from Autobytel.com in combination with transaction data from J.D. Power and Associates, the answer is yes. Conditional on the car, consumers that submitted a purchase request pay on average 2% less than an offline customer. Of these savings, one quarter (0.5%) stem from the fact that Autobytel.com steers consumers to low-price dealerships. Conditioning on the dealership (in addition to conditioning on the car), online consumers pay another 1.5% less than offline consumers. We also show evidence that suggests that we are seeing more than just good bargainers choosing the web channel and paying the same low prices they paid before the Internet.

Secondly, we examine the level of dealer gross margins from the vehicle and from other products and services sold using the web. Dealer margins on the sale of a vehicle through Autobytel.com are significantly lower than

margins earned selling the vehicle the traditional way; prices are lower and the costs of acquiring the vehicle from manufacturers are not. Nonetheless, web-affiliated dealers may be better off working with Autobytel.com. First, they may be able to adjust their business processes to realize potentially lower cost of selling to Internet consumers and, second, online sales are likely to be new business for the dealerships.

To our knowledge this is the first study of the effect of Internet purchase referral services on car pricing using transaction data. This paper contributes to a small body of empirical literature analyzing the effect of the Internet on firms' product market behavior (Brown and Goolsbee [2000], Brynjolfsson and Smith [1999], Carlton and Chevalier [2000], Clay, Krishnan, and Wolff [2000], Iyer and Pazgal [2000]). There are a number of papers that analyze automobile pricing (Ayres and Siegelman [1995], Berry, Levinsohn, and Pakes [1995], Goldberg [1995], Goldberg [1996], Pashigian, Bowen, and Gould [1995], Verboven [1999]). Goldberg [1995] and Berry, Levinsohn, and Pakes [1995] estimate structural models of demand for automobiles. Pashigian, Bowen, and Gould [1995] investigates within-season pricing patterns for automobiles, and Verboven [1999] tries to determine whether pricing practices on base cars differ from those of cars with options. Ayres and Siegelman [1995] and Goldberg [1996] analyze race and gender discrimination by car dealers. Our paper is closest to these in that we are explicitly interested in the differential pricing introduced by dealerships. However, in contrast to all previous studies, our focus is on the Internet and how it affects the level and distribution of auto prices.

We proceed as follows: Section II describes Internet car retailing. Section III describes the data. Section IV presents the results, and Section V concludes the paper.

II. INTERNET CAR RETAILING

Informational sites do not have direct relationships with car dealers. Perhaps the best known example of a for-profit information site is edmunds.com, providing rich editorial content similar to a car magazine (reviews, long term tests, etc.) supplemented with information on exact options and invoice pricing for new and used vehicles. Sources of revenues are advertising, detailed reports on vehicles, and commissions from other Internet sites such as peoplefirst.com and geico.com.

Purchase referral sites were the first Internet companies to contract directly with car dealers. They try to capitalize on the dissatisfaction of consumers with the car buying process by shielding consumers to some degree from direct interaction with dealers and 'arming' them with information. In light of the regulatory constraints in all US states that give dealers the exclusive right to sell cars to consumers, Autoweb.com,

Autobytel.com, and Carpoint.com (to name some of the largest sites) focus on referring consumers to car dealers with which they have contracted. Autobytel.com, the first of these sites, has been in operation since March of 1995. Referral sites offer consumers detailed information about individual cars, including current market conditions and invoice pricing as well as editorial content. This information is more timely, more detailed, and much more comprehensive than the information that was available in offline publications. For example, on many of these sites consumers can select two cars and compare their features side-by-side.

At any given point a consumer may submit a free purchase request in which the consumer specifies which car she is interested in (including options and color), within which time frame she intends to purchase, and where she can be reached. While there are some small differences between referral sites, these purchase requests are generally e-mailed to a salesperson within a dealership who is assigned to responding to Internet purchase requests. Consumers are typically called back or e-mailed within 24–48 hours with a (supposedly) fixed price.[1] They can buy the car, if they wish to do so, without setting foot in the dealership until they pick up the car. The referral service sends an e-mail to the customer a few days after the referral asking whether she has been contacted by the dealer. Two weeks later another survey is sent to consumers, inquiring whether a car has been purchased and whether the consumer was satisfied with the dealership.

Referral sites send submitted referrals to dealers with whom they have contracted. Out of the approximately 22,000 dealers in the US (as of 2000), Autobytel.com contracted with 5,000 dealerships.[2] Dealers pay an annual fixed fee based on the size of the dealership which is \$1607/month on average for Autobytel.com. The closing ratio (sales/referrals) is about 13% for Autobytel.com (though lower for some of the other referral services). Dealers that are affiliated with Autobytel.com, for example, sell on average 8% of their vehicles through the service. This results in a mean payment of \$135 per sold vehicle to Autobytel.com.[3]

In exchange, dealers are assigned exclusive territories such that all customers in a specified geographic area that submit a purchase referral for a particular nameplate get referred to the same dealer. What this implies is that dealers will draw consumers to their dealership that would otherwise have purchased elsewhere. The potentially better vehicle match

[1] Note that the price offer from the dealer (by phone or email) is not a binding commitment.

[2] Autobytel.com interview. The Carpoint.com website reports they have 3,700 dealers under contract, while the Autoweb.com website reports 5,000 contracting dealers in 2000. These self reported numbers should be used with caution since each service has a different way of counting affiliated dealerships.

[3] Youngme Moon [1999], 'Autobytel.com,' HBS Case Study, and J.D. Power and Associates [2000a].

or lower price offered by the Autobytel.com dealer compensates the consumer for a potentially less convenient dealer location. If Autobytel.com dealerships are larger (which implies a greater selection of cars) and lower cost (and can thus can charge a lower price), then they will succeed in changing the dealership choice of some of their referrals. This will result in incremental sales for the dealerships and explains why dealerships are willing to pay a referral fee to Autobytel.com.

The contract usually also specifies that the sales person who responds to requests by the referral service does not sell to walk-in customers. This 'Internet sales consultant' is usually trained directly by the referral service. In addition, dealerships are usually required to compensate this sales person by units sold, not as a percentage of car gross margin. These rules are intended to ensure that consumers receive a 'no-haggle' price but are not always adhered to by dealerships.[4] Dealerships are encouraged to give referral consumers the lowest final price they would normally give a consumer on the sales floor. Purchase referral sites cannot monitor dealer performance directly. Instead they rely on customer satisfaction surveys and on their estimates of 'conversion rates,' i.e., how many referrals are sent to the dealer per sale made. For dealerships, the average conversion rate is about 8 leads from an Internet referral service to 1 sale. If conversion rate numbers are low for a dealership, or many consumers complain about the dealership, a referral site will terminate the contract with the dealer. Autobytel.com, for example, terminated over 250 dealers between 1995 and 1999.[5]

III. DATA

Our data come from two sources. The first is a major online car referral service, Autobytel.com.[6] We have obtained the purchase requests submitted by consumers on Autobytel.com during 1999, yielding slightly over 2 million observations. An observation consists of customer information, desired car, the date the request was made, the dealer to whom Autobytel.com sent the referral, and the time frame within which the consumer is interested in buying the vehicle.

[4] In a J.D. Power and Associates [2000a] survey 43% of dealerships state that the quoted price leaves room for additional discount, 41% state that the quoted price leaves no room for additional discount, and 16% say that they prefer not to quote a price until the customer comes in.

[5] See Lehman Brothers, 'Autobytel.com,' 5/12/1999 and J.D. Power and Associates [2000a] for information about conversion rates and dealer turnover.

[6] Autobytel.com had between 45 and 50% market share of online car shopping in 1999 (LA Times, 3/28/2000, 'Mergers and Acquisitions Report,' Securities Data Publishing 6/12/2000). According to J.D. Power and Associates [2000b], Autobytel.com is the most visited purchase referral site. It is visited by 33% of consumers that researched online to shop for a car, followed by Autoweb.com (18%), and Carpoint.com (17%).

The second dataset we employ comes from J.D. Power and Associates (JDPA). JDPA collects transaction data from a sample of dealers in the major metropolitan areas in the US. We have data from California dealerships, containing every new car transaction at a sample of 1,101 dealerships from January 1, 1999 to February 28, 2000. We take the extra two months of JDPA data to allow for referrals in late 1999 that result in a purchase in early 2000. Each observation in the JDPA data contains customer information, the make, model and trim level of the car, financing, trade-in information, dealer-added extras, and the profitability of the car and the customer to the dealership.[7]

We consider a match between observations from Autobytel.com and JDPA when the make and model and either the geocoded address or the phone number associated with the referral and the purchase transaction are the same. Each observation in the new dataset is a transaction from the JDPA data, augmented with the information from the Autobytel.com data if there was a match. Hence, customers who get an Autobytel.com referral but subsequently do not purchase will not be in the combined dataset. We are not interested in these customers as there is no observable transaction. If the customer purchases a car, but not from a dealership in the JDPA sample, she will not be in our dataset either. Finally, there are online referral services other than Autobytel.com. The customers in the combined dataset who are not identified as using Autobytel.com may have used one of their competitors. This strengthens our test since we will be comparing a group that used Autobytel.com to a group that may include users of competing services.

The combined dataset contains 324,936 vehicle purchases between January 1, 1999, and February 28, 2000. There are 9,545 customers in this dataset who submitted a purchase request in 1999 through Autobytel.com (2.94% of all transactions). The new variables created after the match between datasets are (1) an indicator for Autobytel.com customer (ABT) indicating that the customer who purchased the car submitted a purchase request using Autobytel.com (irrespective of whether this purchase request went to the dealer that sold the car), (2) an indicator for Autobytel.com franchise dealer ($ABT\,Franchise$) indicating that the dealer who sold the car is an Autobytel.com affiliated dealer, i.e., is under contract with Autobytel.com and receives purchase requests, (3) an indicator for same dealer ($SameDealer$) marking cases when the dealer that sold the car is the same dealer to whom the purchase request was submitted (given that $ABT = 1$).

[7] The dealer can sell the customer extras like service contracts and life insurance that make car and customer profits different.

III(i). *Car Definition and Selection*

Since consumers who use Autobytel.com may prefer different cars than those who do not, it is very important to control for the exact car that was purchased. We define a 'car' as every combination of *make and model* (e.g., Honda Accord, Toyota Camry), *body type* (e.g., convertible, coupé, hatchback, sport utility), *doors* (e.g., 2 door, 4 door, 4D Ext Cab), *trim level* (for Honda Accord, e.g., DX, EX, LX etc.), *drive train type* (e.g., 2WD, 4WD), *transmission type* (automatic, manual), *cylinders* (e.g., 4 cyl, 6 cyl), *displacement* (e.g., 3.0 liters, 3.3 liters), and *model year* (e.g., 1999, 2000).[8] This results in many different 'cars,' some of which comprise very few observations. We drop cars for which there are fewer than 500 observations, unless that car has more than 1.5% share of its segment (as defined by JDPA). In addition, because there are so few units sold, we drop the 'Large' segment. This reduces the sample size by 42% to 186,965. The dataset now contains 204 different 'cars' sold through 779 dealers.

Car segment	Obs.	%	Examples	Final Obs.
Compact	52512	14.6	Honda Civic, Toyota Corolla	33120
Large	839	0.2	Ford Crown Victoria, Chevrolet Impala	0
Luxury	40203	11.2	BMW 323i, Lexus GS400, Volvo S80	16204
Midsize	72726	20.2	Honda Accord, Toyota Camry, Oldsmobile Cutlass	43465
Pickup	61458	17.1	Ford F150, GMC Sierra 1500, Toyota Tundra	25803
SUV	78172	21.7	Jeep Cherokee, Ford Explorer, Lexus RX300	42868
Sporty	21965	6.1	BMW Z3, Honda Prelude, Mitsubishi Eclipse	8176
Van	32380	9.0	Dodge Caravan, Ford Club Wagon, Ford Windstar	17329

III(ii). *Dependent Variables*

There are many different kinds of vehicle 'prices' in the data. We define *Price* as the price that the customer pays for the vehicle, factory installed accessories and options, and dealer-installed accessories contracted for at the time of sale, subject to two adjustments. First, we subtract the *ManufacturerRebate* given directly to the consumer, if any, since this is simply cash that reduces the price of the car to the customer. Secondly, we subtract what is known as the *TradeInOverAllowance*. This is the difference between the trade-in price paid by the dealer to the consumer and the estimated wholesale value of the trade-in vehicle (as booked by the dealer). We adjust for this amount to account for the fact that dealers

[8] 46.5% of the cars in the full dataset have a transmission coded as 'N/A' rather than automatic or manual. Most cars have observations in all three groups or in just automatic and N/A. We treat the N/A group as another type of transmission, since it is likely to be composed of both automatic and standard shift cars, but we cannot tell in what mix.

may offer consumers, for example, a low price for the new car because they are making a profit from the trade-in.[9] Table III(ii) shows an example where consumers A and B are contracting with the dealer for a different new car price but end up paying the same *Price* according to our definition.

	Consumer A	Consumer B
Contract price of new vehicle	21,000	19,000
Actual cash value of trade-in	9,000	9,000
Trade-in vehicle price	10,000	8,000
TradeInOverAllowance	1,000	−1,000
Price	20,000	20,000

All factory and dealer accessories that contribute to the resale value of the car are contained in *Price*; so called 'after market options (AMOs)' such as rustproofing or wax are not included, nor are service contracts and other sources of income for the dealer.

The *VehicleCost* is the retailer's 'net' cost for the vehicle and includes the cost of accessories added by the factory and/or retailer and included in the customer's contract that add to the vehicle's book value. The measure takes into account holdback and includes transportation charges.[10] We do not have any information on the existence or extent of nonlinear, volume-based discounts.

III(iii). *Controls*

To control for time variation in prices, we define a dummy *EndOfMonth* that equals 1 if the car was sold within the last 5 days of the month. Dealers who want to meet volume targets for the month often have sales or other inducements to purchase near the end of the month. A dummy variable *Weekend* specifies whether the car was purchased on a Saturday or Sunday for the same reason. In addition, we introduce dummies for each month in the 14-month sample period to control for other seasonal effects and inflation.

We control for the number of months between when a car was sold and its introduction. This proxies for how 'hot' a car is and what the

[9] If online buyers have unobservably higher quality trade-ins than offline buyers (which we have no reason to think is true), our measure will systematically overstate the price paid by online buyers.

[10] 'Holdback' is the trade's term for additional profit to the dealer built in to the invoice for the car, but not itemized as such on the invoice. If a dealer sells a car at his invoice price, his profit on the car will equal the holdback amount on the car.

opportunity cost of not selling it is for the dealer. Judging by the distribution of sales after car introductions we distinguish between sales in the first four months, months 5–13, and month 14 and later and assign a dummy variable to each category.

Finally, we also control for the region in which the car was sold according to the definition of a region by JDPA as 'Northern California' or 'Southern California.'

IV. RESULTS

We will first describe who uses Autobytel.com and how their behavior differs from others. Next we analyze whether use of Autobytel.com alters the average price a consumer expects to pay for her car using hedonic regressions. We end with a discussion of trade-ins, dealer gross margins, and dealer costs.

IV(i). *Unconditional differences across customers and dealers*

Table I contains descriptive statistics for the entire dataset. Three percent of customers submitted a purchase request through Autobytel.com on average. Twenty-eight percent of purchase requests resulted in a sale at the referred dealer. Notice that this means that 72% of the consumers who use Autobytel.com eventually buy elsewhere and, in principle, need never have visited their referred dealership. Thirty-three percent of transactions went through dealerships who are signed up with Autobytel.com. The average car in the sample sold for $23,741 and earned the dealer almost $1700 in gross margin. There is not enough difference between Northern and Southern California to report the data separately, except that Southern Californians are more likely to purchase from their referred dealer (32% vs. 26%).

TABLE I
SUMMARY STATISTICS (ALL SEGMENTS)

Variable	Obs	Mean	Std. Dev.	Min	Max
ABT	186965	0.032	0.18	0	1
ABTFranchise	186965	0.331	0.47	0	1
SameDealer	186965	0.009	0.10	0	1
Price*	186965	23.7	8.7	6.1	101.9
VehicleCost*	186965	22.3	7.9	7.4	93.6
VehicleProfit*	186965	1.7	1.4	−4.9	13.4
TradeInOverAllowance*	53943	0.75	1.7	−10.0	18.5
EndOfMonth	186965	0.21	0.41	0	1
Weekend	186965	0.32	0.47	0	1

* Reported in thousands of dollars

TABLE II
TRANSACTION SUMMARY STATISTICS BY AUTOBYTEL.COM USE

	Obs	Mean	Median	Std. Dev.	Obs	Mean	Median	Std. Dev.
	ABT = 0				ABT = 1			
Price*	180989	23.7	22.5	8.8	5976	23.8	23.5	7.0
VehicleCost*	180989	22.3	21.2	7.9	5976	22.3	21.8	6.2
VehicleProfit*	180989	1.69	1.42	1.40	5976	1.63	1.25	1.37
TradeInOverAllowance*	52809	0.76	0	1.76	1134	0.21	0	1.18
EndOfMonth	180989	0.21	0	0.41	5976	0.23	0	0.42
Weekend	180989	0.32	0	0.47	5976	0.30	0	0.46
	ABTFranchise = 0				ABTFranchise = 1			
Price*	125071	23.2	22.1	8.1	61894	24.8	23.1	9.9
VehicleCost*	125071	21.8	20.8	7.3	61894	23.3	21.8	8.8
VehicleProfit*	125071	1.68	1.43	1.34	61894	1.71	1.40	1.52
TradeInOverAllowance*	37088	0.73	0	1.73	16855	0.79	0	1.79
EndOfMonth	125071	0.21	0	0.41	61894	0.22	0	0.41
Weekend	125071	0.32	0	0.47	61894	0.33	0	0.47

* Reported in thousands of dollars.

We expect to see a difference between customers who use Autobytel.com and those who do not if Autobytel.com users have higher incomes or education levels and therefore behave differently and buy different kinds of cars. Table II contains means and medians for a subset of variables according to whether the customer used Autobytel.com or not (upper panel) and according to whether the customer purchased at an Autobytel.com franchise dealer or not (lower panel).

Median price is higher for the Autobytel.com customers ($1,017), as we expected (since we are not conditioning on car type in any way). Mean gross margins on the vehicle earned by the dealers are very similar between the two groups; however, unconditional median gross margins are $171 higher for non-Autobytel.com customers. *TradeInOverAllowance* is lower for the Autobytel.com group, which is something we will discuss in detail later in the paper.

Table III contains demographics according to whether the customer used Autobytel.com or not. These are census data from the block group where the customer lives.[11] Unfortunately, we do not have individual buyer demographics. Autobytel.com customers come from census block groups that are more male, more professional, and higher educated. Average income in these census blocks is $67,000 rather than $58,000 for the non-Autobytel.com group. Census blocks with Autobytel.com customers also have fewer farm workers and fewer African-American residents.

[11] A block-group (approximately 700 people) is an aggregation of several census blocks. It is smaller than a census tract (approximately 4,000 people).

TABLE III
DEMOGRAPHICS BY AUTOBYTEL.COM USE

	Obs	Mean	Std. Dev.	Min	Max	Obs	Mean	Std. Dev.	Min	Max
	ABT = 0					ABT = 1				
%Professional	178162	16	9	0	100	5913	19	9	0	66
%Farmer	178162	2.6	5.8	0	100	5913	1.6	3.7	0	59
MedianHHIncome*	177469	58.3	25.6	10.5	150	5895	66.9	26.0	13.2	150
MedianAgeHeadHH	177469	45	7	13	70	5895	45	7	18	70
%CollegeGrad	177569	32	17	0	100	5899	39	17	0	95
%<HighSchool	177569	13	13	0	100	5899	9	9	0	87
%Black	177576	4.4	10	0	100	5899	3.2	6.4	0	99

* Reported in thousands of dollars.

As we would expect from the fact that 143 out of 779 dealers in our sample are affiliated with Autobytel.com, the average minimum distance between Autobytel.com franchise dealers of the same nameplate is 3.2 times larger than the average minimum distance between all dealers in our sample (80.8 vs. 25.3 miles). Table IV shows how Autobytel.com franchise dealers differ from non-Autobytel.com franchise dealers. Even looking at our restricted sample they are much larger: on average they sell 767 sample cars during our time period as compared to 360 for non-Autobytel.com franchise dealers. *ABTFranchise* dealers sell about 3.6% of their cars through Autobytel.com. Non-*ABTFranchise* dealers pick up some Autobytel.com customers who switch dealers; their percentage is lower at 1.6%. *ABTFranchise* dealers accept slightly fewer trade-ins and turn around their cars four days quicker than non-*ABTFranchise* dealers.

IV(ii). *Conditional Prices*

Table V presents the regression of the log of *Price* on an ABT indicator, 'car' dummies, month dummies, an indicator for southern California, and indicators for whether the car was sold on a weekend or at the end of the

TABLE IV
DEALER SUMMARY STATISTICS BY ABTFRANCHISE

	Obs	Mean	Std. Dev.	Min	Max	Obs	Mean	Std. Dev.	Min	Max
	ABTFranchise = 0					ABTFranchise = 1				
DealerVolume	636	360	551	1	4,280	143	767	825	11	4,496
DealerSales*	636	8,075	12,400	12	105,000	143	18,700	20,700	194	149,000
Dealer%ABT	636	1.6	2.4	0	25	143	3.6	3.0	0	13.9
Dealer%TradeIn	569	33.7	15.9	1.5	100	143	26.4	10.0	3.9	51.3
Dealer%SameDealer	636	0	0	0	0	143	2.2	2.7	0	13.1

* Reported in thousands of dollars.

month. The coefficient on the ABT indicator coefficient is −0.015 and it is significantly different from zero at below the 1% level. This means a typical Autobytel.com consumer pays 1.5% less for his or her car, a savings of $355 for the average car purchased by this group. In another regression (not reported) we include fixed effects for each dealer and the ABT indicator. In this specification the coefficient on *ABT* drops to negative 1.02%, which is a savings of $234 on the average Autobytel.com car. Again, these are conservative estimates because the comparison group contains both 'street' and Internet sales.

The second column of Table V shows the coefficients from our main

TABLE V
OLS OF LN(PRICE), TRADEINOVERALLOWANCE, AND LN(VEHICLECOST) ON REPORTED VARIABLES

	(1) log(Price)	(2) log(Price)	(3) TrOverAll.	(4) log(Price)	(5) log(Veh.Cost)
ABT	−0.015 (0.001)**	−0.015 (0.001)**	−345.72 (57.86)**	−0.014 (0.001)**	−0.001 (0.0008)*
ABTFranchise		−0.006 (0.003)*		−0.006 (0.003)*	0.002 (0.002)
SameDealer		0.002 (0.003)	−34.02 (91.37)	0.002 (0.003)	
AnyTrade				0.008 (0.001)**	
ABT*AnyTrade				0.0008 (0.0017)	
EndOfMonth	−0.003 (0.0004)**	−0.003 (0.0004)**	−62.15 (18.81)**	−0.003 (0.0004)**	
Weekend	0.002 (0.001)**	0.002 (0.001)**	−77.32 (17.17)**	0.002 (0.001)**	
SouthernCal	0.0002 (0.003)	0.0005 (0.003)	dropped	0.0002 (0.003)	0.010 (0.002)**
ModelMonth5-13	−0.0006 (0.001)	−0.0005 (0.001)	90.35 (24.64)**	−0.0004 (0.001)	−0.0002 (0.001)
ModelMonth14+	−0.006 (0.002)**	−0.006 (0.002)**	203.15 (36.46)**	−0.006 (0.002)**	0.002 (0.002)
VehicleCost	0.00004 (1.4e-06)**	0.00004 (1.4e-06)**		0.00004 (1.4e-06)**	
Constant	9.194 (0.031)**	9.195 (0.031)**	723.63 (40.35)**	9.192 (0.031)**	9.947 (0.002)**
Month Fixed Effects	Y	Y	Y	Y	Y
Car Fixed Effects	Y	Y	N	Y	Y
Dealer Fixed Effects	N	N	Y	N	N
Observations	186965	186965	55832	186965	186965
R^2	0.97	0.97	0.13	0.97	0.97

Robust standard errors in parentheses; * significant at 10%; ** significant at 5%.

specification with *ABT*, *ABTFranchise*, and *SameDealer* included. Purchasing a car from an Autobytel.com dealer saves the purchaser 0.59% (significant at the 7% level) of the cost of the car, regardless of whether the customer ever went online. Note, however, that the *ABT* coefficient is unchanged. There are two explanations for these results. One is that Autobytel.com chooses the lower-priced dealerships to partner with or, alternatively, lower-cost dealerships gain more from partnering with Autobytel.com. The second reason might be that dealerships interested in the web often sign up with more than one referral service. Thus the lower-than-average prices at these dealerships may represent sales to all online customers. However, because Autobytel.com had a large market share in 1999, one would expect the *ABTFranchise* coefficient in a regression by itself to be higher than when the ABT indicator is also included. However, the coefficient only differs by 0.0003, hardly any change at all. This favors the selection explanations above. Hence we will include the savings from being sent to an Autobytel.com dealer as part of the customer's total gain from the service. The combined savings from the two effects is \$488 per Autobytel.com customer. Since the average customer has a 33% likelihood of shopping at an Autobytel.com dealer at random (weighted by sales volume), the overall expected savings to an average customer is \$355 plus \$93 (two-thirds of 0.0059 times average price), or \$448.

SameDealer's coefficient is insignificant, indicating consumers do not systematically gain or lose by staying with the dealer referred to them by Autobytel.com. This result is likely due to a mix of consumers who choose to leave their referral: some find a better price elsewhere, others accept a higher price from a better (e.g., more conveniently located) dealer.

We run our main specification on each segment of the data separately, to see if the *ABT* coefficient is sensitive to the class of car being sold. The results are reported in the appendix. There is considerable variation across segment in the price discount for online consumers. Consumers who purchase a pickup truck on the Internet pay 2.9% less than offline consumers. Both 'sporty,' 'compact,' and 'midsize' customers pay about 2% less. In contrast, the discount for luxury and van customers is below 1%, and the SUV discount is below average also.[12]

Figure 1 represents the distribution of 'bargaining outcomes' after we control for the characteristics of the car and transaction in the regression. We use the specification from Table column 2 (controls, *ABT*, *ABTFranchise*, and *SameDealer* indicator). The upper graph is a histogram of the residuals after we add the coefficient estimates times the variables for

[12] The anecdotal explanation for the 'luxury' result is that due to the economic boom in 1999, demand was outstripping the supply of these cars so that no matter what channel a customer used, the price was not discounted significantly. A similar story might hold for the SUV segment also.

ABT, *ABTFranchise*, and *SameDealer* and omit the residuals of customers who used Autobytel.com. The lower graph shows the same measure only for consumers who used Autobytel.com. The very small residuals represent consumers who received large discounts, while the long right tail represents customers who paid more than others. The vertical line on the figure shows the mean of the lower distribution and indicates where the average Autobytel.com customer appears in the 'bargaining outcome' distribution. The figure shows that a consumer receiving the mean Autobytel.com price does better than 64.6% of the customers who do not use the service. The distribution of 'bargaining outcomes' for Autobytel.com users is pictured in the lower part of the figure. It is clear that this distribution has a lower mean and smaller variance than the other.

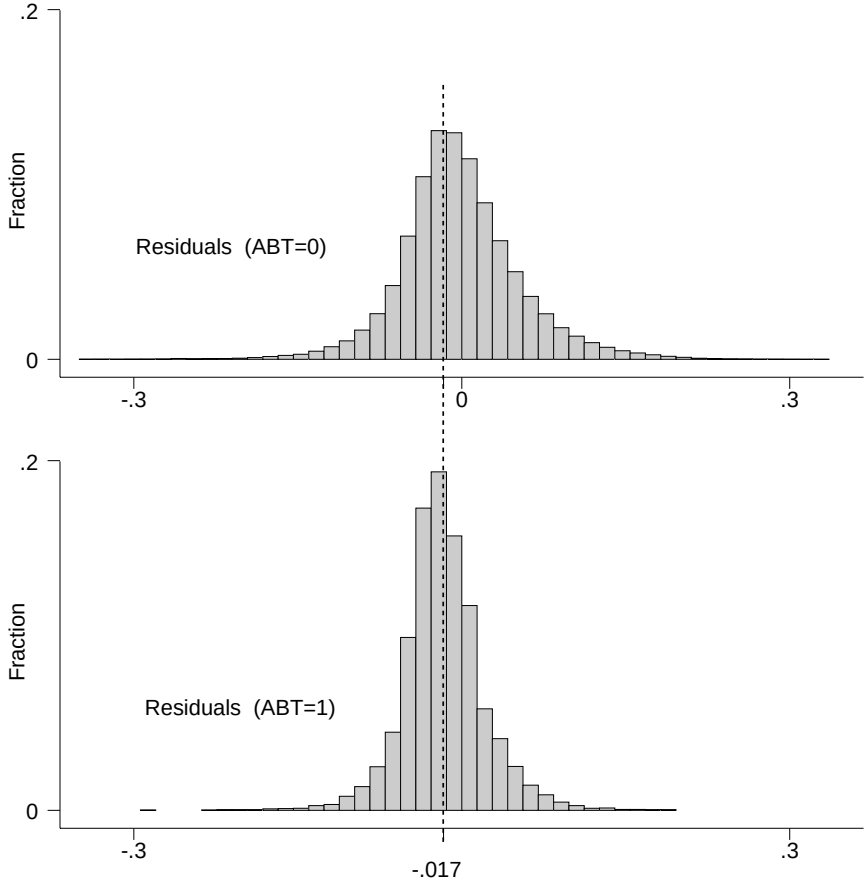

Figure 1
Histogram of residuals from Table V column 2 for ABT= 0 and ABT = 1

Our results may be driven by a strong selection effect: using the Internet does not change the price any consumer pays for a car, but good bargainers are more likely to use the Internet. While a formal treatment of selection with instrumental variables is a topic for future research, we try a simple test to look for patterns consistent with selection. We divide the Autobytel.com customers by quarter and estimate a separate ABT coefficient for each quarter of the sample.[13] We assume that because market growth has been so large, Autobytel.com customers at the beginning of 1999 were less representative of the general population than those at the end of that year, and therefore any selection effect should be falling over time. Declining coefficients would provide evidence for a strong selection effect. Instead we find that the coefficients are nearly constant over time, with a small increase in absolute value in the last quarter of 1999. This is not what we would expect if our results were driven purely by a selection effect.[14]

IV(iii). *Trade-in Pricing, Dealer Gross Margins, and Financing*

An interesting difference between Autobytel.com customers and others appears in the distribution of gross margins between the new vehicle and the trade-in. Autobytel.com customers receive less for their trade-in, although they end up better off overall. As reported in column 3 of Table V, dealers offer on average $346 less for an Autobytel.com customer's old vehicle, conditional on the value of the vehicle (this specification does not contain car dummies, but does contain dealer dummies). Therefore, the average trade-in customer in our dataset sees an Autobytel.com contract price that is about $800 less than non-Autobytel.com offline prices, but receives about $350 less for his or her trade-in, making the net savings about $450. An interesting industry fact is that dealers pay on average $770 above the 'actual cash value' (wholesale price) of the trade-in. This does not represent a loss to the dealer because average new car prices are adjusted up to compensate, resulting in cross-subsidization from the new car to the trade-in. (For example, Goldberg [1996] analyzes contract prices and finds that customers with trade-in vehicles pay on average $600 more for their new car.) Since an online consumer is more likely to be comparing

[13] We exclude the last quarter of the sample because it has few observations.

[14] In addition, we try including in our specification the distance between an Autobytel.com consumer's home zip code and the zip code of the referred dealer. The idea is that consumers who are further away require bigger price discounts to be induced to buy. If Autobytel.com is bargaining on behalf of consumers, we should see this pattern. Instead the distance variable has a small positive coefficient (not reported). This is likely due to the correlation of distance with measures of market structure or geography; so while the results do not support bargaining by Autobytel.com, it is difficult to interpret them. We plan to explore this relationship in future work.

prices before the trade-in is discussed, dealers may be unwilling to quote a high price for the new car in order to be able to subsidize a trade-in later.

It is also interesting to note that consumers who trade in their old vehicles pay on average 0.77% more for their new vehicle than consumers that have no trade-in (see the dummy *AnyTrade* in column 4 of Table V). Recall that the dependent variable *price* above is adjusted for the *TradeInOverAllowance*. The interaction coefficient in the table shows that the price premium that Autobytel.com customers pay when they trade in their old car is no more or less than the premium paid by others whose transactions include trade-ins. It appears that a small overall price premium is charged to both online and offline consumers who take advantage of the convenience of trading in their old car to the dealer.

Column 5 of Table V shows that the cost of acquiring the vehicle is only marginally lower ($33) for Autobytel.com customers and non-Autobytel.com customers. In addition, Autobytel.com franchises face the same cost of vehicles than non-Autobytel.com franchises. If Autobytel.com sales are lower cost than conventional sales, this must be due to lower overhead costs such as salespeople's salaries and real estate rents.

What do these estimates say about dealership profitability? It depends crucially on both how many of the Autobytel.com sales are incremental and the dealership's costs of selling over the Internet. If a dealer can use the low prices he offers through Autobytel.com to capture share from other dealers, the lower margins on those cars are simply additional profit, not losses. In addition, it may be cheaper for a dealer to sell to Internet consumers.[15] The difference between the average total cost of a traditional sale in our data ($1,575) and the average cost for one efficient Autobytel dealer we could get data from ($940) is over $600. Comparing this number to our estimate of the expected price drop for Internet sales of between $300 and $500 provides one potential reason for some dealers to sign up with Autobytel.com: their margins may increase. However, a survey (J.D. Power and Associates [2000a]) finds that 51% of dealerships report that it takes *more* man-hours to complete a sale with Internet customers from an Internet referral service. The report states that this may be because the 'invoice up' style of negotiating is unfamiliar to many dealerships who are used to 'MSPR down' negotiations. The study also finds that dealers that have more experience with referral services are much more likely to

[15] However, our evidence is mixed. Autobytel.com provided us access to the accounting data from one dealer in the midwest who breaks out his Internet sales separately. We compared his costs (phone, computer, managerial time, office supplies, delivery charges, etc.) to average dealer costs from the National Auto Dealers Association (NADA) and our estimates. We use the profit from NADA of $198 and the JDPA median gross margin of $1,773 to infer average cost.

report that Internet customers are more enjoyable to work with (and presumably more profitable) than traditional customers. This suggests that dealerships need to adjust their business processes to realize potentially lower cost of selling to Internet consumers.[16] These calculations show that referral services could be beneficial to affiliated dealers, especially if they gain incremental sales. Non-affiliated dealerships are likely to be worse off as a result of referral services.

V. CONCLUDING REMARKS

This paper investigates the effect of the Internet referral service Autobytel.com on retail auto prices and dealer profits. We show that this Internet referral service has changed the product market behavior of dealerships: they offer lower prices to consumers who arrive via the Internet than to other consumers. We capitalize on the fact that Internet purchase referral services do not control the prices offered by the dealer. This means that prices that are offered to consumers continue to be the choice variable of dealers, whether or not a consumer walked into the showroom or was referred by an Internet service.[17]

We show that conditional on the dealer, consumers that submitted an online purchase request pay about 1.5% ($355) less than other consumers. Including the effect of being sent to a low-price dealer chosen by Autobytel.com, the difference increases to 2% (or $448 for the average car when accounting for the probability of randomly visiting a low-price dealership). In the current paper we are unable to distinguish among several potential causes of the lower prices: better consumer information, bargaining by the referral service, and potentially lower costs of selling online. All are likely to contribute to lower online prices. Our results imply (and we show) that dealer gross margins on the sale of a vehicle to a customer with an Internet purchase referral are significantly lower than gross margins earned selling the vehicle the traditional way. Nonetheless, web-affiliated dealers may be better off from working with Internet purchase referral services. This is because dealerships may be able to adjust their business processes to realize potentially lower cost of selling to Internet consumers, and online sales are likely to be new business for the affiliated dealerships. The evidence we present here suggests that a strong selection effect (good bargainers use Autobytel.com) may not be driving our finding of an Autobytel.com discount. Future work will formally distinguish between a causal role for Autobytel.com and a selection effect.

[16] The dealership described above may have done just that.

[17] We expect dealer pricing to continue as long as manufacturers rely on dealers to estimate demand and to bear the risk of selling their inventory.

APPENDIX

TABLE A
REGRESSIONS OF LN(PRICE) ON REPORTED VARIABLES BY JDPA SEGMENT†

JDPA Segment	Compact	Luxury	Midsize	Pickup	SUV	Sporty	Van
ABT	−0.021	−0.003	−0.019	−0.029	−0.013	−0.019	−0.009
	(0.002)**	(0.002)+	(0.002)**	(0.005)**	(0.001)**	(0.003)**	(0.002)**
ABTFranchise	−0.005	0.00004	−0.006	−0.012	−0.007	−0.010	−0.002
	(0.005)	(0.005)	(0.005)	(0.005)*	(0.003)*	(0.004)*	(0.003)
SameDealer	0.003	−0.008	0.003	0.009	0.006	0.014	−0.002
	(0.003)	(0.004)+	(0.004)	(0.008)	(0.003)+	(0.007)+	(0.004)
Observations	33120	16204	43465	25803	42868	8176	17329
R^2	0.86	0.97	0.88	0.95	0.95	0.97	0.88

Robust standard errors in parentheses.
+ Significant at 10%; * significant at 5%; ** significant at 1%.
† Unreported are *EndOfMonth, Weekend, SouthernCal, VehicleCost, Constant*, car and month fixed effects (see the Journal's editorial Web site for complete results).

REFERENCES

Ayres, I. and Siegelman, P., 1995, 'Race and Gender Discrimination in Bargaining for a New Car', *American Economic Review*, 85(3), pp. 304–321.

Berry, S., Levinsohn, J. and Pakes, A., 1995, 'Automobile Prices in Market Equilibrium', *Econometrica*, 63(4), pp. 841–890.

Brown, J. R. and Goolsbee, A., 2000, 'Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry', Mimeo, Harvard University University of Chicago, GSB.

Brynjolfsson, E. and Smith, M., 1999, 'Frictionless Commerce: A Comparison of Internet and Conventional Retailers', Mimeo, MIT, Cambridge, MA.

Carlton, D., and J. Chevalier (2000): Free Riding and Sales Strategies for the Internet', Mimeo, University of Chicago, Chicago, IL.

Clay, K., Krishnan, R. and Wolff, E., 2000, 'Pricing Strategies on the Web: Evidence from the Online Book Industry', Mimeo, Carnegie Mellon University, Pittsburgh, PA.

Goldberg, P. K., 1995, 'Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry', *Econometrica*, 63(4), pp. 891–951.

Goldberg, P. K., 1996, 'Dealer Price Discrimination in New Car Purchases: Evidence from the Consumer Expenditure Survey', *Journal of Political Economy*, 104(3), pp. 622–654.

Iyer, G. and Pazgal, A., 2000, 'Internet Shopping Agents: Virtual Co-Location and Competition', Mimeo, University of California, Berkeley, Berkeley, CA.

J.D. Power, and Associates 2000a, 'Dealer Satisfaction with Online Buying Services', Management report, Agoura Hills, California.

J.D. Power, and Associates 2000b, 'New Autoshopper.com Study', Management report, Agoura Hills, California.

Pashigian, P. B., Bowen, B. and Gould, E., 1995, Fashion, Styling, and the

Within-Season Decline in Automobile Prices', *Journal of Law and Economics*, 38(2), pp. 281–309.
Verboven, F., 1999, 'Product Line Rivalry and Market Segmentation–With an Application to Automobile Optional Engine Pricing', *Journal of Industrial Economics*, 47(4), pp. 399–425.