

The Effect of Information on Auction Outcomes: A Large Scale Field Experiment*

Steven Tadelis
UC Berkeley
Haas School of Business

Florian Zettelmeyer
Northwestern University
Kellogg School of Management

September 9, 2009

Abstract

When choosing auction mechanisms sellers can decide how much information to reveal to buyers regarding the quality of the goods sold. Using a field experiment in a market for wholesale automobile auctions we are able to measure the effects of information on auction outcomes. We create random assignments of information about quality, and manipulate the availability of information over time. Our results suggest that, as the theoretical literature generally predicts, more information increases expected revenues. Furthermore, by measuring the effects on different quality grades of automobiles it seems like the increase in revenues are due to more competition for any given vehicle. Finally, we quantify the value of gathering information and releasing it to potential buyers in this setup. *JEL* classifications C93, D44, D82, L15

** VERY PRELIMINARY AND INCOMPLETE **

*We are grateful to the management and employees of the firm that provided the data and worked cooperatively with us to implement the experiment. We also thank ... for helpful discussions. Tadelis thanks the National Science Foundation for financial support.

1 Introduction

The past two decades have seen tremendous growth in the use of auctions both in the public and private sectors. (examples, \$\$, etc.) As a consequence, many academics and practitioners have focused attention on the design of auctions to achieve efficient allocations, to maximize revenues, or both.

Though much of the discussion and analysis has been focused on the rules of the auction (open or sealed, first or second price, free entry or invited bidders, etc.), one dimension of auction design that has received less attention is what information to make available to the bidders.

The often celebrated “linkage principle” (Milgrom and Weber, 1982) proves that under rather common conditions, if the seller releases more information the his expected revenue from the auction will increase. Two policy conclusions emerge: First, for many items it is in the seller’s best interest to commit to release as much credible information as he can. Second,

In this paper we document the effect of releasing more information in a large number of business auctions.

Due to the challenge of testing how variation in information affects auction outcomes, there have been few empirical studies of this important question. Two notable ones are Kagel and Levin (1985) and De Silva et al. (2008). ...[KL: indirect] [De Silva: policy change] Our paper contributes in an important way by actually implementing a randomized field experiment in which the amount of information differs yet all other aspects of the auction remain fixed.

Our result strongly support the hypothesis that more information will increase the expected revenues obtained by the sellers. More interestingly, the increase in expected revenues holds true across all quality levels. That is, it is not the case that sellers of low quality vehicles will obtained lower expected revenues, while on average expected revenues will increase. Instead, all quality levels show at least a weak increase in expected revenues, and this is strongest for the very best and very worse quality levels.

Our results suggest that the main driver behind he increase in expected revenues is an increase in competition, which may be a result of better matching between buyers and types of vehicles.

The paper proceeds as follows: Section 2 describes the industry and discusses the implications of the current theoretical literature. Section 3 describes the data and the experimental design, and section 4 presents the experimental results. Section 5 offers a discussion of the results and concludes.

2 Background

2.1 Wholesale Auto Auctions

The U.S. retail market for used cars is sizeable. Estimates place used car sales at more than 30 million cars in 2008, most of which are sold by franchise or independent dealers.¹ Dealers of used cars sell on the retail market and generally purchase their inventory of used cars either from trade-ins, or from the wholesale market for used automobiles.

A prominent source of used vehicles comes from wholesale automobile auctions. In fact, according to the National Automobile Dealer Association (NADA), 35% of all used vehicles sold by new car dealers in 2008 were sourced in such auctions (see NADA DATA, 2009)². Most auctions are administered by a few prominent auction houses that specialize in this market, one of which provided the data for this study.

The buyers in our auction are exclusively dealers, while the sellers mainly belong to one of three categories: dealers who wish to change their inventories; owners of large fleets such as rental car agencies who periodically turn over their inventories; and financial lease agencies who sell vehicles for which a lease contract had ended. Sellers bring their vehicles to the auction one or more days in advance of the actual auction where they are registered and assigned a “lane” number and a “run” number. On the day of the auction the vehicles are lined up in several (up to 12) lanes, according to the registered numbers.³ Several thousands of vehicles will be auctioned off during a sale day.

Before the auction day begins, potential bidders receive information about the lane and run numbers of each car that will be sold at the auction, as well as basic information about the vehicle such as make, model, year, color, mileage and other features. This allows buyers to determine which cars they want to bid on. The information is available online before the auction commences, and a printout is prepared for buyers on the morning of the auction.

At the end of each lane is an auction block where the auctioneer conducts the auction, one car at a time for that lane, so that up to 12 auctions can occur simultaneously. The vehicle which is next in line to be sold is slowly driven⁴ to the auction block where it stops, amid several potential buyers, and is left idling as the auctioneer begins the auction. The auction is an ascending oral (English) auction that lasts for about 45 seconds.⁵ The auction ends when no bidder is willing to raise the price, and if the price exceeds the seller’s reserve price then the sale is consummated.

¹**** refer to NADA, NIADA and other sources if possible.

²This is available at <http://www.nada.org/Publications/NADADATA>.

³For example, a vehicle with a lane-run number of 9-132 will be auctioned in lane 9, and will be the 132nd vehicle in the lane.

⁴Some cars that are not in driving condition are towed.

⁵Interestingly, the auctioneer begins at a very high price, often above the winning bid, and then works his way down until some bidder signals his willingness to buy. This sounds like a Dutch auction but it is not: the first bid is not the winning bid, but instead determines the start of the ascending bid process. This procedure has been in place for decades (see Genesove, 1995 p.26), and we have been told that it is also common in livestock auctions. We were unable to get an answer as to why this procedure is used.

Otherwise, the vehicle either returns to the seller’s lot or is left at the auction lot for a subsequent sale day.

There is a major difference between the way fleet-sellers and dealer-sellers set reserve prices. Fleet-sellers will sell a large number of cars in one sale day (we witnessed one lease agency bringing in over 800 cars), and will have a representative sitting with the auctioneer and determining in real time whether or not to accept the highest bid. This suggests that the reserve price may have some real-time input. Dealer-sellers, however, bring in a handful of cars and are often not present at their cars’ auctions. They determine their reserve prices in advance and convey it secretly to the auction house. The auction house will then inform the high bidder if the sale is accepted.

There are two distinct classes of bidders at the auction. “Lane” bidders are those bidders who are physically present at the auction and visually inspect the car from up close. Prior to the bidding, vehicles are parked outside so that potential bidders who arrive early enough can examine their exterior condition. The second class of bidders are “online” bidders who are able to participate in the auction through an Internet webcast, which provides streaming audio and video of the auction in real-time. These bidders have online access to basic information about the vehicle, e.g., make, model, year, color, mileage and other features.

2.2 Information and Standardized Condition Reports

As the description above suggests, buyers have some information about the vehicle at the time of the auction, including both basic information and, for the lane bidders, a close visual inspection of the car (including listening to the engine of those cars that can be driven.) Since it is not possible to perform a serious inspection of the vehicles by the potential buyers (not to mention the disadvantage of the online bidders who cannot themselves see the vehicles in any detail), there is residual uncertainty about vehicle’s quality. As a response, many auction houses offer some form of condition reports that describe in more detail what the condition of the vehicle is. Historically, fleet-sellers have requested some tailor-made condition reports for the vehicles they sell, but dealer-sellers have not followed suit. Also, the output from these tailor-made reports was not standard, and buyers were not always pleased with the representation of the information.

In response, the auction house from which this paper’s data originates had developed a Standard Condition Report (SCR) that is aimed at offering a standard set of inspections, and a standard way in which to present the information. The SCR is based on a detailed inspection of the vehicle’s exterior, documenting all imperfections (including whether there is an additional layer of paint that implies some previous damage.) The interior condition is also carefully documented, as is any visual damage to the chassis. The inspections do not include the mechanical condition of the car but the technician who executes the inspection documents whether the engine has any peculiar sounds. The technician enters all of the information through a computerized hand-held device that registers the information on a central computer, and creates a standardized report.

The SCR is then posted online in a standard one-page format. Aside from documenting

a detailed summary of the inspection, two other summary statistics are generated. First, a “condition grade” (CG) is calculated based on the input of the inspection.⁶ The grading system is from 1 through 5, with increments of 0.1, where a CG = 1 is considered “rough”, and a CG = 5 is considered “clean”. Second, the SCR calculates the expected number of labor hours needed for a body-shop technician to correct the reported damage, as well as the cost of the materials needed. Using a standard hourly labor rate this translates into the cost of bringing the vehicle to a condition where exterior and interior damage are no longer noticeable. Hence, both the CG and the estimated costs are standardized measures of vehicle quality.

2.3 The Effect of Quality Information: Theory

A question that has been troubling the auction house is whether the information revealed by the SCR is valuable to the buyers, and if in turn, it translates into more consummated sales and higher expected revenues for the sellers. (The fees collected by the auction house depend both on whether a sale occurred and on the sale price.) In theory, this question can be answered by the well known “linkage principle” that was demonstrated in the seminal work of Milgrom and Weber (1982) (henceforth, MW).

In an affiliated values auction model, MW show that if the seller releases more information to the bidders about the good he is selling then the expected revenue of the auction will go up.⁷ The intuition for this result is not straightforward, yet it mostly follows from creating more competition (see, e.g., Klemperer 2004). Namely, by offering more information, the assessments of the bidders become more congruent, resulting in lower information rents for the winning bidder. The linkage principle derived in MW, however, does not always extend to an auction environment in which multiple units of a good are being sold at once, as shown by Perry and Reny (1999). In many ways, the auction setting that we study is closer to environment of Perry and Reny, and hence it is not clear what the effect of more information will be on auction outcomes.

The empirical implications of the linkage principle go beyond simply stating that more information will increase expected revenues, or that, as a consequence, the revenue ranking of auction formats will relate to the amount of information revealed during the auction. Another implication of the linkage principle is that, given a fixed set of bidders, if the information revealed is favorable then the expected revenue should increase, while if the information revealed is unfavorable then expected revenue should decrease.⁸ Hence, if our setting satisfies the assumptions under which the linkage principle is true, then when SCRs are revealed two empirical predictions emerge:

LP1 Expected revenues should increase.

⁶**** refer to Genesove and Overby who also mention CRs. Explain how SCRs differ due to standardizations.

⁷To be precise, the information revealed must be affiliated. That is, once it is revealed, the valuations of the bidders move closer to each other in a statistical sense.

⁸Still, expected revenue increases when information is revealed. Above and beyond shifting willingness to pay, which as the above discussion suggests is a wash, affiliation causes updated valuations to move closer, creating more competition and less information rents.

LP2 For vehicles with high (low) CGs revenues should be higher (lower) than for vehicles who do not have reported SCRs.

It may be, however, that the environment of our study is one where the bidders who appear at any given vehicle auction depend on the results of the SCR, in which case the assumption of a fixed set of bidders is violated. In discussion with industry participants, it seems that used car dealers specialize in the type of vehicles they sell. For example, a dealer in a neighbourhood with low income may specialize in older vehicles with worse CGs, while a dealer in a neighborhood with higher income residents specializes in newer cars with higher CGs.

If this is the case then the information revealed by the SCRs may increase expected revenues for a different reason than that behind the linkage principle. Namely, the information will reveal to buyers the CGs of cars that are selling in different lane/run numbers, and will help buyers focus on the cars that they are interested in. This in turn implies that for any given CG, there will be better matching of buyers to vehicles, and in turn, the vehicles will be more likely to sell at the market price perceived by the seller (which will of course determine the reservation price.) This simple “market matching” theory offers a straightforward prediction:

MM1 Expected revenues should (weakly) increase for any given CG level

3 Data

3.1 Experimental Design

The purpose of the experiment was to measure the treatment effect of SCRs on expected auction revenue, probability of sale, and auction price for cars that were consigned to the auction by used car dealers. Our basic approach was as following: We inspected a subset of all dealer-consigned cars at one auction location over the course of 19 weeks using the SCR inspection procedure. Inspected cars were randomly assigned to one of two conditions. In the treatment condition, the SCR of an inspected car was made available to buyers (and sellers). In the control condition, the SCR was withheld from buyer and sellers; only the researchers knew that these cars had been inspected.

Due to a limited number of certified vehicle inspectors we did not inspect all dealer-consigned cars during the 19 weeks. Specifically, out of approximately 1500 dealer consigned vehicles each week, we inspected between 150 and 600 cars per week, depending on the number of inspectors who were available to us (see Table 11). In total, we inspected 8096 cars, 3977 of which were in the control group (SCR not reported) and 4119 were in the treatment group (SCR reported).

The number of inspected cars depended on the number of available inspectors during that week (between 3 and 12). For an auction that was conducted on Wednesday of a given week, we designating for inspections all cars that were checked in starting Friday morning of the prior week. Since we knew the number of available inspectors for that week, we could estimate how many

cars they could inspect by Tuesday (the day before the auction). We kept designating all cars that were checked in until we reached the number of cars that we estimated could be inspected in time for the auction. Once we reached that number we no longer designated checked-in cars for inspection. On days with many inspectors, we inspected all cars that were checked in until mid-day Tuesday. On days with few inspectors we stopped assigning cars to be inspected at some time on Monday.

We assigned cars to treatment and control groups during the check-in process. Cars whose VIN (Vehicle Identification Number) ended in an even digit were assigned to the treatment group. Cars whose VIN ended in an odd digit were assigned to the control group. While the first digits of a VIN number designate manufacturers, country of origin, make, model, model-year, as well as some trim-level information, the later digits are assigned sequentially as vehicles are produced. Hence, the last digit of the VIN is a good randomization device: Whether the digit is even or odd is unrelated to the type of car sold and to the condition of the vehicle. Also, even and odd digits are equally represented in the population of produced cars. We thus expected an approximately even split between treatment and control groups. Consistent with this, the randomization procedure assigned 49.1% of cars to the control group and 50.9% to the treatment group.⁹

Our experiment covers two periods: Weeks 21-30 and weeks 31-39 of 2008. These periods differ in how buyers were made aware of SCRs. During weeks 21-30 the wide availability of SCRs was not explicitly publicized. As discussed in the previous section, SCRs are only available online, but not on the vehicles as they run through auction lanes. Hence, during the first half of the experiment a dealer would only learn about the availability of SCRs if that dealer used the auction house's website to preview cars that would be offered for sale on auction day. A dealer who learned about available cars only on-site on the day of the auction would not know that some of the cars had SCRs. Moreover, if dealers who logged on the day before to see which vehicles are available for sale did not have a habit of searching for SCRs (since these rarely existed for dealer-consigned cars) then they too would not be aware of the SCRs.

As we analyzed auction outcomes after the first eight weeks of the experiment we found little evidence that cars with SCRs were more likely to sell or sold at higher prices (we will present these results in detail later). This could mean that the information contained in SCRs had no effect, however, it could also mean that dealers did not know that SCRs were now available for a significant number of dealer-consigned cars. Hence, starting in week 31 we sent an e-mail to all buyers informing them that they could find SCRs for some of the dealer-consigned cars on the auction house's website prior to the auction day. These e-mails were sent once a week until the end of the experiment.

⁹We cannot reject the hypothesis that our randomization procedure assigned an equal proportion of cars to treatment and control groups (at a 5% significance level).

3.2 Auction and inspection data

The experiment yielded data on 8096 dealer-consigned and inspected cars, 3977 of which were in the control group and 4119 were in the treatment group. For each consigned car we have detailed information on the car, the outcome of the inspection (the SCR), the outcome of the auction, and some data about the auction participants.

Specifically, we observe the specific car that was consigned at the level of a model – model-year – body type – engine – trim level (e.g. a Honda Accord, 1999, 4-door, V6, EX trim). We also observe the mileage of the car. More detailed information about the condition of the car comes from the SCR as described in section 2.2. We use two key measures. The first measure is the CG, a number between 1 (rough) and 5 (clean). The second measure is the estimated cost to fix the damage detailed in the SCR. This includes the auction house’s estimates of both part and labor costs and is reported in dollars.

We observe a unique seller ID that allows us to identify whether different cars were consigned by the same seller. The data reports whether a car was sold during the auction. If the car was sold, we observe the auction price and a unique buyer ID that allows us to identify whether different cars were purchased by the same buyer. Finally, we know the average auction price for cars of the same car type that sold at any of auction house’s locations nationwide during the prior week (henceforth “National Auction Price” or NAP). This allows us to construct a useful normalization of price that is independent of the type of car. Summary statistics are reported in Table 12.

3.3 Randomization check

We compare the treatment and control groups on a variety of observable characteristics. Specifically, if the randomization worked as intended, the distribution of condition grades, repair costs, mileage, vehicle age (model year), and national auction prices in the prior week should be comparable across control and treatment groups. We use a Kolmogorov-Smirnov test for equality of distribution functions. The results are reported in Table 1.

Table 1: Kolmogorov-Smirnov test for equality of distribution functions

Variable	D	p-value
Condition grade	0.0142	0.81
Repair costs	0.0302	0.05
Mileage	0.0170	0.60
Model Year	0.0165	0.64
National Auction Price	0.0183	0.94

For all five measure we fail to reject the hypothesis that the distribution functions are the same. However, the test statistic for repair costs is just above the critical level, indicating that

repair cost may have a different distribution between control and treatment groups. To investigate this further, we compare the means of repair costs across the two conditions. Repair costs for the control group are on average \$1382, for the treatment group the cost are \$1316. We will take account of this \$66 difference when interpreting our auction price results.

4 Results

The results are organized into three parts. First, we report the basic findings of our experiment. Second, we report some supporting evidence for our interpretation of the results by looking at the behavior of online bidders. Finally, we show how our results vary by condition grade, which can help distinguish between the simple market matching story and the linkage principle.


4.1 Basic findings

We begin this section by reporting expected auction revenues for cars in the two experimental conditions. Next we decompose revenues into the probability of sale and auction prices. We then address two alternative explanations for our results.

4.1.1 Expected revenues

The weekly expected auction revenue during our experiment is shown in Figure 1. The figure shows the expected revenue per dealer-consigned car, which varies around a mean of \$3700 (the probability of sale for the entire sample is 43% and the average auction price for a sold car is \$8600).

Figure 1: Expected auction revenue by experimental condition and week



The figure distinguishes between weeks 21 to 30 (during which we did not publicize the existence of SCRs), and weeks 31-39 (during which we sent weekly e-mails mentioning that SCRs for dealer-consigned cars could be found online). We have also highlighted two shocks that

occurred in weeks 37 and 38: First, the auction was affected by the landfall of Hurricane Ike which disrupted Texas, the key market for the auction location where we did our experiment. Second, the auction was affected by the collapse of Lehman Brothers, which decreased auto sales nationally. While these events affected sales at our auction location, we do not have a strong a-priori reason to believe that cars in our two experimental conditions would have been differentially affected. For this reason we treat weeks 31 to 39 as one sample.

Our first findings are summarized in Table 2. Before we publicized the existence of SCRs (weeks 21-30), the expected revenue per consigned car was \$3677 for cars without a posted SCR and \$3690 for cars with a posted SCR. The difference of \$12.8 is not statistically significant at any conventional level (using a t-test). Our results changed once we started sending weekly e-mails mentioning that SCRs for dealer-consigned cars could be found online (weeks 31-39). During this period, the expected revenue per consigned car was \$3262 for cars without a posted SCR and \$3888 for cars with a posted SCR. The expected revenue difference of \$626.1 (or 19.2%) is highly statistically significant (p-value < 0.01). This suggests that SCRs increased expected auction revenues once buyers became aware that SCRs could be found online.

Table 2: Auction revenue by experimental condition for weeks 21-30 and weeks 31-39

	No posted SCR	Posted SCR	Difference	% Difference	t-statistic	p-value
Weeks 21-30	\$3,677.3 2,602 cars	\$3,690.1 2,798 cars	\$12.8	0.003%	0.082	0.94
Weeks 31-39	\$3,261.8 1,375 cars	\$3,887.8 1,321 cars	\$626.1	19.2%	2.89	0.004

4.1.2 Probability of Sale and Prices

We can decompose the difference in expected revenues between cars with and without posted SCRs into differences in the probability of sale and differences in auction prices. Table 3 shows that during weeks 21-30, cars with and without a posted SCR were equally likely to sell; approximately 43% of cars sold in either condition. During weeks 31-39, cars with a posted SCR were 6.3 percentage points (or 16%) more likely to sell than cars without a posted SCR. This difference is highly statistically significantly different from 0 (using a test of proportions with p-value < 0.01).

Prices in the two experimental conditions were not significantly different, in either period. Table 4 shows these results.

One problem in concluding that transaction prices did not differ between experimental conditions is that the variance of prices of sold cars is very high. This is because the auction location sells everything from 11 year old small cars to current model year luxury cars. Ideally, we would like to specify prices relative to the typical price for cars of the same car type, i.e. of the same make, model, and model-year. Fortunately, we know the average auction price for cars of the

Table 3: Sales probability by experimental condition

	No posted SCR	Posted SCR	Difference	% Difference	z-statistic	p-value
Weeks 21-30	43.1% 2,602 cars	43.6% 2,798 cars	0.6%	0.01%	0.41	0.68
Weeks 31-39	39.2% 1,375 cars	45.5% 1,321 cars	6.3%	16.1%	3.31	0.001

Table 4: Transaction prices by experimental condition

	No posted SCR	Posted SCR	Difference	% Difference	t-statistic	p-value
Weeks 21-30	\$8651.4 1,106 cars	\$8582.7 1,203 cars	-\$68.7	-0.01%	-0.279	0.78
Weeks 31-39	\$8446.0 531 cars	\$8704.7 590 cars	\$258.7	3.1%	0.75	0.45

same car type that sold at any of the auction house’s locations during the prior week, what we refer to as the National Auction Price (NAP). We use this measure to construct a normalized price for each car in the sample, specifically, the price of the car divided by the NAP. This normalized price allows us to reevaluate whether there are price differences between experimental conditions. Table 5 shows these results.

Table 5: Transaction prices / NAP by experimental condition

	No posted SCR	Posted SCR	Difference	% Difference	t-statistic	p-value
Weeks 21-30	1.064 1,106 cars	1.059 1,203 cars	-0.005	-0.5%	-0.5	0.62
Weeks 31-39	1.035 531 cars	1.055 590 cars	0.02	1.9%	1.61	0.11

The findings suggest that after week 31, prices were higher by 1.9% for cars with a posted SCR relative to cars without a posted SCR. The difference, however is only marginally significant (p-value 0.11).

Overall, the decomposition suggests that most of the effect of SCRs on expected auction revenue come from an increased probability of sale; transaction prices did increase, but only by little.

4.1.3 Alternative explanations and robustness

To conclude that the increase in auction revenue due to SCRs is attributable to the information revealed in the reports we need to rule out an alternative explanation for why SCRs increased

revenues. We also want to revisit our randomization procedure by checking whether our findings are robust to the inclusion of fixed effects.

Substitution

Suppose that the e-mails that buyers received from week 31 onwards focused buyers' attention on cars with SCR without affecting their willingness to pay for cars with posted SCRs relative to cars without posted SCRs. This might have led to an increase in the number of bidders for cars with posted SCRs and a decrease in the number of bidders for cars without posted SCRs. The larger number of buyers for cars with posted SCRs could have made it more likely that reserve prices were met and this would have increased the probability of sales. Specifically, consider the sales percentages in Table 3. The probability of sale was 43% in weeks 21-30 for both conditions. In weeks 31-39 the *average* probability of sale remained at 43% but cars without a posted SCR sold 39% of the time while cars with a posted SCR sold 45.5% of the time. Could it be that the SCR simply made buyers substitute from cars without posted SCRs to cars with posted SCRs without changing their willingness to pay? Or could it be that SCRs *did* increase buyers' willingness to pay *relative to* cars without a posted SCR but that overall demand for cars at the auction fell from weeks 21-30 to weeks 31-39?

To answer these questions we need an estimate of the secular trend in the probability of sale for our sample period, i.e. we need to know what the sales of cars without a posted SCR would have been if demand for cars at auction had stayed the same from weeks 21-30 to weeks 31-39. This is particularly important because during weeks 21-30 the stock market was declining steadily (the DOW dropped by about 15%) and during week 38 Lehman Brothers crashed. Arguably, market demand may have reacted to these events.

To find such a trend we use cars that were not part of our experiment, namely the cars offered for sale by fleet-sellers. For these cars there was no change in available information due to our experiment. In using fleet-seller consigned cars to establish a secular trend for the probability of sale of dealer consigned cars, we assume that their demand conditions are affected similarly to the demand for dealer consigned cars. This assumption is not unreasonable: While fleet-seller consigned cars are on average somewhat newer, the overlap in age and condition between fleet-seller and dealer consigned cars is high.

{XXX Add evidence of the overlap of these two groups of cars. }

{XXX try to check for a trend in the probability of sale before week 21 if we can get the data.}

The probability of sale for fleet-seller consigned cars is 66.94% in weeks 21-30 (14,161 cars) and 59.75% in weeks 31-39 (13,332 cars). This means the sales probability for fleet-seller consigned cars decreased by 7 percentage points, suggesting that demand for cars at auction decreased from week 21-30 to 31-39. Adding fleet-seller consigned cars to our sample allows us to use a difference-in-differences linear probability regression that estimates the change over time in the probability of sale for cars with and without a posted SCR relative to fleet-seller consigned cars.

The results are in Table 6.

Table 6: Linear probability model: diff-in-diff specification

Dependent Variable: Sold	
Dealer-consigned car, no posted SCR	-.24** (.01)
Dealer-consigned car, posted SCR	-.23** (.01)
Week 31-39	-.07** (.0058)
Week 31-39 * Dealer-consigned car, no posted SCR	.031+ (.017)
Week 31-39 * Dealer-consigned car, posted SCR	.089** (.018)
Constant	.67** (.004)
Observations	35589
R-squared	0.034

* significant at 5%; ** significant at 1%; + significant at 10% level. Robust SEs in parentheses.

The constant in this regression is the probability of sale for fleet-seller consigned cars during weeks 21-30. The coefficient on *Week 31-39* is the change in the probability of sale for fleet-seller consigned cars relative to weeks 21-30 and is our estimate of the secular trend. The variables of interest are the interaction between *Week 31-39* and the two dealer consigned car conditions. The coefficient on *Week 31-39 * Dealer-consigned car, no posted SCR* is 0.031 and is not significantly different from 0 at a 5% level. This means that we cannot reject the hypothesis that the probability of sale for dealer consigned cars was unchanged between weeks 21-30 and weeks 31-39. In contrast, the coefficient on *Week 31-39 * Dealer-consigned car, posted SCR* at 0.089 is significantly different from 0 (p-value < 0.01). The interpretation of these results is as follows: Under the maintained assumption that the demand conditions of fleet-seller consigned cars change similarly to the demand conditions for dealer consigned cars, we find no evidence that the e-mails sent out starting in week 31 led dealers to substitute from cars without posted SCRs to cars with posted SCRs. Instead, it seems that the probability of sale for cars without posted SCRs was unchanged (relative to fleet-seller consigned cars) while the probability of sale for cars with posted SCRs increased.

One remaining concern is that there may have been substitution between fleet-seller consigned cars and dealer-consigned cars with a posted SCR. If so, controlling for the secular trend by using

the change in probability of sale of fleet-seller consigned cars would no longer be valid. To address this concern we constructed a sample of buyers who only purchased fleet-seller consigned cars during weeks 21-30. 616 dealers fall in this category, a large fraction of the 1670 dealers who purchased at least one car (fleet-seller or dealer consigned) during our experimental period. If there is substitution between fleet-seller consigned cars and dealer consigned cars with a posted SCR, we should find that these 616 dealer—if they purchased *any* dealer consigned cars during weeks 31-39—should be more likely to buy cars with a posted SCR than without a posted SCR. We find no evidence of such behavior: Dealers who only purchased fleet-seller consigned cars during weeks 21-30 purchased 48 dealer consigned cars with a posted SCR and 53 dealer consigned cars without a posted cars after we started publicizing SCRs by e-mail (i.e. during weeks 31-39).

We conclude that substitution is unlikely to explain why SCRs increase expected auction revenue.

Randomization check

Previously, we compared the treatment and control groups on a variety of observable characteristics to make sure that the randomization worked as intended. A second approach to checking whether our procedure yielded a random assignment to treatment and control groups is to analyze whether our basic results change as we control for a variety of fixed effects. Specifically, we regress auction revenue on the treatment, controlling for seller fixed effects (267), week fixed effects (9), model year fixed effects (12), and condition grade fixed effects (5).

Table 13 shows the regression results. For comparison, we also report the treatment effect without fixed effects (which is also in Table 2). The point estimate of the treatment effect drops from \$626 to \$439. However, we can't reject the hypothesis that the treatment effect is unchanged by the inclusion of the extensive set of fixed effects. The point estimate in the fixed effect specification is somewhat less precisely estimated but is still significant at a 6% level. These results provide no evidence that our procedure yielded a non-random assignment to treatment and control groups.

4.2 Online Transactions

We have argued that SCRs did not increase expected auction revenue during weeks 21-30 because dealers during that period were not aware that SCRs had been posted for many dealer consigned cars. One way to test this argument is to look at the behavior of dealers for whom we know that they must have been aware of SCRs even during weeks 21-30. If these dealers behave no differently before and after week 31, this supports our argument that the effectiveness of SCRs during weeks 31-39 was tied to dealers knowing about them.

To identify a set of dealers who must have been aware of SCRs even during weeks 21-30 we make use of the auction house's online bidding feature. Clearly, dealers who bid online must have

know about SCRs because the SCRs are listed on the page that is used to start online bidding. Furthermore, for these online dealers this is the only source of information that puts them on some equal footing with the on-site lane bidders.

Online bidding was relatively rare at the time of our experiment. Of the 8,096 dealer consigned cars that were up for auction between week 21 and week 39, only 243 (3%) received an online bid. The 8,096 cars up for auction yielded 3,482 sales. Of these sold cars, only 137 (3.9%) received the winning bid from an online bidder.

In the following we look at three measures of online behavior as a function of whether we posted a SCR or not. First, what percentage of vehicles received an online bid? Second, for what percentage of sold vehicles was the winning bid placed online? Third, what is the expected number of online bidders? We will compare all three measures for weeks 21-30 and 31-39.

Table 7 shows the percentage of vehicles that received an online bid by week and by whether a SCR was posted. Over the entire experimental period, 3.45% of cars with a posted SCR received an online bid, compared to 2.54% of cars without a posted SCR. This 36% difference in the probability of receiving a bid is statistically significant (using a test of proportions and a 5% significance level). The key comparison is whether a similar difference already existed in weeks 21-30 or whether it was mostly driven by dealer behavior in weeks 31-39. We find that a posted SCR increased the probability of receiving an online bid by 30% during weeks 21-30. This suggests that an SCR had a meaningful effect on dealer behavior during weeks 21-30 for dealers who knew about its existence.

Table 7: Percentage of dealer-consigned cars which received an online bid

	No posted SCR	Posted SCR	Difference	% Difference	z-statistic	p-value
All weeks	2.54% 3,977 cars	3.45 % 4,119 cars	0.91%	35.8%	2.39	0.017
Weeks 21-30	2.69% 2,602 cars	3.50% 2,798 cars	0.81%	30.2%	1.72	0.086
Weeks 31-39	2.25% 1,375 cars	3.33% 1,321 cars	1.08%	47.7%	1.70	0.089

We find a similar result in Table 8, which shows the percentage of sold vehicles for which the winning bid was placed online. Over all weeks, the winning bids of 4.7% of cars with a posted SCR was places online, compared to 3.07% of cars without a posted SCR. This 53% difference is statistically significant (using a test of proportions and a 5% significance level). As before, much of the SCR effect is already present during weeks 21-30 (although the SCR effect is a bit smaller and statistically weaker than in the overall sample).

Our final online result is in Table 9, which shows the expected number of online bidders per 100 auctions. We find that over all weeks of the experiment, more online bidders participated in auctions for cars with a posted SCR (4.73 per 100 auctions) than for cars without a posted SCR

Table 8: Percentage of sold dealer-consigned car with where winning bid was placed online

	No posted SCR	Posted SCR	Difference	% Difference	z-statistic	p-value
All weeks	3.07% 1,660 cars	4.72 % 1,822 cars	1.65%	53.6%	2.50	0.01
Weeks 21-30	3.21% 1,121 cars	4.5% 1,221 cars	1.29%	40.3%	1.62	0.11
Weeks 31-39	2.78% 539 cars	5.15% 601 cars	2.37%	85.3%	2.03	0.04

(3.65 per 100 auctions). As for the previous two measures, the SCR effect seems to be present already in weeks 21-30 (although the SCR effect is a bit smaller and statistically weaker than in the overall sample).

Table 9: Expected number of online bidders per 100 auctions

	No posted SCR	Posted SCR	Difference	% Difference	t-statistic	p-value
All weeks	3.65 3,977 cars	4.73 4,119 cars	1.08	29.8%	2.21	0.027
Weeks 21-30	3.77 2,602 cars	4.72 2,798 cars	0.95	25.3%	1.57	0.12
Weeks 31-39	3.42 1,375 cars	4.77 1,321 cars	1.35	39.5%	1.60	0.11


Given that online dealers knew about SCRs from the beginning (week 21) of the experiment, and given that the effect of a posted SCR barely changes between weeks 21-30 and 31-39, we conclude that the effectiveness of SCRs we observe offline during weeks 31-39 is most likely to be tied to dealers knowing about SCRs.

4.3 Transactions by Quality Grades

We now investigate whether the effect of a posted SCR on auction outcomes differs by the condition of the vehicle. The basic auction expected revenue results are shown in Figure 2. Not surprisingly, expected revenues generally increase with condition grade. Notice also that a posted SCR does not affect expected auction revenues in weeks 21-30; this is consistent with our earlier findings. The two most important results in the figure are that an SCR has the largest effect for cars high condition grades, i.e. cars that are in very good condition, and that SCRs do not have a negative effect for the very worst grades. In fact, the effect is slightly positive. This suggests that market matching is likely to be an important consequence of the SCRs.


Another way to look at expected auction revenues is to normalize revenues by the National Auction Price of the consigned car. This reduces the variance in the expected auction revenue

Figure 2: Expected auction revenue by condition grade and experimental condition



measure. Figure 2 shows the result. A similar pattern emerges, except that SCRs seem to have a non-negative effect across *all* condition grades, again consistent with a market matching story.

Figure 3: Expected auction revenue / NAP by condition grade and experimental condition



To assess the statistical significance of these results, we use a t-test to compare auction revenues (normalized by NAP) by condition grade during weeks 31-39 (see Table 10). The small cell sizes for each condition grade allow us to conclude only for grades 1 and 5 that a posted SCR is associated with higher expected auction revenues. There is weak evidence that a posted SCR is associated with higher expected auction revenues for grades 2 and 4. The effect for condition grade 3 is clearly too small to consider as different from 0.

We can decompose the auction revenue effect into a sales probability and price effect. As Figures 4 and 5 show, most of the difference in expected auction revenues comes from differences in sales probabilities not from differences in auction prices. This is true regardless of condition grade.

Table 10: Expected auction revenue / NAP by condition grade, weeks 31-39

Condition Grade	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	t-statistic	p-value
1	1070	0.361	0.425	0.064	17.7 %	2.02	0.044
2	483	0.428	0.496	0.068	15.9 %	1.39	0.16
3	644	0.436	0.472	0.035	8.0 %	0.82	0.41
4	254	0.504	0.609	0.104	20.6 %	1.56	0.12
5	245	0.293	0.483	0.191	65.2 %	2.98	0.003

Figure 4: Sales probability by condition grade and experimental condition






Figure 5: Price/NAP by condition grade and experimental condition



SCRs contain an estimate of the labor and parts cost required to fix damage on the inspected vehicle. These yield a different estimate of the condition of a vehicle than the condition grade. For example, as a rule the AH will not award a car a condition grade above 3 if any sheet metal of the car has been repainted. Now suppose that a car has had some parts of its sheet metal repainted but the car has no damage otherwise. Then the repair cost estimate is zero but the condition grade is 3.

Since the measure are not perfectly correlated, but the repair costs give dealers useful information about the condition of the vehicle, we can also investigate whether the effect of a posted SCR on auction outcomes differs by the estimated repair costs of the vehicle. We create repair cost quintiles (to stay with the condition grade ordering we define 1 as high repair cost and 5 as low repair cost). Figure 6 shows the effect of SCR on expected auction (divided by NAP) by repair cost quintile.

Figure 6: Expected auction revenue/NAP by repair cost and experimental condition



Similar to our findings about condition grades, we find that expected auction revenues are generally higher for cars with a posted SCR in weeks 31-39. However, we don't find evidence that the effect is largest for cars with low repair costs (as we did not cars with high condition grades). There is also no evidence that a posted SCR reduces expected auction revenues for cars with high estimated repair costs. In fact, the results point to the opposite.

5 Discussion

- Linkage may be there but market matching seems to overcome the negative effect that low CGs may have
- Our sellers set their reserve price without knowing that their cars will have SCRs. If the sellers respond with changing reserve prices then by revealed preferences it must be that expected revenue is higher still.

6 References

- De Silva, Dakshina G., Timothy Dunne, Anuruddha Kankanamge and Georgia Kosmopoulou (2008) “The impact of public information on bidding in highway procurement auctions,” *European Economic Review* **52(1)**:150–181
- Genesove, David (1993) “Adverse selection in the wholesale used car market,” *Journal of Political Economy* **101(4)**:644-665.
- Goeree, J.K., Offerman, T. (2003) “Competitive bidding in auctions with private and common values,” *Economic Journal* **113 (4)**:598–613.
- Kagel, John .H., R.M. Harstad and Dan Levin (1987) “Information impact and allocation rules in auctions with affiliated private values: A laboratory study.” *Econometrica* **55(6)**:1275–1304.
- Milgrom, Paul R. *Putting Auction Theory to Work*. 2004. Cambridge University Press, New York, NY.
- Milgrom, Paul R. and Robert J. Weber (1982) “A theory of auctions and competitive bidding,” *Econometrica* **50 (5)**:1089–1122.
- Overby, Eric and Sandy Jap (2009) “Electronic and Physical Market Channels: A Multi-Year Investigation in a Market for Products of Uncertain Quality,” *Management Science*, **55(6)**:940-957.
- Perry, Motty and Phillip .J. Reny (1999) “On the failure of the linkage principle in multi-unit auctions,” *Econometrica* **67(4)**:895–900
- Porter, Robert H. (1995) “The role of information in US offshore oil and gas lease auction,” *Econometrica* **63(1)**:1–27.

Appendix

Table 11: Dealer-consigned and inspected cars by week[†]

Sale Week	Dealer-Consigned Total	With SCR	
		Not reported	Reported
21	1,442	237	223
22	1,709	195	186
23	1,438	324	330
24	1,606	281	365
25	1,249	303	344
26	1,408	229	250
27	1,170	290	305
28	1,462	245	245
29	1,440	267	281
30	1,621	231	269
31	1,533	233	247
32	1,590	214	215
33	1,329	237	154
34	1,555	225	185
35	1,526	150	140
36	1,474	73	85
37	1,418	90	107
38	1,554	71	84
39	1,639	82	104
Total	28,163	3,977	4,119

[†]Weeks are of 2008.

Table 12: Summary Statistics

Variable	N	mean	p50	sd	min	max
Model Year	8096	2003.509	2004	2.653825	1997	2009
Mileage	8096	76091.4	71340	45527.27	0	999999
Condition Grade	8096	2.419466	2	1.312998	1	5
Repair Costs	8096	1348.07	1025.12	1236.787	0	16110.8
Sold	8096	.4300889	0	.4951189	0	1
Sales Price	3430	8604.672	7300	5861.903	500	59000
National Auction Price	3430	8395.569	6975	5810.735	200	62000
Sales Price/National Auction Price	3430	1.056347	1.030303	.2430727	.2358974	5.6

Table 13: Randomization check: Expected auction revenue for weeks 31-39

	Base result	Fixed Effects
Posted SCR	626** (217)	439+ (226)
Week 32		-277 (390)
Week 33		90 (406)
Week 34		511 (399)
Week 35		489 (435)
Week 36		318 (530)
Week 37		-272 (499)
Week 38		-1120* (536)
Week 39		-1083* (545)
Condition grade=2		643+ (329)
Condition grade=3		767* (316)
Condition grade=4		2344** (432)
Condition grade=5		1821** (469)
Model year 1998		1338 (1380)
Model year 1999		1162 (1226)
Model year 2000		1120 (1212)
Model year 2001		1540 (1184)
Model year 2002		2179+ (1160)
Model year 2003		1980+ (1150)
Model year 2004		2627* (1152)
Model year 2005		2103+ (1155)
Model year 2006		2867* (1156)
Model year 2007		3623** (1184)
Model year 2008		2682* (1265)
Model year 2009		1035 (2245)
Seller fixed effects	no	yes
Constant	3262** (146)	451 (1120)
Observations	2696	2696
R-squared	0.003	0.227

* significant at 5%; ** significant at 1%; + significant at 10% level. Robust SEs in parentheses.