SEARCH AT WHOLESALE AUTO AUCTIONS*

DAVID GENESOVE

Wholesale trade in used cars is conducted by ascending bid auctions, with sale subject to the seller's acceptance of the winning bid. *One out of three times* the seller rejects the bid, and no trade takes place. I model the seller's decision as the outcome of search, and thus determined by the winning bid distribution, and a reported retail price. This market is ideal for testing search theory since all offers, whether or not accepted, are observed. Qualitative predictions of the theory, in particular the role of the variance, are confirmed. The quantitative results are more ambiguous.

One out of every three times, seller and buyer fail to trade at wholesale auto auctions. Instead, the seller exercises his right to refuse the winning bid. This paper shows how search theory can explain the sale rate and its variation across different types of automobiles.

Few readers will find intrinsic interest in wholesale auto auctions. The value to studying them lies instead in the opportunity to observe all offers to trade in an environment that neatly matches that envisaged by the standard search model. A dealer who brings a car to the auction is offered a price (the winning bid of an oral, ascending bid auction) which he can accept or reject. If he rejects it, he can return the next week to receive a new one. Because the set of bidders varies from week to week, and each bidder's needs also change, the winning bid is initially uncertain. Finally, bringing a car to the auction is costly; it either consumes the dealer's time or prevents him from bringing a different car instead. Section I of this paper describes auto auctions in further detail, and Section II reviews the basic search model.

Having both accepted *and* rejected offers allows me to test search theory in new ways. Previous search studies have observed the duration of search but not unaccepted offers. They have been effective at studying the impact of benefits and costs that flow over time during search, such as unemployment insurance benefits for job seekers or rental income for home owners. My data contain unaccepted bids but do not follow searches over time. Accordingly,

*This is a revision of a chapter of my dissertation, "Coconuts, Lemons and Pies: Search, Adverse Selection and Bargaining at Wholesale Auto Auctions." Funding from Princeton University's John M. Olin Program for the Study of Economic Organization and Public Policy, the Sloan Foundation, and the Social Sciences and Humanities Research Council of Canada is gratefully acknowledged. I am especially grateful to Mr. Clyde Hilwig of Automotive Market Report for allowing me access to the early editions of that publication.

The Quarterly Journal of Economics, February 1995

I have the opportunity to test the implications of the stochastic part of search theory. In particular, I can examine how shifts in the mean and variance of the offer price distribution affect searchers' acceptance decisions. These predictions are set forth in Propositions 1 and 2.

I undertake a series of increasingly more structural investigations. First, I examine the reduced-form probit for bid acceptance (Section IV). This tests whether (for example) sellers more often reject the winning bid when offering a make and model whose winning bid variance is high. They do. To get closer to the quantitative implications of the theory, I then condition the seller's acceptance decision on the winning bid as well (Section V). Doing so identifies the seller's reservation price. We see that a higher winning bid variance increases the reservation price, which is the principal stochastic implication of search theory. Finally, I estimate the fully structural model implied by the definition of the optimal reservation price and a distributional assumption on search costs (Section VI). This model is shown to perform worse than the less structural model.

The seller's problem at a wholesale auto auction does differ from the search problem in one significant way. The alternative to accepting the winning bid at the auction is not merely to wait until the next auction is held. It is to return the car to the dealer's lot, and perhaps sell it to a consumer before returning to the auction. This suggests a model in which the seller faces alternating search opportunities and must form a reservation price for each one. Section III shows that the optimal auction reservation price can be expressed as the solution to a search problem in which the seller faces the offer distribution at the auction only, but where the cost of search and the discount rate are amended to reflect the retail opportunity. Unfortunately, information on trading opportunities with consumers is incomplete. Average transacted retail prices are recorded in used car guides, but no source has information on unconsummated trades. Consequently, I treat the probability that the seller will consummate a trade on the retail lot as a parameter, common to all cars. The average retail price is included as a regressor in each of the probits, as the model in Section III suggests.

Aside from this complication, these auctions provide a convenient testing ground for search theory. Comparing the markets for different types of cars holds constant several potentially complicating factors. The mechanism of trade is the same for each car. The set of potential sellers and buyers, which is the population of new and used car dealers, is also the same.[1] The dealers are frequent participants in the auction, so that rationality, knowledge of the underlying distributions, and risk neutrality are all reasonable assumptions.

I emphasize that this paper examines the behavior of the seller at the auction, not the bidders. It does not test auction theory. Quite the contrary, it makes use of the structure of the auction to identify demand. It is because the seller acts only once the buyers have bid that the entire offer distribution is observed.

The data are extracted from several issues of Automotive Market Report (AMR) in the spring and summer of 1951, and include attributes (make, model, model year, body style, options, a quality assessment, and usually mileage), winning bid, and an indicator for sale or no-sale, for all cars consigned at auctions that AMR's reporters visited. The data are described more fully in Appendix A.

I. The Wholesale Auto Auction

The following description of a wholesale auto auction is based on visits to three different auction houses in 1989 and 1992. Discussions with various auction owners and Lawrence [1984] suggest that the manner of trade has not changed over the years.

Wholesale auto auctions,[2] which on the buyers' side are limited to dealers and on the sellers' side to dealers and large fleet owners, serve mainly as a means by which dealers can adjust the composition of their stock of used cars. Having a well-balanced inventory of cars is viewed as good business practice in the used car industry. The auction provides a market in which a dealer can, in effect, trade one car for another, and thereby transform the portfolio of cars received as trade-ins to one nearer to his retail needs. These markets are sufficiently thin relative to the heterogeneity of the cars that the winning bid will vary from one visit to the next.

Bidding on several different cars is conducted simultaneously. Cars are lined up in lanes, at the end of which is an auction block,

1. The retail public is explicitly excluded from these auctions.

2. Wholesale auctions are the principal means by which used cars are traded among car dealers. The 1989 National Automobile Dealers Association Economic Survey reported that, of used cars obtained by new car dealers from other than consumers or noncar dealer firms, 60 percent were purchased at the auction. There are currently almost 300 auctions in the United States and Canada, with yearly sales of six and one-half million vehicles [Hart 1989; National Auto Auction Association 1989].

where the auctioneer stands, with the seller beside him. Before the bidding on the previous car has concluded, the next car is driven up to the auction block. Dealers, some of whom have had a cursory look at the car while it was in the lane or, earlier, parked outside, now examine it more carefully. They open the hood to peer at the engine and listen to it, examine the body of the car, and may take a look inside the passenger area as well. Mileage is chalked on the car or announced, and options are likely to be chalked as well.

The auction lasts about a minute and a half. It is fundamentally an oral, ascending bid (English) auction, but there is a twist. And that is that the auctioneer begins at a high price,[3] and then works his way down until a bid is made. So in its first part, it *sounds* like a Dutch auction, with the called-out prices declining; however, the first bid made is not the winning bid, but the signal for the bids to start increasing.[4] After that, bidding proceeds in the usual manner, with the auctioneer calling out higher and higher values. The auction ends when the auctioneer has solicited a bid that no other bidder is prepared to exceed (signified, in lieu of a gavel, by the whipping of a rubber tube across the table).

With a bid[5] in hand, the auctioneer then turns to the seller and asks for the seller to respond—an agreement to sell the car at that price, or a rejection of the bid. Rejection of the bid is not necessarily the end of the matter, as the seller and the winning bidder may bargain over a final price, a process that is generally mediated by the auctioneer. Presumably, the bidder and seller bargain because the bid did not exceed the opportunity cost of the seller's parting with the good or because the seller is attempting to expropriate part of the bidder's surplus, which, in this sort of auction, we know to be the difference between the valuation of the highest and second highest bidders. Figure I is a schematic representation of the entire bidding-bargaining process. Genesove [1992] discusses the postbidding bargaining.

In principle, the opportunity to the seller to bargain, rather than simply to reject the winning bid transforms what would otherwise be a simple search problem into a more complicated mechanism. From the econometric standpoint, if the final price is

3. The auctioneer's initial price almost always exceeds the winning bid. What effect it has on the subsequent bidding is an open question. One auction official, otherwise quite forthcoming about the workings of the auction, avoided discussion of the initial price, aside from describing its choice as an important part of the auctioneer's art.

4. This method is apparently common in antique auctions as well.

5. Where the meaning is clear, I shall refer to the winning bid as the bid.

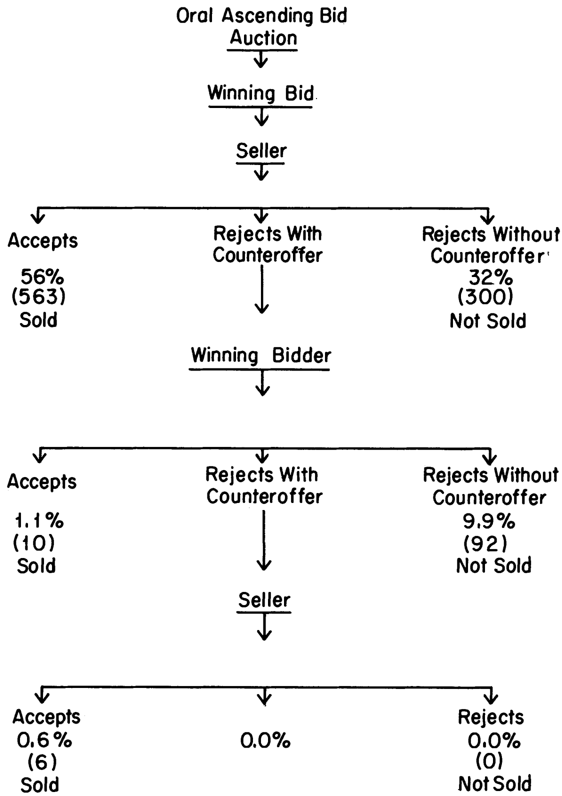

FIGURE I

determined by bargaining, it would reflect the seller's reservation price in part and so not be exogenous, and, hence, will not identify the reservation price. However, in practice, postbidding bargaining is rarely successful. Unfortunately, the 1951 data fail to record when bargaining took place; but data I collected over several weeks in the summer of 1989 (see Genesove [1993]) substantiate the claim. The percentages displayed at each node in Figure I indicate the fraction of the 1989 consigned cars for which the process ended there. Although postbidding bargaining follows one out of every four rejections of the winning bid, there are only sixteen cases (1.7 percent) in which the seller rejected the winning bid and yet the car was nevertheless sold.

The very rarity of successful bargaining suggests that bargaining is inefficient. If so, it might be argued that buyers would overbid in order to clear the seller's reservation price. But such a buyer must then believe that a seller would accept from the auction a winning bid that the seller would have rejected, if the buyer had offered it in bargaining following a lower winning bid. That seems an unlikely belief. In any case, bias introduced by buyers bidding according to their conjectures about the seller's reservation price would operate much like the bias that follows from unobserved quality differences, which is discussed in Section V of this paper. For these reasons, I ignore the possibility of bargaining in the rest of the paper.

For our purposes, the most notable aspect of wholesale auto auctions is the sale rate. Only 63 percent of the consigned cars in the 1951 data[6] were sold. This figure is by no means out of line with sale rates at other auto auctions at other times. As Figure I showed, only 58 percent of the consigned cars in the 1989 data were sold. From September 1970 to August 1971 the average sale rate for auctions reported in AMR, was 68 percent, and never exceeded 74 percent, nor fell short of 58 percent. These rates compare with typical auction sale rates of 67 percent for Impressionist paintings, 90 to 95 percent for wine [Ashenfelter 1989], and 92 percent for fish [Frappier 1992].

Definite patterns in the sale rate across different types of cars are evident as well. Bids on more expensive models are less likely to be accepted than bids on cheaper models. Figure II plots the average bid against the fraction sold for each of 43 models (e.g., Plymouth Deluxe, Studebaker Champion). It is clear that the higher is the average bid (among all bids, both accepted and rejected), the smaller is the fraction sold. The results are similar when cars are grouped by model year, as in Table I, which indicates that, beyond the first year, the sale rate tends to increase with the age of the car. However, when cars are ordered by AMR reporters' rating of their physical condition, as in Table II, we see that "better" cars are more likely to sell.[7]

6. Cadillacs are dropped from the data set. They are clearly outliers, with an average bid in 1951 of $1980, an amount one and a half times as high as the next highest make, and a variance of almost $600,000, or three times as large as the next largest variance. At 50 percent the sale rate is the lowest of all makes. Much of the work here is concerned with distinguishing between the effect of the variance and that of the mean on the sale rate. Clearly, any contribution that Cadillacs might make to distinguishing between the two effects would rely exclusively on the chosen functional form.

7. The 1989 data confirm these conclusions. In fact, the regression of the sale rate on the average bid (in 1951 dollars) for makes is almost exactly the same for the two years. The fraction sold increases with the age of the car in the 1989 data. There are no quality ratings for the 1989 data.

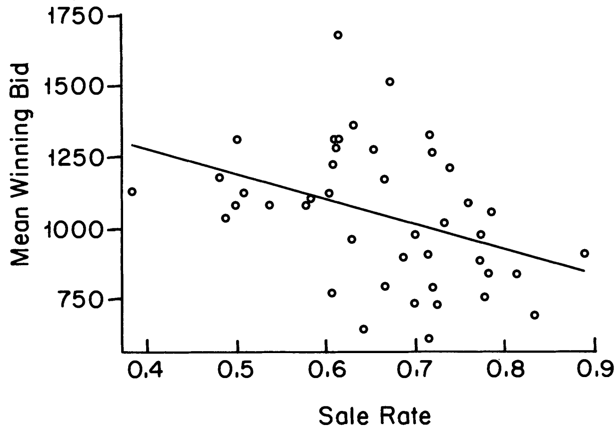

FIGURE II
Mean Bid versus Sale Rate, by Make

II. SEARCH THEORY

This section outlines the simple model in which the seller can offer the car on the wholesale market only. It uses the basic search model [Mortensen 1986]. Section III incorporates the retail market as well.

A seller has a single automobile that he wishes to sell. Each period that the seller continues to search, he incurs a cost of search c and receives an offer y (the winning bid), which he can either accept or reject. If he rejects the offer, he continues to search in the next period. He discounts the future at rate ρ. The seller is assumed to know the distribution of winning bids for a car of his type, H. He adopts a reservation-price strategy, which Kohn and Shavell [1974] have shown to be optimal; that is, he sets a price R, above which he will accept any offer, below which he will refuse and continue to search. The value of search using reservation price R, $V(R)$, satisfies

$$V(R) = (1 + \rho)^{-1}\left\{\int \max\,(V(R),y)\; dH(y) - c\right\}. \tag{1}$$

The optimal choice of a reservation price is the unqiue R that satisifes $V(R) = R$. Totally differentiating (1) shows this to be true, but it is also follows from stationarity: if R is to be used in the future, then $V(R)$ is the opportunity cost to accepting a bid today, and thus today's optimal reservation price as well.

TABLE I
SALE RATE BY MODEL YEAR

Year	Number	Percent sold
1951	115	69
1950	827	60
1949	920	68
1948	498	65
1947	507	74
1946	342	71

TABLE II
SALE RATE BY CONDITION

Condition	Number	Percent sold
Sharp	388	86
Clean	424	75
Good	195	54
Fair	132	64
Rough	21	48

Assume that the shape of the distribution of winning bids is the same for all cars. The distributions differ only in a location parameter and a scale parameter, which, without loss of generality, may be assumed to be the mean, μ, and standard deviation σ. Thus, if z is the vector of attributes of the car, the distribution of winning bids is

$$H(y|z) = H_0([y - \mu(z)]/\sigma(z)) \tag{2}$$

for all z and some common, though not necessarily known, distribution H_0.[8] For most of this paper, no assumption is made about the form of H_0.[9]

8. If the number of bidders N is sufficiently large, and bidders have private valuation, the winning bid will be approximately distributed in this manner. For in an English auction, private valuations imply that

$$H(\cdot|z) = NF(\cdot|z)^{N-1}[1 - F(\cdot|z)] + F(\cdot|z)^N,$$

where F is the distribution of a single bidder's valuation and N is the number of bidders. When N is large, and if the limit of the inverse hazard of $F(\cdot|z)$ is the same for all z, $(y - \mu)/\sigma$ is approximately distributed as $Q(\cdot) - [Q(\cdot) \log Q(\cdot)]$, where $Q(\cdot)$ is one of the three extreme-value types and μ and σ are functions of N and F [Reiss 1989, pp. 154–61].

9. This is in contrast to previous empirical search papers in which the lack of information on unaccepted offers has required that the form of the distribution be specified [Finn and Heckman 1982].

This assumption implies, in turn, that the reservation price may be written as

$$R(z) = \mu(z) + \sigma(z)r([c + \rho\mu(z)]/\sigma(z)), \tag{3}$$

where $r(x)$ is the optimal reservation price for a zero mean, unit variance, and search cost equal to x [Balvers 1990]. The probability of sale is the probability of receiving a bid greater than R, and so equals $1 - H(R) = 1 - H_0(r)$. Note that

$$\frac{\partial r}{\partial x} = -\frac{1}{1 + \rho - H_0(X)} < 0. \tag{4}$$

PROPOSITION 1. An increase in the variance (σ^2) increases the reservation price (R) and decreases the probability of sale [Balvers 1990]. A one-dollar increase in the mean (μ) increases the reservation price by less than one dollar and increases the probability of sale [Mortensen 1986].

An increase in the variance, being here an increase in scale, is a general inflation of the deviation from the mean of every possible winning bid. If the seller were to accept all bids, the value of search would be unchanged; but because the seller accepts only bids exceeding the reservation price, the value of search, given the original reservation price, is increased, and so the optimal reservation price will increase as well.

An increase in variance has both a direct and an indirect effect on the probability of sale. The indirect effect acts through the increase in the reservation price, which further truncates the interval of accepted bids. This decreases the probability of sale. But the increase in spread will also directly affect the probability of a bid originating above the initial reservation price. Although the sign of the direct effect can be either positive or negative (according to whether R is greater than or less than μ[10]), the sum of the two effects on the probability of sale is negative [Balvers 1990].

The mean bid plays only a supporting role. If the reservation price were to fully match a dollar increase in the mean, it would leave the expected duration of search unchanged, while increasing the expected accepted bid by a dollar. Since future receipts are

10. The condition that the direct effect be negative, $R < \mu$, may be restated as $1 - H(R) > 1 - H(\mu)$, which states that the probability of sale exceeds the probability that a winning bid is greater than the mean winning bid. In our data, however the automobiles are aggregated, this condition is almost always satisfied. Thus, not only is the sum of the two predicted effects negative, but so, too, are the components.

discounted, this would increase the present value of search by less than a dollar. As $V(R)$ is concave in R, the optimal reservation price must therefore rise by less than a dollar. However, it is clear from (3) and (4) that unless the discount rate is very high (and as auctions are held weekly that is unlikely), the increase in the reservation price will be close to the full dollar, and the effect on the probability of sale small [Mortensen 1986].

III. Sale of the Lot

The alternative to selling the car at the auction is not merely to wait until the next consignment opportunity. Rather, it is to return the car to the dealer's lot, and perhaps sell it to a consumer before the next visit to the auction. If the range of offers (final prices net of any reconditioning and other costs) in the consumer market lies everywhere below the value of search at the wholesale market alone, nothing is lost by omitting the option of sale to a consumer. But if a consumer might make an offer above that value, then the overall value of continued search at the auction market is understated. It is tempting to argue that the seller has revealed a preference for the auction over the lot by bringing the car to the auction, but implicit in that argument is a restriction that the seller must choose between the two. Rather, he might bring the car to the auction, while entertaining offers from consumers between auction visits, until the car is sold at one of the two venues.

Where between each visit to the auction the seller receives an offer,[11] drawn from distribution M, from a consumer visiting the dealer's lot, the pair of optimum reservation prices at the auction (R) and the lot (L) is the unique pair (R,L) such that

$$R = (1 + \rho_0)^{-1}\left\{\int \max(L,x)\, dM(x) - c_0\right\} \tag{5}$$

$$L = (1 + \rho_0)^{-1}\left\{\int \max(R,y)\, dH(y) - c_1\right\}, \tag{6}$$

where ρ_0 is the interest rate, c_0 is the cost of moving the car from the auction to the lot, and c_1 is the cost of bringing the car to the auction. Substituting (6) into (5) yields

$$R = (1 + \rho)^{-1}\left\{\int \max(R,y)\, dH(y) - c\right\}, \tag{7}$$

11. No offer can be interpreted as an offer of price zero.

where

$$c \equiv c_0 + (1 + \rho_0)c_1 - \phi\mu_L$$
$$1 + \rho \equiv (1 + \rho_0)^2/1 - \lambda$$
$$\mu_L \equiv \int_L x\, dM(x)/1 - M(L)$$
$$\lambda \equiv 1 - M(L)$$
$$\phi \equiv (1 + \rho_0)\lambda/1 - \lambda.$$

That is, λ is the probability of sale at the lot, and μ_L is the average selling price on the lot. Equation (7) expresses the optimal reservation price strategy at the auction as if future sale at the auction were the only alternative, with the search cost and discount rate reinterpreted to reflect the opportunity for retail sales. I will treat λ as common to all car types, and thus an (unidentified) parameter, and μ_L as data. This approach is ad hoc and far from ideal, but it is necessitated by the limited availability of data on the retail market. Published average retail prices exist, but there is no information on the likelihood of sale on the lot.

It will be useful to note that

$$\begin{aligned} c + \rho\mu &= \{c_0 + (1 + \rho_0)c_1\} + \rho\mu - \phi\mu_L \\ \rho &> \phi, \quad \text{when } \rho_0 > 0 \\ \rho &= \phi, \quad \text{when } \rho_0 = 0. \end{aligned} \tag{8}$$

Proposition 1 remains true. In addition, we have the following:

PROPOSITION 2. A one-dollar increase in the average retail price (μ_L) increases the auction reservation price by less than one dollar and decreases the auction probability of sale. Furthermore, a simultaneous increase of one dollar in both the mean wholesale bid and the average retail price increases the auction reservation price by less than one dollar and increases the auction probability of sale.

Increases in μ_L are decreases in c, and thus increase the reservation price [Mortensen 1986] and, since the bid distribution is unchanged, decrease the probability of sale. This is intuitive: increased opportunities elsewhere makes sale at the auction less attractive. As for a simultaneous increase in μ and μ_L, consideration of (3) and (8) reveals two effects. The increase in μ increases the reservation price one for one through the first term on the right-hand side of (3). Together, μ and μ_L increase $c + \rho\mu(z)$ by $\rho - \phi$; this decreases the reservation price by $(\rho - \phi)/(\rho + 1 - H_0(R))$, which is nonnegative but less than one, and is zero if and only if ρ_0 is zero. This too seems reasonable: when the mean opportunity at both the auction and elsewhere increase equally, the seller's optimal sale rate remains essentially unchanged, unless interest rates are high.

Table III summarizes the two propositions. It contains predictions on both the probability of sale (unconditional on the bid) and the reservation price. Section IV tests the first set of predictions by estimating a probit when the dependent variable is an indicator for sale, and the regressors include a predicted mean, predicted variance, and average retail price only. Section V tests the second set of predictions by conditioning on the bid as well, thus identifying the reservation price.

IV. The Unconditional Sale Rate

As no two used cars are alike, it is necessary to parameterize the mean and variance. The following three equations specify the model:

MEAN:

$$E[y_i | z_i] = z_i\beta \tag{9}$$

VARIANCE:

$$E[(y_i - z_i\beta)^2 | z_i] = \exp(z_i\alpha) \tag{10}$$

PROBABILITY OF SALE:

$$E[I_i | z_i] = \Phi(\gamma_1 + \gamma_2(z_i\beta) + \gamma_3(z_i\alpha) + \gamma_4\mu_{Li}), \tag{11}$$

TABLE III
EFFECT ON UNCONDITIONAL PROBABILITY OF SALE AND RESERVATION PRICE OF A ONE-DOLLAR INCREASE IN MEAN, VARIANCE, AND RETAIL PRICE

	Probability of sale		Reservation price	
	$\rho_0 > 0$	$\rho_0 = 0$	$\rho_0 > 0$	$\rho_0 = 0$
Mean bid (μ)	+	+	+	+
Variance (σ^2)	−	−	+	+
Retail price (μ_L)	−	−	+	+
Mean bid + retail price	+	0	+	1

where y_i is the winning bid on the ith car, z_i is a row vector of characteristics of the car, I_i equals one if the car is sold, and zero otherwise, and Φ is the standard normal distribution.

Equations (9) and (10) specify the mean and the log-variance as linear functions of the attributes (model year, model, body style, mileage, various options, physical condition, and auction site). The exponential form of (10) guarantees a positive variance. The right-hand side of (11) is an approximation to the unconditional probability of sale. In general, the probability of sale is determined by the distributions of the bid and the reservation price. The optimal search theoretic reservation price is, in turn, determined by the bid distribution, the discount rate, and the cost of search. Since I assume that the bid distribution can be parameterized by its mean and variance, as in (2), and that, apart from the retail price, the search cost and the discount rate do not vary systematically across types of cars, the unconditional probability of sale must depend on mean, variance, and retail price only.[12]

For an estimate of the average selling price on the lot, I use the "average retail price" of the relevant make, model, and model year, as reported in the June 1951 National Automobile Dealer's Association (NADA) Used Car Guide, Region "A" [1951].[13] Region "A" is the Northeast, where the reported auctions were located. It is difficult to know how well NADA's sampling scheme approximates the geographical distribution of the sellers in our data. Also, NADA samples only new car dealers, not used car dealers. However, this last fact may work to our advantage, as it is reasonable to suppose that most of the sellers at the auction were new car dealers selling to used car dealers. It might not be too much of an exaggeration to imagine that the distribution of the winning bid is generated by the pool of used car dealers, and the NADA retail price by the pool of new car dealers.

I estimate β by applying ordinary least squares to (9). I then regress the logarithm of the squared OLS residuals from that regression on the car attributes, z. Let α_0 denote the constant term in α. The resulting coefficient estimates $\hat{\alpha}$ are consistent (though not efficient) estimates of all but α_0 [Harvey 1976]. As an estimate of α_0 I use the slope coefficient from the regression of the squared

12. Using the log-variance ensures that attributes enter linearly into (11). A further advantage to using the log-variance will become evident from equation (16) in the next section.

13. Retail prices reported in the more recent NADA guides are not direct estimates but are imputed from wholesale auction prices. Those reported in the early 1950s were based directly on retail sales.

residuals from (9) on exp $(z_i\hat{\alpha} - \hat{\alpha}_0)$. Finally, I substitute for the mean and log-variance in (11) their predicted values from the previous two steps, and by a probit regression, obtain estimates of γ. Because the system (9)–(11) is exactly identified and lower block triangular, these estimates are consistent though not efficient. Standard errors are calculated by the General Method of Moments (GMM), as in Hansen [1982], and are discussed in Appendix B.

Column (1) of Table IV presents estimates for the mean equation (9). The R^2 of .89 indicates that most of the variance in the bid is captured by the included attributes. The signs of the coefficients are what one would have expected. Price is declining in model year, the ranking of the models seem right (not shown), and cars with high mileage or in poor physical condition fetch a lower price. ("Sharp" indicates the best physical condition, and "rough," the worst.) The coefficients on the auction dummies increase with the auction's distance from Detroit, with particularly high figures around the Philadelphia-New Jersey area. Whitewall tires add an extra \$34 to the bid, and a radio adds another \$29. The one abnormality is the negative coefficient on HEATER, though it is not significant.

Column (2) of Table IV presents estimates for the log-variance equation (10). On the one hand, the newer the car and the better the model (not shown), the higher is the predicted variance. On the other hand, the better the physical condition, the lower the predicted variance. Table V presents summary statistics for the winning bid, the constructed mean, the constructed log-variance, and the average retail price.

The predicted values from the mean and variance regressions are inputs to Table VI, whose first column presents estimates of (11).[14] A greater mean wholesale bid increases, while a greater retail price decreases, the probability of sale at the auction. (Both variables are measured in hundreds of dollars.) An increase in the log-variance is associated with a lower probability of sale. All of this is consistent with search theory, and is significant at the 1 percent level. A simultaneous increase in both the mean wholesale bid and the average retail price decreases the probability of sale, in contradiction to the theory. Yet with a t-statistic of 1.25, this is insignificantly different from zero (which is the model's prediction under a zero interest rate) at the 10 percent level for a one-sided

14. Linear probability and logit models produced near identical estimates of the coefficients (up to a scale factor).

TABLE IV
MEAN AND VARIANCE EQUATIONS

		Bid (1)	ln (u^2) (2)
MODEL YEARS	1951	728	1.5
		(13)	(0.2)
	1950	216	0.5
		(6)	(0.1)
	1948	−271	0.4
		(8)	(0.2)
	1947	−376	0.4
		(9)	(0.2)
	1946	−451	0.4
		(9)	(0.2)
MILEAGE	ln (miles)	−13	−0.1
		(6)	(0.1)
	missing mileage	−65	0.2
		(23)	(0.4)
BODYSTYLE	two-door	2	−0.1
		(6)	(0.1)
	convertible	87	0.2
		(7)	(0.1)
	closed coupe	11	0.3
		(9)	(0.2)
	business coupe	−106	−0.2
		(26)	(0.5)
	sedanette	−17	−0.3
		(12)	(0.2)
	station-wagon	−16	1.0
		(17)	(0.3)
	aero	28	−0.2
		(26)	(0.5)
	belair	250	−0.5
		(28)	(0.5)
	riviera	302	−1.8
		(27)	(0.5)
	cat	392	−1.1
		(38)	(0.7)
OPTIONS	odometer	20	−0.6
		(7)	(0.2)
	whitewalls	34	0.1
		(7)	(0.1)
	hydromatic	41	0.0
		(16)	(0.3)
	dynomatic	99	0.5
		(14)	(0.3)
	radio	29	−0.0
		(5)	(0.1)

TABLE IV
(CONTINUED)

		Bid (1)	ln (u^2) (2)
	heater	−9	−0.0
		(10)	(0.2)
	gyromatic	−4	0.9
		(31)	(0.6)
	power glide	74	0.5
		(21)	(0.4)
	"wf"	−40	−1.5
		(33)	(0.6)
	visor	16	−0.3
		(31)	(0.6)
CONDITION I	sharp	275	−1.3
		(27)	(0.5)
	clean	197	−1.3
		(26)	(0.5)
	good	178	−1.2
		(27)	(0.5)
	fair	104	−1.3
		(27)	(0.5)
	rough	8	−0.1
		(36)	(0.7)
CONDITION II	sharp/sharp	187	−1.0
		(21)	(0.4)
	sharp/clean	173	−0.9
		(22)	(0.4)
	sharp/fair	35	1.5
		(40)	(0.7)
	clean/sharp	167	−1.0
		(25)	(0.5)
	clean/clean	151	−0.9
		(20)	(0.4)
	clean/fair	87	−0.4
		(22)	(0.4)
	fair/clean	86	−0.2
		(29)	(0.5)
	fair/fair	120	−0.5
		(23)	(0.4)
TIRES	excellent	36	−0.2
		(19)	(0.4)
	good	1	−0.3
		(17)	(0.3)
	fair	9	−0.1
		(18)	(0.3)
AUCTIONS	Akron	30	0.8
		(18)	(0.3)

TABLE IV
(CONTINUED)

	Bid (1)	ln (u^2) (2)
Aptco	23	−0.2
	(14)	(0.3)
Arena	−7	0.1
	(14)	(0.3)
Belair	55	0.4
	(16)	(0.3)
Danville	10	0.6
	(16)	(0.3)
Ebensburg	24	0.6
	(16)	(0.3)
Emlenton	42	0.4
	(15)	(0.4)
Fort Wayne	72	0.1
	(13)	(0.2)
Gilbert	57	0.6
	(17)	(0.3)
Manheim	96	0.4
	(17)	(0.3)
Owosso	−10	0.3
	(14)	(0.3)
Plainfield	66	1.0
	(17)	(0.3)
Toledo	8	0.6
	(21)	(0.4)
R^2	.89	.10
Root MSE	117	2.2

The number of observations is 3209. The independent variables also include 50 model-specific dummies. The omitted regressors are 1949 model year, four-door body style, rough/rough condition, poor tires, and Simpson (Detroit) Auction. Standard errors are in parentheses.

test. However, one must be somewhat leery of this last test. It requires that NADA's average retail price be an especially good proxy for its theoretical counterpart: not only must the two be ranked the same, but the one must increase dollar for dollar with the other.

To account for the possibility that the log-variance is mimicking nonlinear terms in the mean, with which it is positively correlated, column (2) adds the square of the predicted mean. The estimated coefficient on the variance is barely affected and remains negative. When the square of the log-variance is added, the variance terms are jointly significant and negative throughout the range of the data, and the remaining terms are uneffected.

TABLE V
SUMMARY STATISTICS

	Mean	Standard deviation	Minimum	Maximum
Bid	1090	352	110	2925
Mean bid	1090	333	340	2512
Log-variance	7.74	0.69	5.44	11.09
Retail price	1406	402	138	2925

Mean bid is the predicted value from the first column of Table IV. Log-variance is the predicted value from the second column of Table IV.

TABLE VI
PROBIT ESTIMATES OF THE PROBABILITY OF SALE

		Unconditional		Conditional on bid				
		(1)	(2)	(3)	(4)	(5)	(6)	(7)
Constant		2.043	2.446	2.033	1.009	1.013	2.045	1.009
		(0.273)	(0.373)	(0.271)	(0.089)	(0.089)	(0.301)	(0.088)
Bid	y			0.150	0.209	0.214	0.147	0.209
				(0.021)	(0.032)	(0.035)	(0.039)	(0.032)
Mean bid	μ	0.072	0.011	−0.064		−0.124	−0.061	
		(0.018)	(0.042)	(0.026)		(0.040)	(0.043)	
Log-variance	$\log \sigma^2$	−0.160	−0.174	−0.153			−0.155	
		(0.036)	(0.038)	(0.037)			(0.043)	
Retail price	μ_L	−0.081	−0.078	−0.095	−0.103	−0.104	−0.096	−0.103
		(0.015)	(0.015)	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)
END $\sigma \int \max(y,x) \times dH(x)$					−0.121			−0.120
					(0.039)			(0.039)
END − mean						−0.137	0.007	
						(0.064)	(0.077)	
Mean bid**2			0.0025					
			(0.0016)					
t-statistic		−1.25[a]		−1.17[b]	−2.24[c]	−1.98[d]	−1.17[d]	−2.25[c]
H_0 assumption		NONE	NONE	NONE	Normal	Normal	Normal	EmpCDF

a. Test that sum of coefficients on Mean bid and Retail price exceeds zero.
b. Test that sum of coefficients on Bid, Mean bid, and Retail price exceeds zero.
c. Test that sum of coefficients on Bid, END, and Retail price exceeds zero.
d. Test that sum of coefficients on Bid, END, Mean bid, and Retail Price exceeds zero.
Standard errors are in parentheses.

V. THE CONDITIONAL SALE RATE AND THE RESERVATION PRICE

I assume that in deciding whether or not to accept a winning bid, each seller adopts a reservation price strategy—accepting all bids exceeding that price, rejecting all below. This reservation price need not be the same as that predicted by search theory; it may reflect only the seller's opportunities elsewhere than at the auction, be strictly proportional to the mean bid, perhaps, or be what the seller deems a "fair value" for the car. What the empirical analysis should reveal is whether or not the search theoretic reservation price and the observed reservation price coincide.

Although the reservation price is not directly observed, its functional dependence on other variables may be inferred by the qualitative choice method.[15] A sale is observed if and only if the bid exceeds the reservation price

$$y > R. \tag{12}$$

Taking a linear approximation to the reservation price and adding on an error term, the condition reads

$$y + \{\delta_0 + \delta_2\mu(z) + \delta_3 \log \sigma^2(z) + \delta_4\mu_L\}/\delta_1 - v > 0. \tag{13}$$

The random variable v is assumed to be distributed normally, with zero mean and variance $(1/\delta_1)^2$, and reflects deviations from the median reservation price that are specific to the seller and uncorrelated with the attributes of the car. These may reflect differential search costs, such as the transportation cost between the seller's lot and the auction, deviations from the average retail prices due to differences in local market conditions, or even private information that sellers hold. Thus,

$$\begin{aligned} E[I\,|y,z] &= \text{prob}\,\{y > R\,|y,z\} \\ &= \Phi(\delta_0 + \delta_1 y + \delta_2(z\beta) + \delta_3(z\alpha) + \delta_4\mu_L). \end{aligned} \tag{14}$$

Because y is observed, the usual inability to identify the scale of the coefficients in a binary choice model does not arise. In fact, the model is identified by the restriction that the coefficient on the bid y, once normalized by scale, equal one.

The full model is (9), (10), and (14). Equation (14) identifies the reservation price through the probability of sale conditional on the bid. It differs from (11) only in the addition of this variable; alternatively, (11) is derived from (14), with the bid "integrated out."

I obtain consistent estimates of $\boldsymbol{\beta}$ and $\boldsymbol{\alpha}$ as in the previous section, and by a probit corresponding to (14), consistent estimates

15. In principle, the distribution of reservation prices evaluated at any price may be consistently and nonparametrically estimated by the sale rate at that bid. If one-third of the time a $2000 bid for some type of car is accepted, it must be that the reservation price of one-third of the sellers is less than $2000. This is the natural extension of using the minimum accepted offer when sellers are homogeneous [Flinn and Heckman 1982]. To be practicable, such an approach would require, for each type of car, many offers at the same price. This is not available here.

of δ. The predictions of the propositions may be restated as δ_2, δ_3, $\delta_4 \leq 0$ (increases in the mean and log-variance of the bid and the retail price all increase the reservation price), and $\delta_1 + \delta_2 + \delta_4 \geq 0$ (a simultaneous increase in the mean bid and retail price increases the reservation price by less than one).

Column (3) displays estimates and GMM standard errors for (14). The estimate of 0.15 for the coefficient on the winning bid (measured in hundreds of dollars) suggests a standard deviation of the reservation price of \$667. This is almost six times the root mean square error from equation (9), and more than half the average winning bid. The difference in the standard errors of the reservation price and winning bid distributions is not unexpected: the winning bid is an extreme order statistic, and thus should exhibit a smaller variance than the variance of its underlying distribution, if the latter is bounded. The result follows if the dispersion of a seller's valuation is similar to any given bidder's. It is surprising that the dispersion in the reservation price is so large relative to the mean winning bid; it suggests that for many sellers, the sale option on the lot must be extremely small, perhaps because they are already overstocked in that type of car.

The coefficients on the mean bid and retail price are negative, as predicted: a one-dollar increase in the mean bid increases the reservation price by (0.064/0.15 =) 42 cents, and a one-dollar increase in the retail price increases the reservation price by 64 cents. This sums to more than one, in contradiction to the prediction that a simultaneous dollar increase in mean bid and retail price raises the reservation price by less than one dollar, although the t-statistic is insignificant.

The sign of the variance coefficient accords with search theory: it is negative and significant. However, the magnitude of the variance coefficient is much greater than the theory would predict. To see this, first rewrite (7) as

$$(1 + \rho)R = H_0(r)R + \int_r [\sigma x + \mu]\, dH_0(x) - c, \tag{15}$$

where $r \equiv (R - \mu)/\sigma$, and then note that

$$\begin{aligned} R'\,(\log \sigma^2) &= \sigma \int_r x\, dH_0(x)/2\{1 + \rho - H_0(r)\} \\ &< \sigma \int_r x\, dH_0(x)/2\{1 - H_0(r)\} \\ &= \{E_H[y \mid y \geq R,z] - E_H[y \mid z]\}/2 \\ &= \{E_H[y \mid I = 1,z] - E_H[y \mid z]\}/2. \end{aligned} \tag{16}$$

The left-hand side of the inequality can be estimated from column (3) by the ratio of the estimated coefficient on the log-variance to (1/100 times) the coefficient on the bid, which is 100. The right-hand side is half the difference between the mean accepted bid and the mean bid, conditional on attributes z. For the set of cars that sold, the average value of the right-hand side can be estimated by half the average difference between the bid and the predicted mean bid, which is a mere five dollars, and has a standard error of one dollar. Estimating an average value over the entire sample is less straightforward. It will not do to take the difference between the mean accepted bid and the mean bid for the entire sample, as sale is not independent of attributes. Instead, the bid was regressed on the car attributes on the subsample of cars that sold, and a predicted value formed from the estimated coefficients for all observations. Half the average difference between this value and the observed bid is five and one-half dollars. Thus, the estimated effect of the variance on the reservation price is nearly twenty times as large as search theory would predict.

Learning is one possible explanation for this discrepancy. Assuming that the offer distribution is normal with an unknown mean and known variance, and that sellers hold normal priors over the mean, Burdett and Vishwana [1988] show that the initial reservation price exceeds the no-learning reservation price by an amount that is itself increasing in the variance of the prior. A coefficient that exceeds the bound in (16) might be consistent with a model of search with learning, if, in addition, the variance of the prior is positively associated with the variance of the true distribution, which is what I estimate. Most learning models (e.g., Rosenfield and Shapiro [1981] and Burdett and Vishwana [1988]) have the further implication that the reservation price declines with the number of previous searches. Unfortunately, the lack of any recorded unique identifier for the cars at the auction makes following them over time impossible.[16]

In contrast, nonstationarity is unlikely to explain the discrepancy. Imagine that all used cars depreciate by A dollars each period in both the mean and retail price. With a zero interest rate, the reservation price would also fall by A dollars each period, thus

16. It bears repeating that the seller's initial knowledge of the offer distribution is much finer than in more common applications of the model, such as worker or consumer search. He will have bid, or have watched others bid, on similar cars; he may have read the NADA Used Car Guide, which provides average wholesale prices, or AMR, which at the time provided both the entire empirical distribution of wholesale prices, and its own hedonic index.

leaving the sale probability unchanged over time. The level of the reservation price would correspond to a stationary model with search cost $c + A$, and so Propositions 1 and 2 would continue to hold. Time would be an omitted variable if the interest rate were positive, but as it would also be small, so would any bias. In any case, there is little evidence of nonstationarity. A time trend added to (9) yields a small, positive, and insignificant coefficient.

Finally, consider the consequences of omitted characteristics. Undoubtedly, the recorded attributes do not encompass all that dealers find relevant about the consigned cars. Some component of the measured variance reflects heterogeneity not among the winning bids for a perfectly defined car, but among the cars themselves. Since there is no reason to suppose that this part of the variance has any effect on the reservation price, its presence will, in general, act like measurement error, and bias the variance coefficient in the unconditional probability estimates of columns (1) and (2) toward zero. For the special case in which the "attribute" variance is proportional to the "bid" variance, the constructed log variance will differ from the desired estimate by a constant only, and the unconditional probability results will remain valid.

Matters are more complicated for the conditional probability of sale estimates of column (3), for where there are unrecorded attributes the error in the reservation price must then be positively correlated with the bid: cars with higher bid residuals are higher quality cars, and so are more highly valued by sellers as well. This should lead to an underestimate of the coefficient on the bid, relative to the other coefficients, and may thus explain why both the log-variance and a simultaneous increase in the mean and retail price are estimated to increase the reservation price by more than the theory would predict. Overbidding by buyers who recognize sellers with high reservation prices should have a similar effect.

The solution to this problem is to instrument the bid with variables that affect the winning bid but not the reservation price. Since the latter reflects future opportunities only, any variable that measures "temporary shocks" is an admissible instrument. Unfortunately, neither date dummies nor dates interacted with auction sites, nor selected functions of the same (weather conditions and the total number of cars consigned at the auction, and deviations of the time from each auction's average) have any explanatory power in predicting the winning bid.

VI. The Exact Linear Form

An alternative approach to the method in Section V is to work with (12) directly, under the assumption that H_0 is known. For any bid y, we can define the search cost for which the bid equals the reservation price. From (7), that value is $-(1+\rho)y + \int \max(y,x)\, dh(x)$. A sale occurs if the true search cost exceeds that value. Say that the cost of search c may be written as $c_m - \zeta$, where ζ is distributed normally and independently of z, with standard deviation σ_ζ: ζ is specific to the seller but not to the car. As before, ζ may be interpreted as a deviation in either the primitive search cost or the expected retail price. Then the probability of sale, conditional on the bid y, is

$$\Phi(\{c_m + (1+\rho)y - \int \max(y,x)\, dH(x) - \phi\mu_L\}/\sigma_\zeta). \tag{17}$$

Define END $\equiv \int \max(y,x)\, dH(x)$ (for "expected next draw").

Estimating (17) requires an assumption on H_0. In column (4), which presents estimates from a probit regression on the bid, END, and the average retail price, H_0 is assumed to be standard normal. As expected, the coefficients on END and the retail price are negative. Also, the coefficient on END is smaller in magnitude than that on the bid, as expected. However, the sum of the coefficients on all three variables, which, by inspection of (17), should equal $\{(1+\rho) - 1 - \phi\}/\sigma_\zeta$ and so be nonnegative, is, in fact, significantly negative. Again, this result depends on NADA's average retail price increasing one for one with the theoretical construct, and so is somewhat suspect.

In Column (5), END is decomposed into two parts, μ and END $- \mu$. This is justified by the identity,

$$\int \max(y,x)\, dH(x) \equiv \mu + \sigma \int \max\{[y-\mu]/\sigma, x\}\, dH_0(x). \tag{18}$$

Although there are any number of different ways to decompose END, this one is particularly attractive as it permits the mean bid to enter the regression independently of the variance. In fact, the hypothesis that the coefficients on the two terms are equal cannot be rejected at any reasonable level of significance (t-statistic of 0.3).

Column (6) adds the log-variance term back into the regression. This is a more meaningful test of the structural version of the model: search theory predicts that the variance should enter the conditional probability of sale in a precise way, as captured by END, and thus predicts a zero coefficient on the log-variance. Clearly, the data would have it otherwise; including the log-variance term reduces the coefficient on END to near zero, and nearly replicates the "nonstructural" estimate of column (3).

To check the robustness of these results to the choice of the bid distribution H_0, column (7) presents estimates under the assumption that H_0 is the actual empirical distribution of $\{[y_i - z_i\beta]/\exp(z_i\alpha)\}$. (The reported errors *do not* take the estimation of H_0 into account.) The results are remarkably similar to those in column (5), especially in light of the fact that standard tests dramatically reject the hypothesis of normality for the empirical c.d.f. Thus, the rejection of the exact linear form is not an artifact of the choice of the bid distribution.

VII. Conclusion

Search theory provides an organizing structure to explain the behavior of market participants when meeting in order to trade is costly and offered terms of trade are ex ante unknown. The theory yields both general qualitative predictions about the probability of sale, both conditional on the offer and unconditional, as well as precise restrictions on the functional form of the conditional probability of sale.

In this paper I have applied the simplest search model to sellers' behavior at wholesale used car auctions. A fair assessment of the theory must be that it fails in the details but succeeds in the broader patterns. Variance, mean bid, and retail price increase the reservation price. The mean bid increases the unconditional probability of sale, while the variance and retail price lowers it. All this is as the theory predicts.[17]

Some of these partial correlations might have been predicted without the aid of search theory. The simple mathematics of probability theory predicts a negative association between the variance of received offers and the unconditional sale rate, when the reservation price is unresponsive to the former. The principle of opportunity cost alone suggests that the retail price be positively associated with the reservation price.

17. The few laboratory search experiments that considered mean-preserving spreads obtained similar results: increasing dispersion increases self-reported reservation prices [Schotter and Braunstein 1981] and search durations [Cox and Oaxaca 1989]. The empirical literature on search by job seekers has tended not to consider distributional implications. One paper [Warner, Poindexter, and Fearn 1980] that considers the effect of the variance of actual (and therefore accepted) wages on a self-reported reservation wage did not find any significant effects.

One would *not* have predicted, absent search theory, that the seller's reservation price be an increasing function of the variance of the bid. Nor that the retail price be negatively associated, and the mean bid positively associated, with the unconditional sale rate.

However, some of the quantitative predictions of the simple search model are not borne out. The predicted functional form is clearly rejected by the data, as shown in Section VI. Also, the reservation price is more sensitive to the variance than the theory would predict. This last result might be explained by a more general search model in which sellers learn about an unknown offer distribution during the search process. The presence of unrecorded attributes might also be responsible for the results.

APPENDIX A: THE 1951 DATA

Automotive Market Report (AMR) was first published in March 1951 with the purpose of reporting on wholesale automotive markets. In its initial issues, AMR recorded the winning bids of each dealer-consigned car at every auction that its reporters attended. As AMR's coverage of auctions expanded, it restricted its listing to accepted bids only, and finally only a selection of these. I use data from the initial issues only.[18]

For each car consigned, the winning bid (y), whether or not the car sold (I), and characteristics of the car (z) are indicated. Unfortunately, the list of characteristics is not always the same. The model year, the make, the model, and the body style are always listed. The mileage is almost always given as well. There is almost always one of two types of descriptions of the condition of the car. The first is a single index which rates the car as a whole as either new, sharp, clean, good, fair, poor, or rough. The second type consists of three indices, one each for the interior, exterior, and tires. The date of the auction and the auction location are also known. Unfortunately, it is impossible to track specific cars across dates.

The editions of May 14, May 28, June 11, June 25, and July 9 yielded 5777 readable observations. The criteria for admittance to the final data set were (i) a make-model type with at least eighteen observations, (ii) a model-year of 1946 or later, (iii) a listing of the options, (iv) one of the two condition ratings (v) a major body style,

18. Fleet cars (of rental car companies and other large companies) are also sold at wholesale auto auctions but are not reported in Automotive Market Report.

and (vi) a make-model type listed in the June 1951 N.A.D.A. Used Car Guide. All Cadillacs were excluded. The natural cutoff year is 1946 because, due to the War, and there are no 1945, 1944, or 1943 model years in the data, although there are 1942 model years and earlier. Since some body types are concentrated among a few model types, conditions (i) and (v) are not independent; an informal, iterative procedure resulted in the choice of eighteen observations as the cutoff for (i). Application of these conditions reduced the number of observations to 3209.

APPENDIX B: STANDARD ERRORS

Standard errors are calculated according to the general method of moments. For the conditional probability model, the estimates of the coefficients (α,β,δ) may be interpreted as solutions to the moment equations:

$$\sum_{i=1,\dots,N} z_i(y_i - z_i\beta) = 0 \tag{B1}$$

$$\sum_{i=1,\dots,N} z_i[\ln (y_i - z_i\beta)^2 - z_i\alpha] = 0 \tag{B2}$$

$$\sum_{i=1,\dots,N} z_i w_i = 0, \tag{B3}$$

where $w_i \equiv [I_i - \Phi_i]/(\Phi_i\,(1 - \Phi_i))$, $\Phi_i \equiv \Phi(t_i\delta)$, and t is the vector of regressors in the probit; all other variables are defined as in the text. Note that (B3) implicitly defines the Probit estimator. Let W_{12} be the White Asymptotic Variance-Covariance matrix for the system of equations (B1)–(B2), L the log likelihood of the probit, $V_3 = \partial^2 L/\partial\delta\partial\delta$, and $B = \partial^2 L/\partial\delta\partial(\beta,\alpha)$. Then the variance-covariance matrix for the estimator of δ is

$$V = V_3 + V_3 B W_{12} B' V_3. \tag{B4}$$

This assumes that the bid residuals are distributed independently of the reservation price residuals—in other words, that winning bidders and sellers are randomly matched. In practice, the estimate of V is little different from the simple probit standard errors, the estimate of V_3.

For the unconditional probability equations, we must take the covariance between the residuals of (9) and (10) and the residuals of (11) into account—large bids are more likely to be accepted. In this case, the variance covariance matrix of γ is

$$V = V_3 + V_3 B W_{12} B' V_3 - V_3 MAB' V_3 - (V_3 MAB' V_3)', \tag{B5}$$

where $A = I_2 \otimes (Z'Z)^{-1}$, $M = (w_i\phi(t_i\delta)t_i'z_ir_{1i} \,|\, w_i\phi(t_i\delta)t_i'z_ir_{2i})$, Z is the matrix of attributes, r_1 is the residual from the bid equation, r_2 is the residual from the variance equation, and ϕ is the normal probability density.

DEPARTMENT OF ECONOMICS, MASSACHUSETTS INSTITUTE OF TECHNOLOGY

REFERENCES

Ashenfelter, Orley, "How Auctions Work for Wine and Art," *Journal of Economic Perspectives,* III (Summer 1989), 23–36.
Balvers, Ronald J., "Variability and the Duration of Search," *International Economic Review,* III (August 1990), 747–51.
Burdett, Kenneth, and Tara Vishwana, "Declining Reservation Wages under Learning," *Review of Economic Studies,* LV (1988), 655–65.
Cox, J. C., and R. Oaxaca, "Laboratory Experiments with a Finite Horizon Job Search Model," *Journal of Risk and Uncertainty,* II (1989), 301–30.
Flinn, C., and J. Heckman, "New Methods for Analyzing Structural Models of Labor Force Dynamics," *Journal of Econometrics,* XVIII (1982), 115–68.
Frappier, Dennis (General Manager, Portland Fish Exchange), personal conversation, May 1992.
Genesove, David, "Bargaining," mimeo, January 1992.
——, "Adverse Selection in the Wholesale Used Car Market," *Journal of Political Economy,* CI (August 1993), 644–65.
Hansen, Lars Peter, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica,* L (July 1982), 1029–54.
Harvey, A. C., "Estimating Regression Models with Multiplicative Heteroscedasticity," *Econometrica,* XLIV (May 1976).
Hart, Bernard, personal correspondence, March 1989.
Kohn, Meir, and Steven Shavell, "The Theory of Search," *Journal of Economic Theory,* IX (1974), 93–123.
Lawrence, Eric G., "The Wholesale Automobile Auction," Unpublished thesis, Johns Hopkins University, 1984.
Mortensen, Dale T., "Job Search and Labor Market Analysis," *The Handbook of Labor Economics,* O. Ashenfelter and R. Layard, eds. (Amsterdam, The Netherlands: Elsevier Science Publishers BV, 1986), pp. 849–919.
National Auto Auction Association, *National Auto Auction Association Membership Directory,* 1989.
National Automobile Dealers Association, *NADA 1988 Economic Survey* (1989).
——, *N.A.D.A. Used Car Guide* (June 1951).
Reiss, R.-D., *Approximate Distributions of Order Statistics: With Applications to Nonparametric Statistics* (Amsterdam: Springer-Verlag, 1989).
Rosenfield, Donald B., and Roy D. Shapiro, "Optimal Adaptive Price Search," *Journal of Economic Theory,* XXV (1981), 1–20.
Schotter, A., and Y. M. Braunstein, "Labor Market Search: An Experimental Study," *Economic Inquiry,* XIX (January 1981), 133–44.
Warner, J. T., J. Poindexter, and R. Fearn, "Employer-Employee Interaction and the Duration of Unemployment," *Quarterly Journal of Economics,* XCIV (1980), 211–33.

LINKED CITATIONS

You have printed the following article:

Search at Wholesale Auto Auctions
David Genesove
The Quarterly Journal of Economics, Vol. 110, No. 1. (Feb., 1995), pp. 23-49.
Stable URL:
http://links.jstor.org/sici?sici=0033-5533%28199502%29110%3A1%3C23%3ASAWAA%3E2.0.CO%3B2-Q

[Footnotes]

[17] **Employer-Employee Interaction and the Duration of Unemployment**
John T. Warner; J. Carl Poindexter, Jr.; Robert M. Fearn
The Quarterly Journal of Economics, Vol. 94, No. 2. (Mar., 1980), pp. 211-233.
Stable URL:
http://links.jstor.org/sici?sici=0033-5533%28198003%2994%3A2%3C211%3AEIATDO%3E2.0.CO%3B2-L

References

How Auctions Work for Wine and Art
Orley Ashenfelter
The Journal of Economic Perspectives, Vol. 3, No. 3. (Summer, 1989), pp. 23-36.
Stable URL:
http://links.jstor.org/sici?sici=0895-3309%28198922%293%3A3%3C23%3AHAWFWA%3E2.0.CO%3B2-5

Variability and the Duration of Search
Ronald J. Balvers
International Economic Review, Vol. 31, No. 3. (Aug., 1990), pp. 747-751.
Stable URL:
http://links.jstor.org/sici?sici=0020-6598%28199008%2931%3A3%3C747%3AVATDOS%3E2.0.CO%3B2-2

LINKED CITATIONS

Declining Reservation Wages and Learning
Kenneth Burdett; Tara Vishwanath
The Review of Economic Studies, Vol. 55, No. 4. (Oct., 1988), pp. 655-665.
Stable URL:
http://links.jstor.org/sici?sici=0034-6527%28198810%2955%3A4%3C655%3ADRWAL%3E2.0.CO%3B2-N

Large Sample Properties of Generalized Method of Moments Estimators
Lars Peter Hansen
Econometrica, Vol. 50, No. 4. (Jul., 1982), pp. 1029-1054.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28198207%2950%3A4%3C1029%3ALSPOGM%3E2.0.CO%3B2-O

Estimating Regression Models with Multiplicative Heteroscedasticity
A. C. Harvey
Econometrica, Vol. 44, No. 3. (May, 1976), pp. 461-465.
Stable URL:
http://links.jstor.org/sici?sici=0012-9682%28197605%2944%3A3%3C461%3AERMWMH%3E2.0.CO%3B2-B

Employer-Employee Interaction and the Duration of Unemployment
John T. Warner; J. Carl Poindexter, Jr.; Robert M. Fearn
The Quarterly Journal of Economics, Vol. 94, No. 2. (Mar., 1980), pp. 211-233.
Stable URL:
http://links.jstor.org/sici?sici=0033-5533%28198003%2994%3A2%3C211%3AEIATDO%3E2.0.CO%3B2-L