

The Effects of Business-to-Business E-Commerce on Transaction Costs

by

Luis Garicano and Steven N. Kaplan*

First Draft: July 2000

This Draft: July 2001

Abstract

This paper studies transaction costs changes arising from the introduction of the Internet in transactions between firms. We divide transaction costs into coordination costs and motivation costs. We classify coordination efficiencies into three categories: process improvements, marketplace benefits, and indirect improvements. For motivation costs, we focus on informational asymmetries. We apply this framework to internal data from an Internet-based firm to measure process improvements, marketplace benefits, and motivation costs. Our results suggest potentially large process improvements and marketplace benefits. We find little evidence that informational asymmetries are more important in the electronic marketplace than in the existing physical ones.

* Graduate School of Business, University of Chicago. This research has been supported by the Center for Research in Security Prices, the Kauffman Foundation, and the Lynde and Harry Bradley Foundation and the Olin Foundation through grants to the Center for the Study of the Economy and the State. We thank Severin Borenstein, Judy Chevalier, David Genesove (the editor), Charles Morris, Rod Parsley, Erik Peterson, Jagadish Turimella, Frank Wolak, an anonymous referee, and seminar participants at IESE (Barcelona), University Pompeu Fabra, and the NBER E-Commerce project for helpful comments and discussions. Address correspondence to Luis Garicano, Graduate School of Business, The University of Chicago, 1101 East 58th Street, Chicago, IL 60637 or e-mail at luis.garicano@gsb.uchicago.edu. ©2000 Luis Garicano and Steven Kaplan.

1. *Introduction.*

By separating the physical and information flows connected with each transaction, the Internet potentially provides new ways to trade goods and services. For each individual transaction, the choice between an Internet-based transaction or marketplace and a physical one depends on the transaction costs of each of the two alternatives.

In this paper, we study the changes in transaction costs that result from B2B electronic commerce. The Internet has the potential to affect both coordination and motivation costs. Coordination costs are those “related to the need to determine prices and other details of the transaction, to make the existence and location of potential buyers and sellers known to one another, and to bring the buyers and sellers together to transact.¹” The Internet may change these coordination costs directly in two ways. First, it may reduce the costs involved in an existing business process, either by reducing the cost of an activity already being conducted, as when a transaction that is currently conducted by phone or fax is automated; or by providing an opportunity to redesign the existing process. Second, it may reduce the cost of matching buyers and sellers as it is potentially less expensive to search for products, compare prices, or find buyers over the Internet than to read catalogs and make phone calls. As a result, buyers will find sellers they might not have otherwise found.

Second, the Internet has the potential to affect motivation or incentive costs, including those related to informational asymmetries and those related to imperfect commitment. First, to the extent that physically observing the merchandise to evaluate its condition is valuable to the buyer, informational asymmetries may be worse in Internet transactions. Second, the Internet may change the cost of commitment. On the one hand, by standardizing processes and by leaving an electronic trail, the Internet has the potential to reduce commitment costs. On the other hand, buyers may avoid intermediary fees by viewing the product over the Internet, but contacting the seller directly.

In estimating the impact of the Internet on these costs, we rely on detailed internal data from one Internet-based firm, Autodaq, to measure process improvements, marketplace benefits, and motivation costs. Autodaq operates in the wholesale used car auction market. The sellers in this market are

automobile rental companies, lease finance companies, commercial fleets, manufacturers, and automobile dealers. Autodaq initially focused on large volume sellers. The buyers in this market are automobile dealers or wholesale buyers, not individuals. For example, rental companies buy new cars, rent them for a specified time or usage, and then sell them. Lease finance companies lease new cars for several year terms. At the end of the lease term, the lease finance company must sell those cars (roughly 50%) that are not purchased by the lessee or the original dealer. In 1999, the wholesale used car auction market included roughly 1.4 million cars from rental companies, 1.9 million cars from lease finance companies, 2.1 million cars from commercial fleets, and 3.5 million cars from dealers and wholesalers.² The wholesale used car auction market, therefore, is roughly 9 million cars per year.

Our results suggest that process improvements and marketplace benefits are potentially large – on the order of 5% of the automobile value and 80% of the economic transaction cost. Moreover, we do not find evidence that the Internet increases adverse selection costs. We suspect that this is at least due in part to the measures taken by Autodaq to avoid this problem.

The paper proceeds as follows. Section II estimates the impact on process improvements. Section III estimates marketplace benefits. Section IV estimates the impact on asymmetric information costs. Section IV concludes.

II. Impact on Coordination costs (1): Process improvements

In this section, we compare the time and economic costs involved in the Autodaq / Internet process with those in the physical auction process.

In the physical world, when a large volume seller needs to dispose of a car, the seller stores the car and then has it transported to a physical auction site. At the physical auction site, the car is described and inspected. The car may also be reconditioned by the auction site operator. Reconditioning involves repairing minor flaws in the car's exterior – dents, scratches, etc. When a sufficient number of cars are physically at the auction site, an auction is held. Dealers travel to the physical auction site and bid on the

car. After the auction, the car is transported again to the winning dealer. The winning dealer performs any necessary maintenance or repairs and any additional reconditioning needed to retail the car.

In the Autodaq / Internet system, Autodaq contracts with an inspector who inspects, describes, and photographs the car. For cars coming off lease, this occurs at the turn-in dealer. For cars coming from rental fleets, this occurs at the fleet marshalling yard. The car is then put up for sale in an online auction. Dealers bid on the car over the Internet from their computers. The car is transported to the winning dealer. The winning dealer performs any necessary maintenance, repairs and reconditioning. If the car does not sell over the Internet, the car continues through the physical auction process.³

Unlike physical auctions, which are run as ascending oral auctions, Autodaq auctions employ a second price auction in the form of a “proxy bidding” mechanism. With a proxy bid, dealers submit the highest price they would be willing to pay and Autodaq automatically increases their bid in the presence of other bids by just enough to become the leading bid. The auction format also allows dealers to directly purchase the car by accepting the ask price given by the seller.

[Table I Here]

Table I compares the physical auction process to the Autodaq Internet process, both in terms of time and money. The comparison is made for a typical car coming off lease or from a rental fleet. The table measures time from the day the car comes off lease or is retired by the rental car company to the day the car arrives at the buying dealer. The table measures costs as the economic costs of the process. It does not measure the benefits to a seller from moving from a physical process to the Internet process. We report both estimated times involved in the physical auction process and in the Internet process and actual times for both processes from a sample provided by one of the sellers that used the Autodaq process.

The estimates for the physical auction process in column (1) and (2) were provided by Autodaq and by Tom Kontos of ADT Automotive. As mentioned earlier, ADT Automotive was the second largest competitor in the physical auction business.⁴ We obtained similar estimates in interviews with other industry participants. Column 1 reports that the physical auction process takes from 28 to 44 days.

We also estimate these costs directly from a sample of cars sold through the physical auction process provided to us by one of the lessors that provided Autodaq cars to sell. Our analysis is in columns (3) and (4) of Table I. The information provided by the seller allows us to calculate time to sale from (1) lessor inspection date and (2) lease end for cars sold through the physical auction and for cars sold through Autodaq.

Neither date is ideal. According to Autodaq, a car was typically inspected before the lessee turned it in. Time to sale from lessor inspection date, therefore, overstates the time from turn-in to sale. A car sold through Autodaq was inspected an average of 9 days before the car was turned in. According to Autodaq, the overstatement is slightly worse for the cars sold through Autodaq because all such cars were inspected before they were turned in. While most of the cars sold through the physical auction were inspected before they were turned in, a small number were inspected at the physical auction. The comparisons between Internet and physical auction processes, therefore, will slightly understate the advantage of the Internet.

Time to sale from lease end also is problematic because cars are sold both well before the lease end date and well after the lease end date. On this dimension, we do not know if there is a bias between the cars sold through Autodaq and cars sold through the physical auction.

In our analysis, we use the time from inspection to sale because (1) it appears to be a more reliable measure of the disposition process and (2) we have inspection dates for all cars, but do not have lease-end dates for all cars. The results are qualitatively similar using both dates.

For all 9,205 cars, we calculate the time that elapsed from lessor inspection date to the date the car was sold. The median time is 35 days. We add two days to this to estimate the delivery time from the auction to the purchasing dealer. As we report in Table I, the median elapsed time is 37 days. This is close to the midpoint of the range provided by ADT Automotive. 37 days also is consistent with the estimates we obtained from other interviews. On the other hand, if these cars were inspected 9 days before they were turned in, the median time would be more like 28 days, which is at the low end of the range in column 1.⁵

Column 5 in the table reports Autodaq's estimates of the time that is involved in the Internet auction process. Autodaq believes that the Internet process should take 7 days compared to the 28 to 44 days in the physical auction process.

The potential time reductions come in several areas. First, it typically takes 9 to 15 days before lessors and fleet owners ship a car to the physical auction site. Part of the reason for the delay is that the physical auction company does not pick the car up immediately. The other reason is that the seller may attempt to sell the car to the original dealer, but must take some time attempting to determine the appropriate price. It is not entirely clear that all of the savings here are Internet specific. It would seem possible for the lessors to contract with a physical auction site to reduce this time as well. It remains to be seen whether Autodaq can reduce this time.

Second, it typically takes 15 to 25 days from the time a car arrives at a physical auction site until it is sold. On the Internet, Autodaq estimates this time can be reduced to 4 days. One reason for the delay in the physical auction process is that the car generally waits some time before it is reconditioned and reconditioning takes some time.⁶ The more important reason, however, is that the physical auction sites try to make each individual auction somewhat homogeneous in terms of the cars available. In other words, they attempt to sell largely Fords in one auction; largely Toyotas in the next. This is done because dealers typically look for particular types of cars. As a result, the physical auctions will wait until they have a critical mass of a particular car type or brand before holding an auction. This is not a consideration for Autodaq because the dealer does not have to physically go to the Autodaq auction site.

Autodaq's estimates make two optimistic assumptions. First, the estimates assume that the cars sell quickly on the Internet, which implies a liquid market. Second, the estimates assume that the cars are listed for sale almost immediately after they come off lease which assumes sophisticated and timely tracking and inspection processes. We interpret Autodaq's estimates, therefore, as the likely process costs of a liquid Internet market.

To obtain a more neutral estimate of the gains generated by the Internet, we calculated the actual time that elapsed from the day a car was inspected by the lessor to the day the car was sold for a sample

of 694 cars sold over the Internet by Autodaq. The median time is 14 days. We add three days to this to estimate the time until the car is delivered to the purchasing dealer. Column 7 reports that the median actual elapsed time is 17 days.⁷

The “Dollars” columns in Table I attempt to value the economic costs of the two processes. The most important costs are the costs of capital, depreciation, and transportation. The cost of capital is relatively straightforward. The typical car (in our sample) sells for \$13,600. Each day the car is not sold, the seller is not able to deploy that capital elsewhere. We assume a cost of capital of 8%. This is essentially a debt cost of capital (and, as such, may understate the true cost of capital for a seller).

The cost of depreciation is based on the fact that sale prices for used cars depreciate with the age of the car. We assume a depreciation rate of 14%. This reflects the fact that in the data provided by Autodaq, the sales price declines by 14.8% per car-year (with a standard error of 1.7%).

Autodaq and the industry experts with whom we spoke estimated that it costs \$110 to ship a car from the lessor to a physical auction and then an additional \$110 to ship a car from the physical auction to the buyer. The transportation cost to ship a car from the lessor directly to a local buyer was estimated at \$137. The difference reflects the absence of economies of scale in shipping directly.

Autodaq estimates that a dealer travels one hour each way to an auction and buys four cars. This translates to one-half hour of travel time per car purchased. Conservatively valuing a dealer’s time at \$40 per hour⁸ this translates into \$20 per car. Autodaq assumes that a dealer spends five hours at the physical auction. We assume, conservatively, that the dealer does not waste any of these hours at the auction.

Finally, we assume that reconditioning costs are the same for the physical auction as for the Internet auction. This also is likely to be conservative in that cars bought in a physical auction are usually reconditioned again by the buyer after they are bought. To account for this, we have not added any extra time to the Internet process for reconditioning.

Based on these assumptions, we estimate in column (2) that the physical auction process has an economic cost of \$540 per car (not including reconditioning) given the industry estimates of the time

costs in column (7). Using our sample results rather than the industry estimates, we obtain an almost identical cost of \$548 per car (not including reconditioning) in column (4).

Under the assumption of a liquid market and using Autodaq's estimates, the Autodaq / Internet process has an economic cost in column (6) of only \$255 per car (without reconditioning) – a \$285 reduction from the industry estimates. Using the costs implied by the actual 17 days elapsed from inspection to delivery in our sample, the Internet process has a total economic cost in column (8) of \$337 (without reconditioning) – a \$211 reduction from the physical auction sample results. The difference would be at least as large if we measured the costs from turn in to delivery because the average time from inspection to turn in is at least as large for the Internet sample as for the physical auction sample.

Both of the estimated reductions (\$211 from the sample and \$285 from industry estimates) are conditional on both markets being liquid. In the Autodaq sample, this was not the case – the probability of a sale was 24% not 100%. As a result, the (conditional) process savings overstate actual savings.

We estimate the actual savings using the following assumptions. The seller attempts to sell a car over the Internet. If a sale occurs, it occurs in a median 5 days.⁹ In the 76% of cases in which a sale does not occur, the seller decides after 5 days to sell the car through a physical auction process. The car then takes 28 days before it is delivered to a purchasing dealer.¹⁰ For these cars, the lessor incurs 5 additional days of interest and depreciation costs that we estimate to be \$41. In our sample, therefore, relative to the physical auction process, the Autodaq / Internet process provides a 24% likelihood of a \$211 reduction in process costs and a 76% likelihood of a \$41 increase in process costs. The net effect is an average decrease in process costs of \$19 per car.

This analysis highlights that liquidity is important in an Internet market not only to deliver attractive pricing, but also to deliver savings in process costs.

Overall, the results in Table I indicate moderate reduction in process costs for cars sold using the Internet in our sample. The results suggest potentially substantial reductions in process costs as the Internet market becomes more liquid. Not including reconditioning, the reductions in a liquid market of

more than \$200 are on the order of 40% of the total economic cost. Multiplied over an annual market of 5 million cars, the analysis implies potential process cost reductions on the order of \$1 billion per year.

III. Impact on Coordination costs (2): Marketplace Benefits

A second potential benefit of B2B is the extension of the market it provides. Both buyers and sellers can search a larger number of counterparts, and, as a consequence, may find goods and services that they would not otherwise have found. In the case of used automobiles, this seems likely to be an advantage. Used automobile dealers require an appropriate mix of inventory in their dealerships. Obtaining that mix is the main reason they purchase at auctions.

In this section, we attempt to estimate the marketplace benefits in the Autodaq / Internet process. Our goal is to assess how much more a dealer would be willing to pay in the Internet market (versus the physical market) for a car that better matches the dealer's desired inventory. This is not possible to estimate directly because it is not observable. On the other hand, because marketplace benefits are potentially large on the Internet, getting some grasp on the magnitude of this gain is important.

In what follows, we propose a simple method that exploits (1) the geographic rollout used by Autodaq and (2) a no-arbitrage argument on the seller side. Under reasonable conditions, this method places a lower bound on the dealer's willingness to pay for access to the larger marketplace.

Autodaq's rollout followed a predetermined pattern. Between the end of October of 1999 and the end of February 2000, the buyers were almost exclusively in California. The sellers, on the other hand, were three large leasing companies that sold cars coming off lease throughout the US. Cars sold in California in that period, therefore, included cars from the Southern, Midwestern and Western U.S.

The type of sellers implies that the cars were, from their perspective, commodities up to their physical characteristics.¹¹ The willingness to pay by buyers for each car differs widely, as it depends on the quality of the match of the particular car with the needs of the dealership. Suppose that a dealer has a choice between two cars that are from the seller's perspective identical, but that are valued differently by the dealer because of the dealer's particular requirements. Suppose, first, that both of these cars can be

purchased over the Internet, but one is geographically further away. If we observe a dealer buying a car that is not from California, the dealer must have viewed that car as a particularly attractive match in order to incur the additional transportation costs. From a seller perspective, cars are indistinguishable. The difference in transportation costs, therefore, provides a lower bound estimate of the difference in willingness-to-pay for cars that are purchased from out-of-state. This provides an estimate of the marketplace benefit for those cars.

[Table II Here]

In Table II, we report the transportation costs for 586 cars sold in the Autodaq Internet auction. The transportation costs are the actual costs paid by the buyers. Table II shows that transportation costs average \$465 for out-of-state cars and only \$223 for California cars, implying a transportation cost differential of \$242 per car.

It is important to note that we cannot say with certainty how much value was created from this improved matching in our sample. In the extreme, it is possible that the buyer values an out-of-state car at exactly \$242 more than an in-state car and pays the entire differential in transportation costs, leaving the buyer with no surplus. It seems reasonable to argue, however, that with a liquid Internet market, dealers in California will be able to buy cars in California over the Internet and capture more, if not all, of the gains from improved matching.

We also can attempt to estimate the marketplace benefits relative to a physical auction. For a sale to take place on the Internet, the Internet price must be at least equal to the physical auction price less the process cost savings. Thus the sum of the marketplace benefits and the process improvements is at least equal to the difference in transportation costs caused by the additional shipping distance of the Internet auction versus the physical auction. As before, it may be the case that, at the current stage of development of the Internet, the total increase in surplus is small, if the transport costs are equal to the efficiency gains.

In our sample, we can estimate differential transportation costs caused by the additional shipping distance. Autodaq and industry analysts we spoke to estimated that the buyer pays roughly \$110 to

transport a car it buys from a physical auction site to its dealership. As we reported above, Table II shows that average transportation costs are \$465 for cars transported from out of state to dealers in California. This suggests that the average car from out-of-state purchased on the Internet is transported a much greater distance than the average car purchased at a physical auction. The extra transportation cost of \$355 suggests that the sum of the marketplace benefit and process cost reductions exceeded \$355 on average for out-of-state cars sold on the Internet.¹² Again, in a more liquid market, the distance required to obtain improved matching should decline, and more of the benefit should accrue to buyers and sellers.

Overall, our results suggest substantial marketplace benefits to the Internet auction in the wholesale used car market. These benefits are potentially of the same order of magnitude as the process improvements.

IV. Asymmetric Information in Physical and Internet Automobile Auctions

(i). Quality Information in Car Auctions

While in a physical auction a buyer can obtain an independent indication of the condition of the car (by self-inspection), the Internet auction relies exclusively on information that can be observed in the database. As a consequence, informational asymmetries between sellers and buyers may be more pronounced in Internet auctions.

In the wholesale used car market, however, the potential informational loss may be small as information in physical auctions is usually restricted. In describing the physical auctions, Genesove [1993] points out that bidders have limited access to the cars:

“Prior to the bidding, the car is parked outside, where potential bidders can examine its exterior. They are prohibited from opening the doors or raising the hood. Mileage and options are chalked on the car’s windows. When the car’s turn approaches, it is driven into the appropriate lane and then, before bidding is concluded on the previous car, driven up to the auction block. Now the hood is raised and dealers are permitted to enter the car. There is time to check the odometer, to ensure that the air conditioner works (but, in the summer months at least, not the heater) and to take a look at the running motor. But there is no opportunity to test the brakes or any number of other things that a consumer might check out in a drive around the block (...)

On top of the auction block stands the auctioneer and, beside him, the seller, who under the rules of the auction must be present. The auctioneer announces any major defects in the car, of which the seller has informed him. Bidding is oral and ascending. When bidding will go no higher, the seller is asked to accept or reject the winning bid. About 60% of the time he accepts. The car will have been driven away before the bidding is concluded. From the time it arrived at the auction block until the time it is driven away, a minute and a half will have passed.”

Internet-based auctions such as those run by Autodaq, on the other hand, do not allow any physical inspection of the cars by the buyer.¹³ Instead, the seller and the third-party inspection made available by Autodaq provide extensive information on the car’s options and all other measurable aspects of the car condition, such as its mileage, the damages suffered, age etc. Importantly, Autodaq does not preclude buyers and sellers from participating in physical auctions. This raises the possibility of sellers offering only those cars that are in a relatively worse unobservable condition through this channel.

Possibly attenuating adverse selection in our data is the fact that Autodaq is primarily directed at lessors and fleet owners. Individual used car dealers have only recently started selling cars. Only 571 out of 3552 cars auctioned in our sample, and 111 out of 864 cars sold where auctioned by a dealership.

To understand the implications of the coexistence of these two markets, we take as our starting point a variant of the simple model of adverse selection of Akerlof [1970]. Suppose that, conditional on all the observable characteristics, there are two types of cars, G (good) valued by consumers at P_G and lemons L, valued at P_L with the proportion of good cars sold in a particular market given by q .¹⁴

First, consider the physical market. Assume there is no asymmetric information in the physical market, so that good cars can be sold at price P_G and lemons at price P_L there. Suppose the higher cost of the physical market mechanism is C , so that the value of the sale to the seller is $P_G - C$ if the car is not a lemon, and $P_L - C$ if it is. The average price of cars sold in the physical market is then $P_p = q P_G + (1-q) P_L$.

Now introduce a competitive electronic market. Here, both classes of cars cannot be distinguished, as consumers cannot physically inspect the cars. There are two types of outcomes in this market, depending on the cost of the informational asymmetries relative to the benefit of using an electronic market medium:

1) When the cost of the physical market mechanism is high enough relative to the asymmetric information costs, so that the average price is higher than the net profit from selling a known good car in the physical market, i.e. if $q P_G + (1-q) P_L > (P_G - C)$ or, equivalently, $C > (1-q)(P_G - P_L)$, both types of cars are sold in the electronic market, at a price $P_e = q P_G + (1-q) P_L$. In this case, the ratio of average physical market price to electronic market price is 1.

2) If the cost imposed by the presence of lemons on the sellers of good cars is higher than the gain from using an electronic market C , or formally if $C < (1-q)(P_G - P_L)$, no transactions of good cars take place, as good cars are withdrawn and sold in the physical market. In this case, adverse selection exists in the electronic market. The observed average market price, reflecting the lower average quality of cars transacted, is $P_e = L$.

Adverse selection translates in this case to the withdrawal from the electronic market of cars with relatively good unobservable characteristics within each class in favor of the physical world auction. If adverse selection is present, we would expect to see a lower average price, conditional on observable characteristics, for cars sold over the Internet. High quality cars for each level of observable characteristics would have a low probability of being sold,¹⁵ given that the seller would demand high average prices for the average condition that buyers expect to find in the market.

Apart from this implication for relative price *levels*, adverse selection also has implications on the price *structure*. If adverse selection is important, Internet prices will be lower relative to physical world prices when adverse selection risk is larger. When a low risk of adverse selection exists, i.e. when the variance in the condition of cars is small, the difference between the physical world and the Internet prices will be small. On the other hand, when the adverse selection risk is high, this spread will be large. This different risk is to a large extent predictable. The variance in the unobservable condition of the car is largely a consequence of the unobservable care by the owner. Thus the more that the quality of care affects the value of the car, the larger the risk of adverse selection.¹⁶

Genesove's [1993] study of adverse selection in used car markets is the most notable precedent for our research. He tests for adverse selection by analyzing the effects of the identity of the seller on

prices. He expects systematic differences between the incentives of used and new car dealers to sell used cars to show up in differences in prices if they are selling different quality cars. Our study differs from his in that, rather than examining adverse selection in one market, our focus is on comparing adverse selection in two different markets where we expect, a priori, to find different degrees of informational asymmetry in them. However, we rely on Genesove's insights to examine the extent of adverse selection in the Internet market in itself.

(ii). *Data*

Our sample consists of 3552 sold and unsold cars on auction at Autodaq for a period in 1999 and 2000. These are all of the cars that were auctioned by Autodaq at least once in the period we study, except for those that were withdrawn by their owners without completing a three-auction cycle.¹⁷ For most of our sample period, all cars were put through a maximum of three one-day auction cycles¹⁸.

The construction and content of most of the variables in our sample is self-explanatory. One exception is the ratio of Internet to physical price, which is intended as a proxy for "how much a buyer would have been willing to pay for this car in the physical world." We construct this as the ratio of two variables: the price at which a car was actually sold; and the price at which an average car with similar observable characteristics was sold in that month in physical auctions. Autodaq provided these estimates using as a complete data set of physical auction sales.

While matching these prices generates rich information, the inferences we can draw from the matches are limited, because the matches are not as precise as we would wish. In particular, the algorithm takes into account motorization, drive (2-wheel, 4-wheel), style, model, and model year, but does not differentiate the matches by mileage and option data. For this reason, part of our analysis also relies on the wholesale Kelley Blue Book (KBB) prices of the cars in the sample. The KBB is an industry guide of wholesale and retail prices for vehicles. The KBB price uses the full physical observable characteristics of the car (year, make, model, series, engine, drivetrain, options, mileage, etc.). We do not use the KBB for any analysis of price levels, but only for our analysis of the changes in relative prices in response to changes in the cars' physical characteristics.

[Table III Here]

Descriptive statistics for our sample are given in Table III. The table shows that 24.3% of the 3552 cars offered for sale on the Internet were sold over the period at an average sell price of roughly \$13,600. Estimated physical auction values for the cars were available for a subsample of 3001 observations. The average physical auction value for these cars was \$14,200, while the average KBB value for these cars was substantially higher at \$15,500. According to Autodaq and the industry sources we spoke to (including one competitor), the roughly 10% differential is an industry standard.¹⁹ Share of damages is the ratio of the estimated dollar value of damages suffered to the book value of the car.

Our main objective is to use the data to measure the importance of informational asymmetries in Internet auctions. We try to do this using three elements in our data: the difference in price levels between the Internet and the physical world, the structure of relative prices, and the actual probability of sale of individual cars on the Internet. We expand on our use of these three pieces of information below.

(iii). *Empirical Specifications and Results*

1. Relative Price Levels. Assuming that the physical and Internet markets are competitive, we expect to see lower prices *relative* to the physical market when the quality of the cars sold on the Internet is worse than the quality of the cars sold in the physical real world. We can directly test this implication by comparing the average prices attained by the auctioned cars in the Internet market with the average price they would have attained had they been auctioned in the physical market.

[Table IV. Here]

Table IV presents this test. The data in the table reject the hypothesis (at the 95% level) that cars attain lower prices over the Internet than they would have attained in the physical world. In fact, the data suggest that the Internet prices are significantly *higher* than the prices in the physical world.

There are two caveats to this interpretation. First, the price in the physical market is for an average car within a model-year-motor-drivetrain cell. A finding that the average price of a car in the Internet is higher than in the physical market could be compatible with adverse selection, if cars in the

Internet have less mileage for a given number of years driven, or if they have more options. Regrettably, the micro-data on the physical market do not allow us to draw such distinctions. Second, the differences in average prices may respond to factors other than the average quality of the car sold. In particular, given the lower transaction costs using the Internet, dealers may be willing to pay higher prices for Internet-based transactions. For this reason, we think the evidence on the structure of relative prices in both markets is a better gauge of the extent to which adverse selection matters. We turn to this issue next.

2. Relative Price Structure. As noted above, adverse selection between the physical market and the Internet is more likely to be a problem the higher the proportion of “lemons” and the lower their value. Accordingly, if adverse selection is a problem, the effect should increase as the variance of a car’s condition increases. As the variance of a car’s condition increases, the adverse selection should manifest itself in a larger decrease in the Internet price than in the physical world price.

We test for adverse selection by assuming that conditional on model characteristics, the variance of the value of a newer car, of a car with low mileage or of one with a good observable condition is likely to be small. In contrast, the variance of the value of an older car, or one with more miles will vary more depending on the care taken by its user. In other words, if the quality of care can only be (partially) gauged from direct observation of the car, the cars in the second category, for which care is likely to matter more, will be subject to relatively more important adverse selection problems over the Internet.

To test this hypothesis, we could turn again to the most direct data available, i.e. the price on the Internet relative to the price in the physical world. Because of the limitations in matching Internet prices to physical auction prices, we turn to the KBB prices as our proxy for the price that the car would have attained in a physical auction; these prices do differentiate cars by mileage, condition and options. As long as the relation between the KBB price and the price that an average car with the same characteristics would have obtained in the physical auction is constant, this is appropriate for our purposes.²⁰ We test the hypothesis that as the variance in the condition of the car increases – as given by mileage, age and percent of damages suffered by the car – the Internet price relative to the physical price decreases.

[Table V Here]

The first set (a) of specifications in Table V presents the OLS evidence conditional on the car being sold. The evidence on the existence of more adverse selection in the electronic auction is mixed. Holding constant the KBB price, a car does not lose significantly more value in the Internet than in the physical market. Similarly, higher mileage does not appear to decrease the value of the car more in the Internet than in the physical market. In fact, the Internet price declines significantly more slowly with miles than the KBB price would predict, rejecting the hypothesis of adverse selection. On the other hand, each car-year reduces the Internet price by around 2 percent points more than the KBB price, suggesting that this could be a mechanism through which adverse selection is observed.

A problem with those specifications is that the Internet sample is censored by the reserve price, as we do not observe transaction prices for unsold cars. For this reason, the second set of specifications (b) in Table V repeats the analysis using all of the observations in the sample, with a censored normal regression where the reserve price is the censoring point. The evidence in favor of adverse selection in these specifications is also weak. Including seller fixed effects, the age and the mileage do not appear to reduce the price that a car could attain on the Internet relative to the price that would attain in a physical world auction as measured by the KBB price.

3. Identity of the Seller. Most cars sold in the Autodaq auctions are sold by leasing companies or rental car companies. For the last part of our sample period, however, individual dealers sold cars on the Internet. Following Genesove [1993], we exploit the difference in incentives between the three types of sellers (dealers, leasing companies and rental car companies) to uncover evidence of adverse selection. We expect adverse selection to be most important for cars sold by individual dealers. Dealers have greater incentives and opportunities (1) to check the quality of care and condition of each individual car, and (2) to select those cars to sell on their lot and those to sell on the Internet. As a consequence, after controlling for physical characteristics of these cars, we expect the identity of the seller of the car to matter. Individual dealers should obtain lower prices for their cars, holding everything else equal.

The effect of rental car companies is more ambiguous.²¹ Holding all else constant, rental car company cars have been through many more users. Their unobservable quality should be lower and, as a result, their average price should be lower. On the other hand, selection should be less important for rental cars, as rental car companies have a policy of selling all their cars after some fixed period of time.

The set of specifications (c) in Table V tests these hypotheses. The regressions reject the hypothesis that individual dealers are perceived to sell lower quality cars over the Internet than institutional sellers. The first column controls only for the book price of the car, and shows a significant effect of individual dealer on price, but exactly of the opposite sign as the one predicted by the theory. Controlling for a car's physical characteristics, the effect decreases, but the dealer effect is still positive and significant. Holding all physical characteristics constant, a car sold by a individual dealer earns a premium over the KBB price that is 4% higher than that of the leasing companies (the excluded category).

The lack of evidence of adverse selection in the dealer market is important for another reason. One might argue that our results for lessor or fleet sales are biased because (1) the buyers know the identity of the sellers and (2) lessor sellers get good prices for their cars in general. The similar results for the dealer market suggest that our results are not biased for this reason.

4. Probability of Sale. A final hypothesis concerns the probability that a car is sold. This contains no information on the comparison of adverse selection in the Internet market relative to physical markets, but may contain some information on whether adverse selection exists at all. Clearly, adverse selection implies that cars with good (online) unobservable condition should be relatively less likely to sell on the Internet market. For this difference to be translated into an actually lower probability of sale, however, it would be necessary that cars in relatively good unobservable condition do 'show-up' in the Internet market, likely with a higher reservation price, only to be later withdrawn from auction. In our data, it is possible to assume that sellers initially attempt to sell all cars that they are planning to sell in auction on the Internet. First, that is the arrangement between the firm in our study and the sellers. Second, the

sellers can choose a high reservation price for even their best-conditioned cars. As a result, there is at worst only a small opportunity cost of trying the Internet.

Accordingly, if adverse selection is important on the Internet, we expect that the greater the variance in a car's condition (which would increase the incidence of adverse selection), the lower the probability that the auction is successful. In our sample, cars with more miles, higher age, and a history of more accidents in the past would have a lower probability of sale if adverse selection exists.

The alternative hypothesis is that adverse selection is not a particular problem in Internet markets. A reason for this in Autodaq's case is that individual dealer sales over the Internet were limited and under stringent conditions. To the extent that such dealers are more likely to try to dump their lemons, it may actually be that the Autodaq market is more, rather than less, efficient than the physical auction market.

The final set of specifications (*d*) in Table V present the analysis of the probability that a sale actually takes place using a Probit model. The evidence is inconsistent with the existence of important adverse selection in these markets. Damaged cars do seem to be somewhat less likely to be sold, but neither older cars nor cars with more miles are less likely to be transacted. In fact, there is a significantly positive effect of the age of the car on the probability of a sale.

Caution must also be exercised in interpreting these results. Only if dealers do not withdraw their 'good condition' cars prior to sale, but rather post them with a higher reservation price, do we expect to see adverse selection manifested in a lower probability of sale for cars more affected by adverse selection.

5. Is Adverse Selection a problem over the Internet? Overall, we find little evidence consistent with the hypothesis that adverse selection is more pervasive in the Internet market than in the physical world, or even that adverse selection is a problem at all in the Internet marketplace we study.²² As we observed before, this conclusion necessarily must be qualified by the measures that Autodaq has taken to reduce the incidence of adverse selection in this particular instance.

V. *Conclusion*

Our results suggest that process improvements and marketplace benefits are potentially large – on the order of 5% of the automobile value and a much large fraction of the total transaction cost. Moreover, we do not find evidence that the Internet increases adverse selection costs. We suspect that this is at least due in part to the measures taken by Autodaq to avoid this problem.

It is worth adding that it is likely that the direct changes in transaction costs that we focus on here understate Internet-based efficiencies. This is so because as the costs of undertaking spot market transactions decrease, participants in these transactions are likely to adjust their behavior and realize further efficiency gains.

Two important issues are left for future research. First, we have not found a way to study the effects of the impact of the Internet on commitment problems. Second, future research must consider the consequences of these transaction cost changes for the behavior of agents – the indirect efficiencies. In particular, it remains an open question as to the effects of Internet markets on vertical integration in ‘bricks and mortar’ companies.

References

Akerlof, George, 1970, "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics* 222: 488-500.

Bond, Eric W, 1982, "A Direct Test of the "Lemons" Model: The Market for Used Pick-Up Trucks" *American Economic Review* 72 (4): 832-80.

Genesove, David, 1993, "Adverse Selection in the Wholesale Used Car Market" *Journal of Political Economy* v101, n4: 644-65

Milgrom Paul and John Roberts, 1992, . *Economics, Organization and Management*. Englewood Cliffs, NJ: Prentice Hall.

Table I
Process Cost Of Physical Auction Versus Internet Auction Process Per Car

Process cost of physical auction versus Internet auction based on used car auctions from 1999 to 2000.

	Physical Auction				Internet Auction			
	Industry Estimate		Sample Results		Autodaq Estimate		Sample Results	
	Time (days) (1)	Dollars (2)	Time (days) (3)	Dollars (4)	Time (days) (5)	Dollars (6)	Time (days) (7)	Dollars (8)
Wait for pick-up	9 to 15		N.A.		0		N.A.	
Ship to auction	2		N.A.		0		N.A.	
Ready for sale	10		N.A.		2		N.A.	
Ready for sale until sale	5 to 15		N.A.		2		N.A.	
Ship to dealer	2		2		3		3	
Total Time	28 to 44		37		7		17	
Capital Cost of Time (36 days)		\$107		\$110		\$21		\$51
Depreciation Cost of Time (36 days)		\$188		\$193		\$37		\$89
Inspection Cost		\$ 5		\$ 5		\$60		\$60
Shipping Cost		\$220		\$220		\$137		\$137
Reconditioning Cost		X		X		X		X
Dealer Travel Cost	0.5 Hours	\$20	0.5 Hours	\$ 20				
Dealer Time Not Bidding	0 to 4 Hours		0 to 4 Hours					
Total Economic Cost		\$540 + X		\$548 + X		\$255 + X		\$337 + X
Seller Fee		\$100		\$100		\$100		\$100
Buyer Fee		\$200		\$200		\$175		\$175

Table I (continued)
Process Cost Of Physical Auction Versus Internet Auction Process Per Car

Assumptions:

Times:

Industry Estimates: Industry estimates for physical auction were obtained from Tom Kontos at ADT Automotive and confirmed by other sources. Autodaq estimates for Internet auction were provided by Autodaq.

Sample Results: Sample median for physical auction is time from inspection to time of sale for 9205 cars sold by lessors through physical auction process augmented by two days for dealer shipment. Sample median for Internet auction is time from inspection to time of sale for 694 cars sold through Autodaq process augmented by three days for dealer shipment.

Wait for pick-up is time from lessee delivery of car to dealer until car is picked-up by physical auction.

Ready for sale is time from delivery at physical auction site to the time car is ready for sale. Includes time to recondition.

Used cars in our sample have an average sale value of \$13,600. Interest rate / cost of capital assumed to equal 8%. Each day, therefore, costs seller $8\% \times \$13,600 / 365 = \2.98 per day in capital costs.

The table assumes that used car values decline or depreciate in value by 14% per year. Each day, therefore, costs seller $14\% \times \$13,600 / 365 = \5.22 per day in depreciation costs / forgone sales price. In the data provided by Autodaq, the sales price declines by 14.8% (with a standard error of 1.7%) per year.

Inspection cost for physical auction assumes 15 minutes at a cost of \$20 per hour; for Autodaq, is the cost to Autodaq.

Dealer travel cost assumes that dealer travels a total of two hours and buys four cars for an average of 0.5 hours per car. Dealer time is valued at \$40 per hour.

Shipping cost is two shipments at \$110 each for the physical auction; one shipment at \$137 for the Autodaq process. Based on Autodaq and industry interviews.

Reconditioning costs assumed to be the same for both processes.

Table II
Transport Costs Incurred by Buyers

Transport costs incurred by buyers in Autodaq Internet auctions from 1999 to 2000. Transport costs are actual transportation costs paid by purchasers. Estimated transportation cost in physical auction provided by Autodaq.com and corroborated by interviews with industry participants.

	Transport Cost	Number of Cars
Cars transported from outside of California in Autodaq Auction	\$465	411
Cars transported within California in Autodaq Auction	\$223	175
Difference	\$242	
Estimated Transport Cost in Physical Auction	\$110	

Table III
Descriptive Statistics

Descriptive statistics for Autodaq Internet auctions from 1999 to 2000. Sold equals one if a car put up for sale in an Autodaq auction was sold. Internet price is the price the car sold for in the Internet auction. Estimated physical auction value is the value the Autodaq estimates the car would have sold for in a physical auction based on data from physical auctions. Book price is the price of a similar car according to the Kelley Blue Book.

Variable	Obs	Mean	Std. Dev.	Min	Max
Vehicle Sold	3552	0.24	0.43	0.00	1.00
Internet Price [\$1000]	865	13.6	5.3	3.1	32.7
Estimated Physical Auction Value[\$1000]	3001	14.2	5.6	0.0	74.9
Kelley Blue Book Price [\$1000]	3147	15.5	6.2	3.6	79.5
Mileage (10000)	3552	3.57	1.43	0.196	10.83
Age (2000 - year)	3552	2.15	1.50	0.00	7.00
Share of damage (damages/KBB)	3147	0.01	0.02	0.00	0.18
Auction Date (100s days)	3552	2.09	1.11	0.01	3.57
Seller = Rental Company	3552	0.30	0.46	0	1
Seller = Individual Dealer	3552	0.05	0.20	0	1
Seller = Leasing Company	3552	0.65	0.48	0	1

Table IV
Price Levels: Internet Auctions Versus Physical Auctions

Price level of internet auction relative to physical auctions Autodaq Internet auctions from 1999 to 2000. Internet price is the price the car sold for. In the Autodaq auction. Estimated physical auction value is the value the Autodaq estimates the car would have sold for in a physical auction based on data from physical auctions. The data correspond to a matched sample of internet sales with the price of an average car of identical model year, motorization and drive in physical auctions. * is significantly larger than 1 at the 5% level.

Model Year of Manufacture	Internet price/ physical auction value	N(number of Internet Cars)	95% Confidence Interval	Interval
1995	1.176* (0.035)	31	1.104	1.247
1996	1.114* (0.022)	80	1.07	1.158
1997	1.018* (0.005)	418	1.008	1.028
1998	1.049* (0.014)	42	1.02	1.078
1999	1.062* (0.021)	23	1.019	1.105

Table V
Performance of Internet Auctions

Specifications are for Autodaq Internet auctions from 1999 to 2000. Internet price is the price the car sold for. Book price is the price of a similar car according to the Kelley Blue Book. Damages is the ratio of the estimated damages to the book value of the car.

<i>Independent Variables</i>	<i>Dependent Variables</i>							
	(a) Log Internet Price (OLS)		(b) Log Internet Price (Censored Normal Regressions)		(c) Log Internet Price (OLS)		(d) Probability Car Sold (Probit)	
Log Blue Book Price	1.05** (0.01)	1.05** (0.01)	1.01** (0.01)	0.99** (0.01)	1.06** (0.00)	1.05** (0.01)		
Mileage (0000)	0.01** (0.00)	0.01** (0.00)	-0.01** (0.00)	0.00 (0.00)		0.01** (0.00)	-0.02 (0.02)	0.01 (0.02)
Share Damages	0.13 (0.19)	0.32 (0.19)	-0.40* (0.16)	-0.28 (0.15)		0.29 (0.19)	-5.35** (1.6)	-1.47 (1.63)
Age	-0.02** (0.00)	-0.02** (0.00)	0.04** (0.00)	0.00 (0.00)		-0.02** (0.00)	0.28** (0.03)	0.17** (0.04)
Auction Date (100s days)	0.03** (0.00)	0.01 (0.01)	0.02** (0.00)	0.01 (0.01)		0.02** (0.00)	-0.22** (0.03)	-0.72** (0.07)
Seller Fixed Effects		Yes		Yes				Yes
Seller = Individual Dealer					0.10** (0.01)	0.04** (0.01)		
Seller = Rental Company					0.05** (0.02)	-0.04 (0.02)		
Constant	-0.30** (0.03)	-0.25** (0.03)	-0.31** (0.03)	-0.11** (0.03)	-0.31** (0.02)	-0.28** (0.03)	-0.81** (0.11)	0.55* (0.22)
R Sq (N)	0.96 (843)	0.96 (843)			0.96 (843)	0.96 (843)		

Standard Errors In parenthesis * Significant at 5% level; ** Significant at 1% level;.

¹ Milgrom and Roberts [1992], p. 28. Chapter 2 also includes a discussion of coordination and motivation costs.

² See the “2000 Used Car Market Report,” produced by ADT Automotive, the number two player in physical auctions.

³ This process is not unique to Autodaq. Several competitors exist. In particular, the largest operator of physical auctions, Manheim, has an Internet based subsidiary – Manheim Online. Manheim Online differs from Autodaq in that it uses the Internet to list the cars that it has for sale at its physical auction site. In its current incarnation, therefore, Manheim Online, potentially reduces buyer transaction costs, but does not change seller transaction costs.

⁴ ADT has subsequently merged with Mannheim, the largest competitor in they physical auction business.

⁵ For 7,221 cars, we can measure the time from lease end to sale. The median time from lease end to sale in our sample is 36 days (compared to 35 days from inspection).

⁶ It is important to note that in both the Autodaq system and the physical auction, buying dealers typically perform reconditioning (despite the fact that some reconditioning is performed by the physical auction). It is possible that the reconditioning time is greater in the Autodaq system, although Autodaq claims that this is not the case.

⁷ For the 270 cars with a lease-end date, the median time from lease end to sale is 20 days.

⁸ This is conservative as mechanics probably cost more than this.

⁹ This assumes that the car is inspected 9 days before it is turned in.

¹⁰ Again, this assumes that the car is inspected 9 days before it is turned in.

¹¹ In this analysis, we assume that adverse selection problems are absent. Our analysis below confirms that adverse selection is likely not an issue here.

¹² This calculation also is conditional on an interstate sale.

¹³ Both Autodaq and the physical auctions do inspect the cars and describe them for buyers. Autodaq argues (and we agree), that the information in electronic form is richer and more useful as it allows buyers to search more efficiently for their desired cars and options.

¹⁴ Consistent with our previous discussion, the entire surplus is captured by the seller. We now ignore the transport costs considerations to simplify the discussion.

¹⁵ Note that the seller could just not bring high quality cars to the auction block. Given that the cost of merely posting the car on the electronic market is very low, and that the reserve price can be used to avoid selling it cheap, we can expect even cars with very good unobservable characteristics to be posted. In fact, the opposite is likely to have occurred. Autodaq screened out (i.e., did not list) cars with excessive mileage and excessive known damage.

¹⁶ There is another type of adverse selection in this market, unrelated to the quality of care, but related to the quality of manufacturing. This is unlikely to be an issue here for two reasons: first, initial defects are not frequent; and, second, all the models are less than 4 years old so manufacturer warranties typically cover such defects.

¹⁷ None of our results are sensitive to including those cars as not sold.

¹⁸ This constraint was lifted later in the sample, but to little effect: only six cars were sold after three cycles.

¹⁹ The difference between book value and selling price and its magnitude are also present in Genesove [1993]. He finds that the book value is an imperfect predictor of the selling price, but does not document any systematic relation between the bias and the age, mileage or other characteristics of the car.

²⁰ Note that the use of the KBB rather than the average physical auction price biases the results in favor of finding that adverse selection is more pronounced in the Internet than in physical markets. This is true because more information goes into non-auction sales in the physical world that are the bases for the KBB than at auctions. Using the estimated physical auction car of an equivalent car rather than the KBB as the counterfactual does not affect our results.

²¹ We thank the editor for pointing out this ambiguity.

²² In turn, previous studies of physical motor vehicle market by Bond [1982] and Genesove [1993], have found little evidence of adverse selection in these markets.