Information Disclosure as a Matching Mechanism: Theory and Evidence from a Field Experiment*

Steven Tadelis
UC Berkeley

Florian Zettelmeyer
Northwestern University and NBER

May 17, 2011

Abstract

Market outcomes generally depend on the quality of information available to its participants. Using a large scale field experiment that randomly discloses information about quality in wholesale automobile auctions we measure the effect of information disclosure on auction outcomes. As the theoretical literature predicts, more information increases expected revenues. However, the biggest gains in revenue are for the best and worst quality cars, which is at odds with conventional information disclosure theories. We argue that information disclosure causes better matching of heterogeneous buyers to different quality cars, an effect that has been ignored so far. This not only rationalizes the empirical results, but further implications of information as a matching mechanism are also verified in the data. Our findings have implications for the design of auction markets, including procurement auctions, and online consumer auctions.

JEL classifications C93, D44, D82, L15

*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 thank Igal Hendel for helpful discussions, and many seminar participants for comments. Tadelis thanks the National Science Foundation for financial support.

1 Introduction

It is generally accepted that a market's efficiency critically depends on its participants having sufficient information about the nature of the goods and services being traded. The potential hazards a buyer faces when trading in markets with information asymmetries often lead to market imperfections and stifles efficient trade.[1] Indeed, in resale markets, housing markets, labor markets, health care markets and markets for corporate securities, sellers may have better information than buyers about the value of the good or service being traded. Furthermore, sellers may have control over how much information to release, and buyers may choose how much information to acquire.

This paper studies the effects of information disclosure on market outcomes. In particular, we investigate the wholesale market for used automobiles where trade between car dealers is facilitated through auctions. In these markets, sellers will typically have more information about the condition of the used vehicle than buyers do, and sellers can often control the amount of information that they choose to release to potential buyers. Using a randomized field experiment, we are able to precisely document how more information affects auction outcomes. We are also able to quantify the changes in consummated trades, in expected revenues, and how these differ across quality levels of the cars that are sold. The results show that ex ante information does indeed affect market outcomes, but in ways that are inconsistent with the standard theoretical auction literature. Rather than reducing information rents to buyers, we suggest that more information plays an important role in matching buyers with goods.

Studies of auction design focus more on the auction's rules (open or sealed, first or second price, free entry or invited bidders, etc.) and less on how much information a seller should release or bidders should acquire. The most notable exception is the celebrated "linkage principle" identified by Milgrom and Weber (1982). They show that under sensible conditions, if the seller releases more information then her expected revenue from an auction will increase, and two policy conclusions emerge. First, the seller's benefits from committing to release as much credible information as she can. Second, auction formats that cause more information disclosure (e.g., open auctions) will generate higher expected revenues compared to auctions that do not reveal information (e.g., closed auctions.)

The intuition of the linkage principle is subtle because information disclosure can either increase or decrease a buyer's willingness to pay, and with correct expectations this should imply a "wash" for each individual bidder. If the information disclosed in a given auction is "bad news" relative to initial expectations then it will cause valuations to drop, just as "good news" will cause them to increase. As a result, relative to the scenario without information disclosure, bids must be lower following bad news and higher following good news. However, Milgrom and Weber (1982) show that when the valuations of the bidders are affiliated[2] then information disclosure

[1] It is also well known that in some cases, more information can hamper the efficiency of markets. See, e.g., the seminal work of Hirshleifer (1971).

[2] explain affiliation

causes bidders to have more aligned views of the object's value. This in turn causes them to bid more competitively, resulting in higher expected revenues for the seller (reduced "information rents" for the bidders). In related work, Ottaviani and Prat (2001) explore the incentives of a monopolist to reveal information, and describe market conditions where a force similar to the linkage principle occurs.

We test the effect of information disclosure on auction markets by implementing a unique randomized field experiment. We manipulate information disclosure in a market where thousands of vehicles are put up for sale each week with an average vehicle value of $8,500, while keeping all other aspects of the auction fixed. Furthermore, the information disclosed has a clear ranking of quality, which in turn allows us not only to test whether average revenues change, but also measure how the average revenues change for *each* quality rank. This level of detail enable us to show that the more refined predictions of standard information disclosure theory regarding the effect of good and bad news are violated by patterns in the data.

Our empirical results strongly support the hypothesis that more information increases average revenues obtained by sellers, consistent with the linkage principle. More interestingly, however, the increase in expected revenues holds true across all quality levels. In fact, the strongest positive effect is for the *very best* and *very worse* quality levels. Further analysis shows that this is more pronounced when considering whether the information revealed is good news or bad news relative to initial expectations. If, given the observable characteristics of the vehicle, the information revealed is consistent with expectations then information disclosure has no effect on auction outcomes. However, if the information revealed is either *better or worse* than expected, then there is a positive effect on expected revenues. These results are not consistent with standard information disclosure theories as described above.

This surprising observation guided the construction of our theoretical contribution, which to the best of our knowledge has not been explored in the literature. We argue that in addition to increasing competition *within* a given auction or market, ex ante information disclosure increases competition *across* auctions or markets. We describe a situation where goods of different quality levels are randomly offered for sale in different auctions, and heterogeneous potential buyers need to choose in which auction to participate. Though higher quality is more valued by all buyers, the type of bidder who values the good most depends on the quality of the good. Thus, though goods are vertically differentiated, buyer heterogeneity imposes horizontal differentiation. It follows that ex ante information disclosure helps buyers choose which auction or market to participate in, and buyers will use the information to match with the goods for which they have a high value relative to other buyers. This in turn intensifies the effective competition in any given auction or market by increasing the number of relevant high-value bidders. As a consequence, both the number of efficient transactions and the expected revenue for sellers will increase. As we discuss later in section 5, the idea of horizontal differentiation for vertically differentiated goods is explored in Board (2009) who shows that information disclosure can changes the type of bidder who wins the auction, which he labels the *allocation effect*.

By uncovering the matching effect of information disclosure we contribute to the theoretical literature described earlier. Moreover, the simple matching model we construct in section 5 explains our initial set of empirical findings. Namely, when the information disclosed coincides with expectations given observables, the information should not affect the composition of bidders who bid on the vehicle, and as a consequence, the outcomes are the same as they are with no information disclosure. However, when the information disclosed is either a positive or negative surprise relative to expectations, this will attract bidders who are relatively strong given the information disclosed, and as a result the seller benefits. Of course, when aggregating across quality, all the positive effects aggregate to higher expected revenues as we initially observe in the data. We conclude the analysis with a series of tests that both confirm the assumptions of our simple model as well as additional predictions derived from it.

With this paper we also contribute to the growing empirical literature on the effects of information disclosure on market outcomes in general,[3] and on auctions in particular. Due to the challenge of testing how variation in information disclosure affects auctions in the field, there have been few such studies. De Silva et al. (2008) exploit a policy change in the laws of the state of Oklahoma that led to the release of internal estimates of the costs to complete highway construction projects. Using a difference-in-difference approach they show that average bids fell after the change in policy, consistent with the prediction of the linkage principle (since this is a "reverse" auction, a drop in cost-bids is like an increase in revenue.) Cho, Paarsch, and Rust (2010) used a field experiment to vary auction formats and shown that, consistent with the linkage principle, the expected revenues of an open-outcry, English auction are higher than those of auction formats which reveal less information. They do not, however, exogenously vary the amount of information that is disclosed to sellers. There is also a body of work, including Kagel and Levin (1986), Kagel et al. (1987) and Levin et al. (1996), which implements laboratory experiments that directly and indirectly test the linkage principle. By manipulating the information that bidders receive, or the auction formats (open versus closed), they show that more information disclosure results in higher average revenues.[4]

The paper proceeds as follows: Section 2 describes the industry, the details of wholesale automobile actions, and the information provided to bidders. Section 3 describes the data and the experimental design while Section 4 presents the experimental results. Section 5 discusses the existing theoretical implications of information revelation in auctions and offers a simple model of how information revelation can better match goods to buyers, with results that are consistent with the data. New implications are also offered and tested with affirmative results. Section 6 concludes.

[3]See, for example, Porter (1995), Jin and Leslie (2003), Cutler, Huckman, and Landrum (2004), Jin (2005), Busse, Silva-Risso, and Zettelmeyer (2006), and Lewis (2010).

[4]For some qualifications, see Goeree and Offerman (2003).

2 Wholesale Auto Auctions

The U.S. retail market for used cars is sizeable. Estimates place used car sales at more than 35 million cars in 2009, most of which were sold by franchise or independent dealers.[5] 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, 35% of all used vehicles sold by new car dealers in 2008 were sourced in such auctions (see NADA DATA, 2009)[6]. Most auctions are administered by a few prominent auction houses that specialize in this market, one of which provided the data for this study.

2.1 The Auction Process

The buyers in our auction are exclusively dealers, while the sellers mainly belong to one of three categories: dealers who wish to sell used cars from their inventory; owners of large fleets such as rental car agencies who periodically turn over their inventory; and financial lease agencies who sell vehicles for which a lease contract had ended. Sellers bring their vehicles to the auction site one or more days in advance of the actual auction. Each registered vehicle is 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.[7] Several thousands of vehicles may be auctioned off during a sale day.

Before the auction day begins, potential bidders receive a list of vehicles that will be auctioned, including the lane and run numbers, as well as basic information about the vehicle such as make, model, model year, color, and mileage. 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 beginning 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 driven to the auction block where it stops, amid several potential buyers, and is left idling as the auctioneer begins the auction.[8] The auction is an ascending oral (English) auction that lasts for about 30 seconds.[9] The auction ends when no

[5] See the National Independent Automotive Dealer's Association (NIADA) website (http://www.niada.com/) for their 2010 annual report. Sales in 2008 and 2009 were similar, down from more than 42 million vehicles sold in 2006 due to the economic downturn.

[6] This is available at http://www.nada.org/Publications/NADADATA.

[7] For example, a vehicle with a lane-run number of 9-132 will be auctioned in lane 9, and will be the 132$^{\text{nd}}$ vehicle in the lane.

[8] Some cars that are not in driving condition are towed.

[9] Interestingly, 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

bidder is willing to raise the price, and if the price exceeds the seller's reserve price then the sale is consummated. Otherwise, the vehicle either returns to the seller's lot or is left at the auction site 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 seldom 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 can 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 has developed a Standard Condition Report (SCR) designed to offer a standard set of inspections, and a standard way in which to present the information. The SCR is based on a detailed inspection that takes about 25 minutes per car. The inspections cover the vehicle's exterior condition, 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, except that the inspecting technician documents unusual engine sounds. The technician enters all of the

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.

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.[10] The grading system is from 1 through 5, with increments of 0.1, where $CG = 1$ is considered "rough", and $CG = 5$ 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.

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. The basic approach was as follows: A subset of all dealer-consigned cars were inspected 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 auction house knew that these cars had been inspected.

Due to a limited number of certified vehicle inspectors not all dealer-consigned cars were inspected during the 19 week period. The number of inspected cars depended on the number of available inspectors during that week (between 3 and 12). Specifically, out of approximately 1500 dealer consigned vehicles that were registered each week, between 150 and 600 cars were inspected per week (see Table 17). In total, 8098 cars were inspected, 3980 of which were in the control group (SCR not reported) and 4118 were in the treatment group (SCR reported).

For an auction that was conducted on Wednesday of a given week, all cars that were checked in starting Friday morning of the prior week were candidates for inspection. Since the number of available inspectors was known for any given week, it was possible to estimate how many cars they could inspect by Tuesday (the day before the auction). All cars that were checked were designated for inspection until the number of cars that we estimated could be inspected in time for the auction. Once that number was reached, no more checked-in cars were inspected. On days with many inspectors, all cars that were checked in until mid-day Tuesday were inspected,

[10]For other papers that investigate the role of condition reports, see Genesove (1993) and Overby and Jap (2009).

whereas on days with few inspectors inspections were performed on cars that were checked in until some time on Monday.

Cars were assigned 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 while cars whose VIN ended in an odd digit were assigned to the control group. The first digits of a VIN number designate manufacturers, country of origin, make, model, model-year, as well as some trim-level information, whereas 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 49.15% of cars to the control group and 50.85% to the treatment group.[11]

Our experiment covers two periods: Weeks 21-30 (5,402 cars) and weeks 31-39 (2696 cars) 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 weak evidence that cars with SCRs were more likely to sell or sold at higher prices (these results are described in Section 4). This could mean that the information contained in SCRs had little effect. It could also mean, however, that dealers did not know that SCRs were made available for a significant number of dealer-consigned cars. Hence, starting in week 31 an email was sent to all registered 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 emails were sent once a week until the end of the experiment.[12]

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

[12] The emails stated that the company is ramping up its capabilities to offer SCRs, and as such technicians were assigned to inspect a subset of vehicles that were chosen randomly based on the availability of inspection technicians. It was made clear that these were not solicited or affected by the sellers.

3.2 Auction and inspection data

The experiment yielded data on 8098 dealer-consigned and inspected cars, 3980 of which were in the control group and 4118 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 limited data about the auction participants.

Specifically, we observe the car that was consigned at the level of a model; model-year; body type; engine and trim level (e.g. a Honda Accord, 1999, 4-door, V6, EX trim) as well as 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 the 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 18.

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	0.0137	0.83
Repair costs	0.0301	0.05
Mileage	0.0172	0.58
Model Year	0.0167	0.61
National Auction Price	0.0246	0.17

For four of the five measure we fail to reject the hypothesis that the distribution functions are the same. However, the test statistic for repair costs is just at 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 (less than 5%) difference when interpreting our auction price results.

4 Results

The results are organized into three parts. First, in section 4.1 we report the aggregate findings of our experiment and show that more information increases the likelihood that cars sell, and that, conditional on selling, they sell for a slightly higher price. Second, in section 4.2 we show how the results vary by condition grade, and by whether the condition grade is better or worse than expected. We then argue that the theory behind the linkage principle is inconsistent with these results Finally, in section 4.3 we report evidence for our interpretation of the results and show that the results are robust to key alternative explanations.

4.1 Aggregate Findings

The expected revenues for any given seller is comprised of the probability that a vehicle will sell (the reserve is met), and the price obtained conditional on a sale. As such, we consider each of these two components separately.

Table 2 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. This suggests either that SCRs had no effect or that buyers were unaware of SCRs.

During weeks 31-39, when the availability of SCRs was announced with a weekly e-mail, 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).

Table 2: Sales probability by experimental condition

	No posted SCR	Posted SCR	Difference	% Difference	z-statistic	p-value
Weeks 21-30	43.0% 2,605 cars	43.6% 2,797 cars	0.6%	1.39%	0.43	0.66
Weeks 31-39	39.2% 1,375 cars	45.5% 1,321 cars	6.3%	16.1%	3.31	0.001

Prices in the two experimental conditions were not significantly different, in either period. Table 3 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

Table 3: Transaction prices by experimental condition

	No posted SCR	Posted SCR	Difference	% Difference	t-statistic	p-value
Weeks 21-30	\$8,742.9 1,121 cars	\$8,616.9 1,220 cars	-\$126.0	-1.4%	-0.51	0.61
Weeks 31-39	\$8,502.2 539 cars	\$8,738.9 601 cars	\$236.7	2.7%	0.68	0.50

sells everything from 11 year old small cars (e.g., Honda Civic) to current model year luxury cars (e.g., BMW 740). 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. To do this we use the average auction price for cars of the same car type that sold at any of the auction house's locations during the prior week, what we refer to earlier 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 4 shows these results.

Table 4: 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	1.058 1,202 cars	-0.006	-0.5%	-0.56	0.58
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 of 0.11).

In summary, an analysis of the probability of sale and prices conditional on sale suggests that most of the effect of SCRs on expected auction revenues comes from an increased probability of sale; transaction prices did increase, but only by a little.

4.2 Decomposing the Effects

We now investigate whether the effect of a posted SCR on auction outcomes differs by the condition grade of the vehicle, or by the degree of "surprise" of the information relative to what would be expected given observables. As before, we decompose the auction revenue effect into a sales probability and price effect.

4.2.1 Transactions by Quality Grades

Similar to our aggregate results, when we analyze the effect of posted SCRs by condition grade, most of the difference in expected auction revenues comes from differences in sales probabilities rather than from differences in auction prices. To assess the statistical significance of the sales probability results by condition grade, we restrict ourselves to weeks 31-39 and use a test of proportions to assess significance (see Table 5). For grades 1, 4, and 5 we conclude that a posted SCR is associated with a higher sales probability. There is weak evidence that a posted SCR is associated with a higher probability of sale for grade 2. The effect for condition grade 3 is clearly too small to be considered different from 0.

Table 5: Sales probability by condition grade, weeks 31-39

Condition Grade	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	z-statistic	p-value
1	1070	0.365	0.419	0.054	14.8%	1.81	0.07
2	483	0.404	0.472	0.068	16.9%	1.51	0.13
3	644	0.428	0.439	0.011	2.6%	0.29	0.77
4	254	0.478	0.593	0.115	24.1%	1.84	0.07
5	245	0.291	0.481	0.191	65.5%	3.03	0.002

We can also assess the statistical significance of prices by condition grade. We use a *t*-test to compare the prices by condition grade during weeks 31-39 (see Table 6). We can conclude only for grades 3 and 4 that a posted SCR is associated with a significantly different auction price, where it is 6.1% higher for grade 3 and 4.3% lower for grade 4.

Table 6: Price/NAP by condition grade, weeks 31-39

Condition Grade	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	t-statistic	p-value
1	409	1.010	1.039	0.029	2.9%	1.26	0.21
2	209	1.059	1.071	0.011	1.1%	0.45	0.65
3	274	1.035	1.098	0.063	6.1%	2.25	0.03
4	134	1.072	1.026	-0.046	-4.3%	2.07	0.04
5	95	1.038	1.020	-0.018	-1.7%	0.62	0.54

Clearly, the fact that the lowest possible grade of 1 has a *positive effect* on sales despite an unchanged average sales price is surprising. The linkage principle states that more information increases expected revenues *unconditional* on the actual quality level. However, *conditional on bad news*, typical disclosure models predict that bidders *lower* their willingness to pay, making this finding puzzling.

Next, we investigate not just whether expected revenues increase for condition grades that are low but whether expected revenues increase for condition grades that are low *relative to expectations.*

4.2.2 Transactions by Informational Content

Information disclosure theory suggests that whether SCRs actually affect buyers' expectations about the condition of a vehicle depends on what they expect without the SCR. Bidders already have some information (regardless of whether an SCR is posted) that is predictive of condition grade, namely mileage and age. This information allows buyers to estimate the condition grade as a function of age and mileage. As one can see in Table 7, the average condition grade varies substantially by vehicle age and by vehicle mileage as one would predict it to: cars that are older or that have higher mileage will, on average, have worse CGs.

Table 7: Average condition grade (CG) by mileage category and vehicle age

Mileage Category	Average CG	Vehicle Age	Average CG
0-20,000	4	1	4.2
20,001-40,000	3.6	2	3.9
40,001-60,000	3.1	3	3.3
60,001-80,000	2.7	4	3.1
80,001-100,000	2.5	5	2.9
100,001-120,000	2.3	6	2.5
120,001-140,000	2	7	2.2
140,001-160,000	1.9	8	2.1
160,001-180,000	1.6	9	2
180,001-200,000	1.3	10	1.9
>200,001	1.4	11	1.8
		12	1.7

As a result, we would like to perform an empirical test that explicitly allows for condition grade expectations that differ with vehicle age and mileage. We proceed as follows. We first estimate the predicted condition grade of each car in our sample based on the vehicle age and vehicle mileage. We make this prediction by regressing condition grade on vehicle age year dummies, a third-order polynomial of vehicle mileage, and vehicle mileage deciles. We take the difference between the actual condition grade and the predicted condition grade from this regression to construct a distance measure from the expected condition grade. Finally, we split this distance measure into terciles, where the bottom tercile contains cars with worse than expected condition grades, the middle tercile contains cars with close to expected condition grades, and the top tercile contains cars with better than expected condition grades.

As Table 8 shows, during weeks 31-39 there is no statistically significant effect of a posted SCR on the probability of sale for cars for the middle tercile where actual CGs are close to

Table 8: Sales probability by difference of expected condition grade (CG), weeks 31-39

Tercile of Difference from Expected CG	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	z-statistic	p-value
Worse than expected	901	0.338	0.403	0.065	19.2%	2.00	0.045
Close to expected	897	0.416	0.425	0.009	2.2%	0.28	0.78
Better than expected	898	0.421	0.53	0.122	29.0%	3.27	0.001

expected CGs. However, in both terciles where CGs have informational content the effect on the probability of sale is positive and significant.[13]

We replicate this analysis for prices. As Tables 9 shows, the results mirror what we found in Table 6: there is no statistically significant effect of a posted SCR on the prices for cars in any of the terciles during weeks 31-39.

Table 9: Price/NAP by difference of expected condition grade (CG), weeks 31-39

Tercile of Difference from Expected CG	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	t-statistic	p-value
Worse than expected	327	0.982	1	0.018	1.28%	0.87	0.39
Close to expected	371	1.04	1.08	0.03	2.9%	1.49	0.14
Better than expected	423	1.07	1.08	0.01	0.9%	0.43	0.66

As before, and more convincingly, the findings fly in the face of predictions made by standard information disclosure models. We observe that both good news and bad news causes bidders to bid more aggressively.

4.3 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 alternative explanations 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.

4.3.1 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 behaved no

[13]Note that vehicles in the top tercile sell much better even without the SCRs being reported. This is consistent with the fact that buyers have the opportunity to cruise the lot before the auction begins, and thus identify characteristics of the vehicle that are informative, but for which we cannot control in our prediction regression.

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, this is the only source of information that puts online dealers on some equal footing with the on-site lane bidders.

Online bidding was relatively rare at the time of the experiment. Of the 8,098 dealer consigned cars that were up for auction between week 21 and week 39, only 243 (3%) received an online bid. The 8,098 cars up for auction yielded 3,481 sales. Of these sold cars, only 137 (3.9%) received the winning bid from an online bidder.

We consider three measures of online behavior as a function of whether an SCR was posted 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 10 shows the percentage of vehicles that received an online bid by experimental periods 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, p-value 0.02). 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 10: 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% 3,980 cars	3.45 % 4,118 cars	0.91%	35.8%	2.40	0.016
Weeks 21-30	2.69% 2,605 cars	3.50% 2,797 cars	0.81%	30.2%	1.73	0.084
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 11, 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).

Table 11: Percentage of sold dealer-consigned car 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 % 1,821 cars	1.65%	53.6%	2.50	0.01
Weeks 21-30	3.21% 1,121 cars	4.51% 1,220 cars	1.29%	40.3%	1.62	0.10
Weeks 31-39	2.78% 539 cars	5.15% 601 cars	2.37%	85.3%	2.03	0.04

Our final online result is in Table 12, 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 (4.74 per 100 auctions) than for cars without a posted SCR (3.66 per 100 auctions). Similar to 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 12: Expected number of online bidders per 100 auctions

	No posted SCR	Posted SCR	Difference	% Difference	t-statistic	p-value
All weeks	3.66 3,980 cars	4.74 4,118 cars	1.08	29.8%	2.22	0.026
Weeks 21-30	3.77 2,602 cars	4.72 2,798 cars	0.95	25.3%	1.58	0.11
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 tied to dealers learning about SCRs.

4.3.2 Substitution

It is important to rule out that bidders were responding to something other than the information disclosed. One concern may be that bidders respond to SCRs, but not in any way that is tied to their informational content. Suppose that the emails that buyers received from week 31 onwards focused buyers' attention on cars with posted SCRs without affecting their willingness to pay for cars with posted SCRs relative to cars without posted SCRs (i.e., a pure "salience" effect.) This can lead 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 makes it more likely that reserve prices were met, thus increasing the probability of sales. Specifically, consider the sales percentages in Table 2. The probability of a 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 is it that SCRs *did* increased 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 perform two analyses. First, we focus on weeks 31-39 and look for evidence of substitution by comparing the effect of SCRs across different types of cars. Second, we use all weeks of the experiment (21-39) in an expanded dataset (that includes fleet-seller consigned cars) to estimate a secular trend in the probability of sale for our sample period.

Our first analysis is based on the results in Table 8. This table offers strong evidence that bidders respond to the informational content of the SCRs and not to their mere presence. In particular, for the close-to-expected middle tercile, where SCRs have no informational content, there is no significant effect of SCRs on sale probabilities. If there were a pure salience effect that influenced sales above and beyond any informational content of the SCRs, then the effect should have been present for the mid-tercile. In contrast, for the worse-than and better-than-expected terciles, where the SCRs do have informational content, there is a positive effect. We conclude that salience is not likely to drive our results.

In our second analysis we estimate a secular trend in the probability of sale for our sample period. Specifically, we would like to know whether the decline in the sales probability of cars without a posted SCR from weeks 21-30 to weeks 31-39 reflected a general market trend (in which case substitution is a less likely explanation) or exceeded the market trend (which is what substitution would suggest).[14]

To estimate a secular trend we need data that were not part of the market in which the experiment was conducted: We choose cars that were offered for sale by fleet-sellers. For these cars there was no change in available information due to the experiment[15] 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

[14]Estimating a secular trend is 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, sales probabilities may have been affected by these events.

[15]More than 98% of fleet-seller consigned cars receive some type of inspection by the auction company. The inspection is generally not as thorough as the inspection that underlies the SCR in our experiment. The exact nature of fleet-seller inspection depends on the requirements of the fleet-seller and thus varies by fleet-seller.

somewhat newer, the overlap in age and condition between fleet-seller and dealer consigned cars is high.

The probability of sale for fleet-seller consigned cars is 67.25% in weeks 21-30 (13,491 cars) and 59.83% in weeks 31-39 (12,864 cars). This means the sales probability for fleet-seller consigned cars decreased by 7 percentage points, suggesting that demand for cars at the auction site decreased from weeks 21-30 to 31-30. 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.[16]

The results are in column 1 of Table 13. 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. To account for correlation in the errors when a car is offered for auction more than once during out sample periods, we cluster the standard errors at the VIN level. 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 change between weeks 21-30 and weeks 31-39 in the probability of sale for dealer consigned cars without a posted SCR was the same as for fleet-seller consigned cars. 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 emails 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.

A concern may be that the type of cars that are sold by fleet sellers are not comparable to cars that are sold by dealers and that therefore fleet-seller consigned cars are not suited for estimating the secular trend. We can (partially) address this concern by re-estimating the specification in column 1 of Table 13 with model-year fixed effects, vehicle segment fixed effects, nameplate fixed effects, sale week fixed effects, and some (non-SCR) measures which proxy for the condition of the car, namely the mileage of the car, and whether the car was offered under a green, yellow,

[16]The maintained assumption in using this difference-in-differences approach is that fleet-seller consigned cars and dealer consigned cars are subject to the same secular trend. While we cannot test whether this is the case during the treatment period, we can test for equality of pre-treatment trends between fleet-seller and dealer consigned cars. Using data from the beginning of the year to one week before the experiment started (19 weeks), we estimate a linear probability model that estimates a linear time trend in the probability of sale for cars, separately for fleet-seller and dealer consigned cars. The results are in Table 19. We cannot reject the hypothesis that the secular trend in probability of sale is the same for fleet-seller and dealer consigned cars.

Table 13: Linear probability model: diff-in-diff specification†

Dependent Variable: Sold	(1)	(2)
Dealer-consigned car, no posted SCR	-.24**	-.27**
	(.012)	(.015)
Dealer-consigned car, posted SCR	-.23**	-.27**
	(.012)	(.015)
Week 31-39	-.07**	
	(.0066)	
Week 31-39 * Dealer-consigned car, no posted SCR	.031	.029
	(.019)	(.02)
Week 31-39 * Dealer-consigned car, posted SCR	.089**	.087**
	(.02)	(.019)
Mileage on Car		1.6e-07
		(1.0e-07)
Green light		.14**
		(.0081)
Yellow light		-.011
		(.01)
Blue light		-.11**
		(.0096)
Sale Week Fixed Effects	no	yes
Model Year Fixed Effects	no	yes
Vehicle Segment Fixed Effects	no	yes
Nameplate Fixed Effects	no	yes
Constant	.67**	.66**
	(.0049)	(.2)
Observations	35287	35287
R-squared	0.034	0.119

* significant at 5%; ** significant at 1%; + significant at 10% level. SEs (robust and clustered at the VIN) in parentheses.

† Notice that our specification does not distinguish between fleet-seller consigned cars with and without inspections. This is because the inspections are not comparable to the inspections that yield SCRs in our experiment. In addition, more than 98% of fleet-seller consigned cars have some form of inspection.

or red light and a blue light.[17] This means that we are now identifying the secular trend and the result of inspections within cars of the same make, model year, segment, and approximate condition. As can be seen in column 2 of Table 13, there is very little change in the estimates.

Another 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 publicizing SCRs by email (i.e. during weeks 31-39).

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

4.3.3 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 factors. Specifically, we estimate a linear probability model of whether a car sold on the treatment, controlling for seller fixed effects (267), model year fixed effects (13), vehicle segment fixed effects (21), nameplate fixed effects (38), sale week fixed effects (9), condition grade fixed effects (5), and some (non-SCR) measures which proxy for the condition of the car, namely the mileage of the car, and whether the car was offered under a green, yellow, or red light and a blue light.[18]

Columns 1 and 2 of Table 20 show the results. As a reference, column 1 reports the treatment effect on the probability of sale without fixed effects (which is also in Table 2). Column 2 contains the treatment effect on the probability of sale controlling for the various fixed effects. The point estimate of the treatment effect drops from 6.3 percentage points to 4.8 percentage points. However, we can't reject the hypothesis that the treatment effect is unchanged by the

[17] The seller of every car sold at auction has to offer their car under some lights. A green light means that the seller declares that the car has no known mechanical problems. A yellow light means that the seller declares that the car has no known mechanical problems other than those are listed (e.g., "rough engine"). A red light means that the seller sells the car "as is" with no assurance to its mechanical condition. The auction company will arbitrate disputes that may arise for cars that were offered under a green and yellow light if the buyer finds undisclosed mechanical problems. A blue light means that the title of the car is not at the auction site.

[18] See footnote 17 for an explanation of lights.

inclusion of the extensive set of fixed effects. Columns 3 and 4 of Table 20 show that controlling for fixed effects does not alter our conclusion that average prices seem not to have significantly increased due to SCRs.

We repeat this robustness test for the results that are decomposed by conditions grades (Tables 5 and 6). Specifically we control for model year fixed effects, vehicle segment fixed effects, nameplate fixed effects, sale week fixed effects, the mileage of the car, and whether the car was offered under a green, yellow, or red light and a blue light.We also cluster our standard errors by VIN. Column 1 of Table 21 contains the treatment effect information disclosure on the probability of sale by condition grade. The relevant comparisons to the effects listed in Table 5 under the "Difference column" are the first five coefficients in the table. The coefficient of condition grade 1 increases slightly from 0.054 to 0.067 but is statistically indistinguishable. The coefficient for condition grade 5 drops slightly from 0.19 to 0.16, also a statistically indistinguishable difference. Similarly for the other condition grades: we can't reject the hypothesis that the treatment effect is unchanged by the inclusion of an extensive set of controls. Column 2 of Table 21 shows that our controls have a somewhat larger effect on prices. Our conclusion that the posted SCR changes transaction prices for condition grades 3 and 4 no longer holds. We now find that posted SCRs have no effect on prices regardless of condition grade.

Finally, we repeat this analysis for the results that are decomposed by the difference from the expected conditions grade (Tables 8 and 9). Column 1 of Table 22 contains the treatment effect of information disclosure on the probability of sale by difference from expected condition grade. The relevant comparisons to the effects listed in Table 8 under the "Difference column" are the first three coefficients in the table. The coefficients vary little from estimates in Table 8: Our conclusion remains that SCRs positively affect the probability of sale for cars with worse and with better than expected condition grades but does *not* change the probability of sale for cars with close to expected condition grades. Column 2 of Table 22 shows that, just as we found in Table 9, prices seem unaffected for any difference from the expected condition grade.

In summary, the conclusion of the key specifications in the paper are unaffected by the inclusion of a large number of fixed effects. In combination, these results provide no evidence that our procedure yielded a non-random assignment to treatment and control groups.

4.3.4 Alternative definition of condition grade

As we discussed earlier, 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 auction house 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 it 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 CG and estimated costs measures 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). 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, consistent with Table 5, the results point to the opposite.

5 Information Disclosure in Auctions Revisited

We highlight the discrepancies between our experimental environment and the standard assumptions, and demonstrate that the empirical findings in Section 4.2 are at odds with the standard implications of information disclosure. We proceed to outline a simple new theoretical framework that both rationalizes the empirical findings in Section 4.2, and suggests additional empirical implications that are borne out in our data.

5.1 The Linkage Principle and the Allocation Effect

The "Linkage Principle" derived in the seminal work of Milgrom and Weber (1982) (henceforth, MR) provides the benchmark of how information disclosure affects auction outcomes. It shows that in a symmetric affiliated values auction setting, the seller can increase expected revenues if he commits to releases all of his information ex ante.[19] The release of information causes the assessments of the bidders to be more congruent, resulting in lower "information rents" for the winning bidder, and as a result, increases competition. 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.

Recall that our experiment introduces SCRs, thus adding more information about a vehicle's condition. Hence, if the assumptions under which the Linkage Principle holds are satisfied, then the introduction of SCRs implies two empirical predictions. First, expected revenues should increase. Second, for vehicles whose reported SCRs reveal high (low) CGs, revenues should be higher (lower) than for vehicles of *the same CG quality* for which SCRs are not revealed.

Our empirical results in Section 4 confirm that expected revenues increase. However, they increase also for very low CG. Even with a CG of 1, the lowest possible quality level, the release of this "bad news" causes higher expected revenues than revealing no information, violating the basic prediction that bad news causes expected revenues to decrease. This empirical finding is

[19] To 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. See Milgrom and Weber (1982).

striking because any rational expectations model in which bidders have monotonically increasing values in quality will imply that disclosing the worse quality information must result in lower values, and hence lower bids than no information disclosure (Milgrom, 1981).

Theoretical studies have shown that the Linkage Principle may fail when the assumptions of MR are not satisfied. As Board (2009) observes, MR impose two simultaneous assumptions. First, bidders are symmetric. Second, their valuations are monotonic in the information (in our case the *CG*). As a result, the order of valuations coincides with the order of types, which in turn implies that the release of information does not change the type who wins the auction but only the expected price. Board (2009) shows that when either of these assumptions are dropped then the Linkage Principle may fail.

As a simple illustrative example, imagine that the seller's item has quality q uniformly distributed over $[0, 5]$, and there are two different bidders. The first bidder (H) has a valuation equal to $v_H(q) = q$ and the second bidder (L) has valuation $v_L(q) = 2 + \frac{q}{5}$. Hence, the H bidder values relatively high quality ($q > 2.5$) more than an L bidder, while the reverse is true for relatively low quality ($q < 2.5$). If the seller discloses no information and uses a second-price auction then each bidder bids his expected value, both equal to $2\frac{1}{2}$, and revenue is $2\frac{1}{2}$. If, instead, the seller discloses the realization of q then bidder 1 bids $b_1 = q$ while bidder 2 bids $b_2 = 2 + \frac{q}{5}$. Revenue is then $\min\{b_1, b_2\}$, which equals q if $q \leq 2\frac{1}{2}$ and $2 + \frac{q}{5}$ if $q > 2\frac{1}{2}$. Expected revenue is then equal to 2, less than the expected revenue without information disclosure. This simple example illustrates what Board (2009) labels the *allocation effect*, where new information changes the type of bidder who wins the good. Simply put, asymmetry implies a kind of horizontal differentiation across bidders.[20]

The potentially negative impact of the allocation effect is inconsistent with our data since revenues post information disclosure increase. Bidder heterogeneity, however, and the implied allocation effect may still be present in our setting. Board (2009) shows that with many bidders, revenues will increase when more information is disclosed. To see this, imagine that there are four bidders, two of type H and two of type L. With no information disclosed, everyone bids $2\frac{1}{2}$ and revenue is $2\frac{1}{2}$. If information is disclosed then *two bids* are equal to $b_H = q$ while *two other bids* are equal to $b_L = 2 + \frac{q}{5}$. The price is then $\max\{b_L, b_H\}$, which equals q if $q \geq 2\frac{1}{2}$ and $2 + \frac{q}{5}$ if $q < 2\frac{1}{2}$. The expected revenue is now 3, consistent with our finding that revenues are higher post information disclosure.

Still, our finding that even disclosing the worse quality information yields higher expected revenues than disclosing no information at all for the same quality level (Tables 9 and 10) remains puzzling. As implied from Milgrom (1981), if preferences are monotonically increasing in quality then revenues obtained with no information disclosure can *never be lower* than revenues obtained following the worse disclosed information.[21] Hence, through the lens of conventional bidding

[20] An earlier example showing the failure of the Linkage Principle with multi-unit auctions was derived by Perry and Reny (1999), yet the underlying forces share much in common.

[21] If some preferences are decreasing in quality then bad information could lead to higher prices. Imagine two types: $v_1 = q$ and $v_2 = 5 - q$, with two bidders of each type and $q \sim U[0, 5]$. Information disclosure cause

models our empirical findings still beg an explanation.

5.2 Information as a Matching Mechanism

Most auction models, including Milgrom and Weber (1982) and Board (2009), assume that one auction is being conducted at any given time and that the set of bidders at the auction is fixed.[22] Both these assumptions are violated in our environment because multiple auctions are conducted at any given moment and bidders have to exclusively choose which of these to participate in. Perhaps, the disclosure of SCRs affects the decisions of bidders regarding *which items* to bid on, if at all. This section develops this idea by constructing a simple two-type, two-good example to analyze a situation in which heterogeneous bidders choose which of two heterogeneous items to bid on.[23]

In discussions with industry participants we learned that used car dealers are indeed heterogeneous and seem to specialize in the condition of the vehicles that they sell. Dealers sell to customers in their geographical vicinity, implying that local demographics will shape their value for different vehicle quality levels. For instance, high income consumers will not be interested in a beaten-up vehicle, while low income consumers cannot afford to be as picky. Hence, for vehicles with low CGs, dealers from low income neighborhoods will outbid their counterparts from high income neighborhoods. Nevertheless, the willingness to pay for high CGs will be higher for high income consumers who have a higher marginal willingness to pay for improved quality. Hence, in reference to the example we use earlier, dealers in low income neighborhoods seem similar to L type bidders whereas dealers in high income neighborhoods seem similar to H type bidders, resulting in horizontal differentiation across CGs.

We also learned that auction bidders are quite experienced in assessing the condition of vehicles. Recall that the CG reflects mostly the exterior and interior condition of the vehicle. By observing the vehicles from up close, a relatively quick visual inspection can identify to a large degree whether the CG ought to be low, high or somewhere in between. As a consequence, once a bidder shows up at a lane and sees a vehicle, they have a pretty good idea of its condition as measured by the CG. Hence, *conditional* on a bidder showing up at an auction, the information revealed by the SCR is not very discriminating.

Our discussions with industry participants suggest that a formal analysis of our environment should include three basic assumptions. The first is that bidders are heterogeneous and horizontally differentiated with respect to CGs (A1). The second is that there are several goods selling

revenues to be "V" shaped: $5 - z$ if $z \leq 2.5$ and z if $z \geq 2.5$, which are highest when $z = 0$ and when $z = 5$. This example violates the whole notion of calling q quality since we typically think of quality as a dimension over which preferences are increasing and monotonic.

[22] Some models consider a random number of bidders (e.g., McAfee and McMillan, 1987) while others consider endogenous entry of bidders (e.g., Levin and Smith, 1994). These studies consider one auction, so that the endogenous choice of *which auctions* to participate in, which is the focus of our analysis, has not been considered.

[23] Developing and analyzing a more general formal model is beyond the scope of this paper as it would be a challenging stand-alone theoretical analysis.

at several *mutually exclusive* simultaneous auctions (A2). The third is that the disclosure of SCRs may help bidders find the vehicles they are interested in, but upon seeing a vehicle, the information content of the SCR is small (A3).[24] To proceed, we develop a simple example based on these assumptions as follows:

Preferences: Consider two types of bidders (following A1), $\theta \in \{L, H\}$, with $v_H(q) = q$ and $v_L(q) = 2 + \frac{q}{5}$ as described above and depicted in Figure 1. The quality of vehicles q is random and uniformly distributed on the interval $[0, 5]$.A bidder i of type θ has a value $v_{\theta i}(q) = v_\theta(q) + \varepsilon_i$ where ε_i is a private shock that is independently and uniformly distributed over $[-\bar{\varepsilon}, \bar{\varepsilon}]$, with $\bar{\varepsilon}$ being very small. Hence, the expected value of a type θ bidder from a vehicle of quality q is $E[v_{\theta i}(q)] = v_\theta(q)$. We assume that there are four bidders, *exactly two* of each type.

Mechanism: There are two open ascending auctions on two lanes that sell vehicles simultaneously and each bidder can only be present at one lane at a time (following A2). Quality is independent across vehicles and lanes.

Information: Sellers can either disclose nothing, or they can disclose perfect verifiable information about their vehicles' quality $q \in [0, 5]$. Once bidders arrive at a lane, they perfectly observe the quality q (following A3), but before choosing which lane to attend, bidders only know what the seller chooses to disclose.[25]

As before, horizontal differentiation in quality is captured by the fact that $v_L(q) > v_H(q)$ for $q < 2.5$, while $v_L(q) < v_H(q)$ for $q > 2.5$. With no disclosure both types have an expected value of 2.5. Let $v^q_{\min} \equiv \min_\theta v_\theta(q)$, the lower of the two expected valuations, and $v^q_{\max} \equiv \max_\theta v_\theta(q)$, the higher of the two.

Timing proceeds as follows. Bidders observe the information (if any) disclosed by the seller and then choose which lane to participate in. An equilibrium will be characterized by a lane choice, followed by the standard dominant strategy of bidding up to one's valuation in an ascending auction. To make things simple, assume that there are two distinct vehicles, one with grade $q < 2.5$ and the other with grade $q' > 2.5$, and their assignment to one of two lanes is random.

5.2.1 No Disclosure

Assume first that the seller does not disclose the grades of the vehicles. We have (proofs appear in the appendix),

[24] That is, "small" means that the CG plays a more important role in matching bidders to cars, and a less important role in revealing information once a bidder arrives at an auction.

[25] Alternatively, information can takes the form of a partition of the quality interval. For example, a grade of $g \in \{1, 2, ..., 5\}$ could corresponds to the vehicles true quality being uniformly distributed in the interval $[g-1, g]$. The qualitative results and comparative statics that follow would persist.

Figure 1: Valuations of two types if bidders

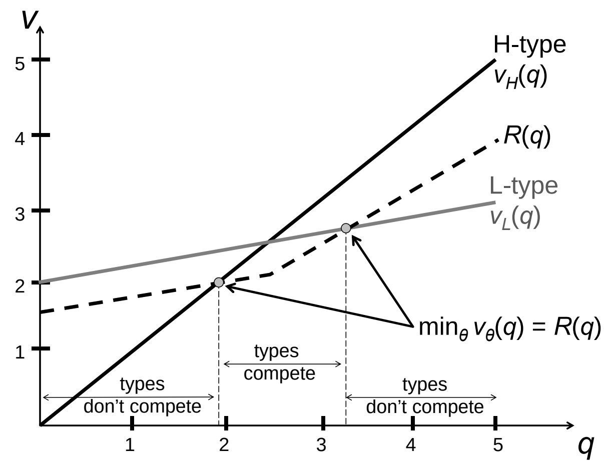

Claim 1: If there is no disclosure then there are two equilibria: a pure-strategy *coordinated equilibrium* where exactly one bidder of each type is in each lane, and a mixed-strategy *random equilibrium* where each bidder chooses each lane with equal probability.

The two equilibria identified in Claim 1 have different outcomes. In the coordinated equilibrium with no disclosure the allocation is efficient, but the expected winning bid of a quality q vehicle is always $ER[q \mid \text{ND coordinated}] = v^q_{\min}$.

The random equilibrium results in sixteen distinct outcomes with equal probabilities and the expected price is the expected second highest value of the bidders. There are three configurations of bidders at a lane that are of interest. First, if no more than one bidder shows up then the price is 0. For any given lane this happens with probability $\frac{5}{16}$. Second, if more than one bidder shows up, but no more than one of them has value $v^q_{\max}$, then expected revenue is $v^q_{\min}$. This happens with probability $\frac{7}{16}$. Last, if more than one bidder shows up, and two of them are of the type with $v^q_{\max}$, then expected revenue is $v^q_{\max}$.[26] This happens with probability $\frac{1}{4}$. We can now write the expression for expected revenues of a quality q vehicle when no information is disclosed as follows:

$$ER[q \mid \text{ND random}] = \frac{5}{16} \times 0 + \frac{7}{16} \times v^q_{\min} + \frac{1}{4} v^q_{\max} .$$

[26]Throughout the analysis we ignore the ε_i's which play a tie-breaking role for identical types by adding some natural idiosyncrasies. The qualitative conclusions are valid as long as $\bar{\varepsilon}$ is not too large. If we don't ignore the ε_i's then we need to calculate the second order statistic of ε_i to correctly determine the expected price in this situation. However, as $\bar{\varepsilon} \to 0$ the expected price goes to $v^q_{\max}$.

5.2.2 Full Disclosure

Now assume that the seller discloses the quality of the vehicle he puts up for sale in the lane. We have,

Claim 2: Given two vehicles with qualities $q < 2.5$ and $q' > 2.5$ auctioned in two lanes, the unique equilibrium has perfect sorting: both L types choose the q-lane and both H types choose the q'-lane.

The intuition for this result is a simple consequence of optimal sorting. Each type will select into the lane where they have a comparative advantage. Hence, information disclosure plays a role as a matching mechanism. Given Claim 2 it is easy to see that with disclosure, the expected revenue of each vehicle is $ER[q \mid \mathrm{D}] = v^q_{\max}$.

5.2.3 Comparing Information Policies

The following corollary follows from the analysis above:

Corollary: Information disclosure increases expected revenues for any given quality. The impact is larger as quality moves farther away from 2.5, the value at which the two types' valuations cross. Furthermore, with information disclosure the variance of winning bids is lower than in the random equilibrium with no information disclosure.

Hence, information disclosure increases expected revenues for *any quality level* regardless of the equilibrium played in the auction with no disclosure. Furthermore, the increase in expected revenues is larger as the grade moves away form the "middle" grade of 2.5, yielding a "U-shaped" effect of information disclosure on expected revenues. The reason that the increase in revenues is U-shaped followed from the fact that $v^q_{\max} - v^q_{\min}$ is increasing as q moves away from the point at which $v_H(q) = v_L(q)$, as illustrated in Figure 1. If the equilibrium play is random then for any level of q, sometimes the seller will receive $v^q_{\max}$, sometimes he will receive $v^q_{\min}$, and sometimes he will receive 0. With information disclosure, however, the seller always receives $v^q_{\max}$, increasing the expected price and reducing the price variance (to zero).

The intuition is similar to that of the allocation effect identified by Board (2009). If heterogeneous bidders are at a lane and a vehicle comes through with a high grade then the H type wins, while the opposite happens for a low-grade vehicle. What differs in our setting is that the ex ante arrival of information on grades causes bidders to *endogenously choose lanes* where they can win, and as a consequence the composition of bidders at lanes is rearranged to create assortative matching. Matching guarantees that two high valuation bidders will be present, intensifying effective competition.

Notice also that for grades in the middle range, close to where $v_H(q) = v_L(q)$, the two types are similar in their valuations. Assortative matching, therefore, has less of an impact when the two types are similar. In contrast, as the grade is closer to the extreme values of 1 and 5, the difference between the types' valuations increases, making the matching-effect stronger. This will play an important role in rationalizing the empirical findings described earlier in Tables 9 and 10.

5.2.4 The Effect of Reserve Prices

What is obviously missing in the example above is the use of reserve prices by sellers. As described in Section 2, about half the vehicles do not sell on any given auction day since their reserve price is not met. In many of these cases the seller keeps the vehicle at the site, which the auction house offers at no charge, to be auctioned later in the week or during following weeks.[27] The ability to keep an unsold car at the auction site at no extra cost results in an "outside option" that is surely not zero.[28] Reserve prices that reflect this option must be considered to correctly predict the effect of information disclosure on auction outcomes.

Imagine that the seller has a small opportunity costs $k > 0$ of keeping the car at the auction house for the next auction, and that there is a discount factor $\delta < 1$ for each period of delay between auctions. Consider the random equilibrium in our example above with no disclosure. The seller expects one of three outcomes: a price of zero (with probability $\frac{5}{16}$), a price of $v^q_{\min}$ (with probability $\frac{7}{16}$) and a price of $v^q_{\max}$ (with probability $\frac{1}{4}$). Given these beliefs, the seller will prefer to reject a bid b and wait for $v^q_{\max}$ if,

$$\begin{aligned} b &< \delta(-k + \frac{1}{4}v^q_{\max}) + \frac{3}{4}\delta^2(-k + \frac{1}{4}v^q_{\max}) + \left(\frac{3}{4}\right)^2 \delta^3(-k + \frac{1}{4}v^q_{\max}) + \cdots \\ &= \frac{\delta}{1 - \frac{3}{4}\delta}(-k + \frac{1}{4}v^q_{\max}) \equiv r(q). \end{aligned} \tag{1}$$

For small opportunity costs (δ close to 1 and k small), the value of $r(q)$ will be somewhat below the "upper envelope" of $v^q_{\max} = \max_\theta v_\theta(q)$, as depicted by the dashed-line in Figure 1.[29]

The observation made earlier that assortative matching has less of an impact when the two types are similar implies that there is an important difference between mid-range quality/grades, where $v^q_{\max} - v^q_{\min}$ is small, and between either very low or very high quality levels, where $v^q_{\max} - v^q_{\min}$ is large. As depicted in Figure 1, there are two values of quality at which $r(q) = v^q_{\min}$. In the interval between these values, the two different types are similar enough so that a $v^q_{\min}$ bid is not rejected. The H and L types are "competitive enough" not to make the option of waiting

[27] Our data is consistent with this since we see vehicles that did not sell in previous auctions being offered at later auctions, often with the same work order. This implies that they were left on the lot for a future auction day.

[28] Two other alternatives are available to dealers whose cars did not sell. They have the option of returning the car to their own lot, where there is some chance it can sell, instead of waiting several days at the auction site. Another way to sell a car is using wholesale buyers who visit dealer lots to buy cars that the dealers have a hard time selling and then relocate those cars to other dealers.

[29] As $\delta \to 1$ (low opportunity cost of delay) the limit of the right-hand side of (1), R, equals $-4k + v^g_{\max}$, and as $k \to 0$ (low opportunity cost of leaving the car at the auction) this equals $v^g_{\max}$. The order of limits is inconsequential.

for $v^q_{\max}$ worthwhile. In contrast, when quality is outside of this interval, $v^q_{\max} - v^q_{\min}$ is large and the option value of waiting for $v^q_{\max}$ makes rejecting $v^q_{\min}$ bids worthwhile.[30]

With disclosure, however, the likelihood of obtaining $v^q_{\max}$ goes up (to probability 1 in our simple example), implying that there will be a larger impact on the probability of sale when the grade is farther away from the middle. Furthermore, with opportunity costs that are relatively low, the curve $r(q)$ will be close to $v^q_{\max}$, and *conditional on selling,* the price difference between revealing or not revealing information will be small. That is, most of the effect will be on the probability of sale.

5.3 Empirical Evidence for Information as a Matching Mechanism

We consider whether our simple example of information as a matching mechanism is consistent with the results from our experiment. To do so we begin by analyzing whether there is evidence in the experimental data for a key premise our model, namely that dealers horizontally differentiate with respect to condition grades. Next, we test two empirical implications from the example: (1) that information disclosure changes auction outcomes in the way hypothesized in section 5.2.4, and (2) that information disclosure changes the variance of condition grades that any given bidder chooses to bid on.

5.3.1 Evidence that Bidders are Horizontally Differentiated in Condition Grades

We begin with verifying that bidders are horizontally differentiated with respect to condition grades. To do so we split each dealer's purchases into "early" and "late" car purchases: "early" car purchases encompass the first 50% of cars purchased by the dealer during our sample period; "late" car purchases encompass the remaining cars that the dealer bought. We test whether the condition grades of cars purchased "early" by each dealer predicts the condition grades of cars purchased "late" by the corresponding dealer.

We begin by calculating for each dealer the average condition grade of cars purchased "early" and "late". If dealers specialize in cars of specific condition grades we would expect that the average condition grades between the two samples are positively correlated. Indeed, for dealers who purchased more than 2 cars during the sample period, we find that the correlation coefficient is 0.45 (p-value <0.01). Another way of analyzing specialization is to calculate a transition matrix between the condition grades chosen for "early" purchases and "late" purchases. Specifically, recall that for each dealer we measure the average condition grade of cars purchases "early" and "late". We split these average condition grades into quintiles and calculate what percentage of dealers who were in a specific quintile for "early" purchases are in the same quintile for "late" purchases. For the 407 dealers who purchased more than 2 cars during the sample period, we

[30]The variation due to the private noise determined by $\bar{\varepsilon}$ will effect the reserve price strategy, but will still result in some bound below the upper envelope of $v^q_{\max}$, and the qualitative comparative static results will continue to hold.

find the following transition matrix:

Condition Grade Quintile		"Late" purchases 1	2	3	4	5	Total
	1	34	14	9	7	2	66
		51.52%	21.21%	13.64%	10.61%	3.03%	100%
	2	28	21	17	10	10	86
		32.56%	24.42%	19.77%	11.63%	11.63%	100%
"Early"	3	13	21	24	15	12	85
purchases		15.29%	24.71%	28.24%	17.65%	14.12%	100%
	4	19	7	15	21	30	92
		20.65%	7.61%	16.30%	22.83%	32.61%	100%
	5	8	9	17	15	29	78
		10.26%	11.54%	21.79%	19.23%	37.18%	100%
	Total	102	72	82	68	83	407
		25.06%	17.69%	20.15%	16.71%	20.39%	100%

Clearly, buyers who choose cars of particular condition grades during "early" car purchases tend to choose cars of similar condition grades during "late" car purchases as well. These findings are consistent with the assumption of our simple example, namely that bidders are heterogeneous and horizontally differentiated with respect to condition grades.

5.3.2 Evidence on how auction outcomes change with information disclosure

Under our sorting theory we should see that information disclosure (a posted SCR) increases the probability of sale for cars in the bottom (worse-than expected) and top (better-than expected) terciles, and not for cars in the middle (close-to-expected) tercile. This is consistent with our earlier results: Table 8 shows that our experimental data for weeks 31-39 follow this prediction. Moreover, since early during the experiment the wide availability of SCRs was not publicized, we should not find the hypothesized pattern during weeks 21-30. Indeed, as Table 14 shows, there is no statistically significant effect of a posted SCR on the probability of sale for cars in any of the terciles.

Table 14: Sales probability by difference of expected condition grade (CG), weeks 21-30

Tercile of Difference from Expected CG	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	z-statistic	p-value
Worse than expected	1802	0.383	0.375	-0.08	-0.2%	-0.36	0.72
Close to expected	1800	0.429	0.452	0.02	4.6%	0.99	0.32
Better than expected	1800	0.477	0.483	0.005	1.3%	0.23	0.82

Overall, our results are consistent with the prediction from our example that there will be a larger impact on the probability of sale when the grade is farther away from the expected condition grade. Furthermore, as predicted when the opportunity costs of not selling are low, we find that most of the effect will be on the likelihood of sale, and conditional on selling, the effect on prices will be small. This latter point can be seen in Table 9 for weeks 31-39 and in Table 15 for weeks 21-30.

Table 15: Price/NAP by difference of expected condition grade (CG), weeks 21-30

Tercile of Difference from Expected CG	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	t-statistic	p-value
Worse than expected	680	0.99	0.98	-0.006	-0.6%	-0.35	0.73
Close to expected	781	1.09	1.08	-0.019	-1.7%	-0.88	0.37
Better than expected	847	1.1	1.1	0.004	0.36%	0.24	0.81

5.3.3 Evidence on Information Disclosure as a Matching Mechanism

Next, we consider the choice of bidders regarding which grades of vehicles to bid on. Ideally, we would observe that after information is disclosed we have less variance in the *CG* of vehicles that any given bidder chooses to bid on. Unfortunately, due to the nature of the auctions we do not know on which cars bidders choose to bid, but only on the cars that bidders successfully won. Using a variance test on the vehicles that bidders win is not informative. The reason is that given the endogenous choice of the reserve price, both with and without information disclosure the right type of bidder should be the winner most of the time.

Instead, we can indirectly see if there is a response of bidders to the disclosed *CG* information. The auction registration process assigns vehicles to lanes and this is done prior to the SCRs being generated. During weeks 21-30 bidders know where vehicles are but they have less information about them. As a consequence, the benefit of switching from one lane to another in search of better matched vehicles is not large. After week 30, however, bidders have more information about the vehicles they are interested in, and know where they are. We expect, therefore, that for any given number of vehicles that a bidder buys, he will have visited more lanes after week 30.

We first regress the number of vehicles purchased by each dealer per week on the number of lanes in which the dealer purchased the cars. We allow this relationship to differ for weeks 21 to 30 and 31 to 39, respectively. To ensure that relationship between number of lanes and number of purchased vehicles is estimated from within-dealer variation in the number of cars purchased over time, we estimate all specifications with buyer fixed effects. The results are in column 1 of Table 16.

As hypothesized, after week 30, buyers on average use more lanes than up to week 30: Up to week 30, for every additional car purchased, dealers purchase cars on 0.47 additional lanes. Starting in week 31, for every additional car purchased, dealers purchase these on 0.64 (0.47+0.17) additional lanes. Notice, however, that this relationship should only hold for cars with an SCR.

Table 16: Number of lanes used by dealers per week†

	All Cars	SCR Cars	Non-SCR Cars
Number of cars	.47**	.42**	.49**
	(.05)	(.075)	(.076)
Week 31-39	-.21**	-.31*	-.17+
	(.067)	(.12)	(.1)
Week 31-39 * Number of cars	.17**	.25*	.13
	(.055)	(.098)	(.082)
Buyer Fixed Effects (837)	yes	yes	yes
Constant	.58**	.64**	.55**
	(.062)	(.097)	(.096)
Observations	2690	1401	1289
R-squared	0.779	0.796	0.843

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

† An observation is a dealer-week conditional on the dealer having made any purchases during a week. If a dealer makes any purchases during a week, on average a dealer purchases 1.47 cars per week.

This is because even after week 31, dealers have no additional information about cars without an SCR. In columns 2 and 3 of Table 16 we thus split our sample into cars with an SCR and cars without an SCR. As can we seen, there is only a difference in how many lanes are used before and after weeks 30 for cars with an SCR: The interaction between the dummy for weeks 31-39 and number of cars is only significant for cars with an SCR but not for cars without an SCR.

6 Concluding Remarks

It is well established that information disclosure can help market participants better evaluate the value of goods and services they are interested in, often resulting in more efficient outcomes and less distortionary information rents. For example, Lewis (2010) shows that by voluntarily disclosing private information on “eBay Motors”, sellers may effectively be offering protection to buyers from adverse selection. This revealing insight helps explains the prevalence of many online transactions that otherwise may seem puzzling due to potential “lemons” concerns.

We have demonstrated that in addition to these important effects, information disclosure can play an important role in providing information that helps buyers choose *which market* to participate in. This simple, yet novel insight has broader applications beyond our market for used automobiles. If heterogeneous participants can sort into markets for heterogeneous goods, then better ex ante information will help them sort into markets for which they have the most value, and in turn, *effective competition* will intensify in all markets. Turning back to eBay's huge marketplace, when sellers reveal more information then buyers can self-select into those auctions that they are most interested in.[31]

The stylized model we offer is tailored to the auction environment we analyze, and was successfully used both to rationalize the empirical results of our paper and to generate hypotheses that could be tested in the particular auctions that we study. Having exclusive simultaneous auctions for similar yet differentiated goods like the ones we study may be close in spirit to online auction sites such as eBay, but other markets will have different institutional details. Developing a general model of information disclosure in markets is beyond the scope of this paper, yet the intuitive driving forces behind our results seem both fundamental and more general. For example, a firm looking to hire people for similar, yet distinct positions may gain from providing more information on its positions, even if the information for some positions may make them seem unattractive relative to others. If a firm posts job vacancies for two positions that share some similarities, each position will receive a more refined and better matched pool of applicants if more information is released that distinguishes the two positions in terms of requirements, skills and job descriptions.

The implications of information disclosure as a matching mechanism may also apply to government procurement. Typically, governments engage in both parallel and sequential procurement of many similar, yet distinct projects. For example, there may be several construction, road or defense acquisition projects that are let out to bid simultaneously, and yet many more are anticipated to materialize within weeks or months. Though these are often thought of as sequential and not exclusive simultaneous auctions, bidders (contractors) with capacity constraints may not be able to bid on later auctions if they win earlier ones. If the procurement authority not only releases information on current tenders, but also releases detailed information on future tenders, then heterogeneous contractors may be able to better select which of the coming auctions to participate in, which in turn may increase effective competition both within and between projects.

[31] As one executive in the company that provided the data commented on this idea, "one man's trash is another man's treasure."

7 References

Board, Simon. (2009) "Revealing Information in Auctions: The Allocation Effect," *Economic Theory*, 38(1):125-135

Busse, Meghan R., Jorge Silva-Risso, and Florian Zettelmeyer. (2006) "1000 Cash Back: The Pass-Through of Auto Manufacturer Promotions," *American Economic Review*, 96(4):1253-1270.

Cho, Sung-Jin, Harry J. Paarsch and John Rust (2010) "Is the 'Linkage Principle' Valid?: Evidence from the Field," mimeo

Cutler, David M., Robert S. Huckman and Mary Beth Landrum. (2004) "The Role Of Information In Medical Markets: An Analysis Of Publicly Reported Outcomes In Cardiac Surgery," *American Economic Review Papers and Proceedings*, **94(2)**:342-346

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.

Hishleifer, Jack (1971) "The Private and Social Value of Information and the Reward to Inventive Activity." *American Economic Review*, **61(1)**:561–574.

Jin, Ginger Z. (2005) "Competition and Disclosure Incentives: An Empirical Study of HMOs," *RAND Journal of Economics*, Spring, **26(1)**:93-113.

Jin, Ginger Z. and Phillip Leslie (2003) "The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards," *Quarterly Journal of Economics*, **118(2)**:409-51.

Kagel John H. and Dan Levin (1986) "The Winner's Curse and Public Information in Common-Value Auctions," *American Economic Review*, **76(5)**:894-920.

Kagel, John H., Ron 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.

Klemperer, Paul (1999) "Auction Theory: A Guide to the Literature," *Journal of Economic Surveys* **13(3)**:227-286.

Levin, Dan and Janus L. Smith, (1994) "Equilibrium in auctions with entry," *American Economic Review* **84(3)**:585-99.

Levin Dan, John H. Kagel and Jean-Francois Richard (1996) "Revenue Effects and Information Processing in English Common Value Auctions," *American Economic Review*, **86(3)**:442-460.

Lewis, Gregory (2010) "Asymmetric Information, Adverse Selection and Online Disclosure: The Case of eBay Motors," *American Economic Review*, forthcoming.

R. Preston McAfee and J. McMillan (1987) "Auctions with a stochastic number of bidders," *Journal of Economic Theory* 43(1):1-19.

Milgrom, Paul R. (1981) "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 12(2):380-391.

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.

Ottaviani, Marco, and Andrea Prat (2001) "The Value of Public Information in Monopoly," *Econometrica* **69(6)**:1673–83.

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

Claim 1: If there is no disclosure then there are two equilibria: a pure-strategy *coordinated equilibrium* where exactly one bidder of each type is in each lane, and a mixed-strategy *random equilibrium* where each bidder chooses each lane with equal probability.

Proof: Consider pure strategies. First, it is easy to see that if one lane has no bidders then any bidder in the other lane would prefer to switch lanes and win at a price of zero. Second, imagine that only one bidder is choosing lane 1. Each of the two identical bidders in lane 2 would have a strict preference to switch lanes *ex ante* since in lane 1 each can win half the time and obtain some rents, while staying in lane 2 they will either lose to the third bidder, or they will win but compete away most of the rents (with the exception of the difference in their ε_i's). This last argument also rules out an equilibrium where each lane has two bidders of the same type since every bidder will have an incentive to switch lanes. The only other pure strategy configuration is where one L-type and one H-type are in each lane, each winning half the time. In this configuration, when winning a vehicle with quality q, the winner obtains expected rents equal to $\max_\theta\{v_\theta^q\} + \varepsilon_i - \min_\theta\{v_\theta^q\}$. Any bidder who switches from this configuration will compete with his own type, implying that he would obtain expected rents strictly less than $\bar{\varepsilon} < \max_\theta\{v_\theta^q\} + \varepsilon_i - \min_\theta\{v_\theta^q\}$. Hence, this is the unique pure strategy equilibrium. Consider mixed strategies. It is easy to see that randomly choosing a lane with equal probability is an equilibrium since a bidder who believes that the other bidders are using this strategy is indifferent between the two lanes. No other mixed strategy profile can be an equilibrium because if some bidder of type θ is choosing a lane with probability greater than $\frac{1}{2}$ then the best reply of the other bidder of the same type would be to choose the other lane with probability 1 to increase the probability of winning with positive rents. ■

Claim 2: Given two vehicles with qualities $q < 3$ and $q' > 3$ auctioned in two lanes, the unique equilibrium has perfect sorting: both L types choose the q-lane and both H types choose the q'-lane.

Proof: We show that in any other configuration, at least one bidder has an incentive to switch lanes. First, it is easy to see that random assignment is not an equilibrium. If all the other bidders are choosing lanes randomly, then an H type bidder has a strict incentive to choose the q' lane since his probability of winning that vehicle is higher, and conditional on winning, he is left with higher rents. (A symmetric argument applies to a L type choosing the q lane.) Second, with pure strategies it is easy to see that if one lane has no bidders then anyone from the other lane would have preferred to switch lanes. Third, imagine that there is only one bidder in lane q and three in lane q'. If the sole bidder in lane q is an H-type then each of the L-types in lane q' has an incentive to switch lanes since they lose in the q' lane and they would win in the q lane and obtain rents. (A symmetric argument holds for a sole L-type in the q' lane.) If the sole bidder in lane q is an L-type then the L-type in lane q' has an incentives to switch lanes. If he stays in lane q' then he loses for sure against the two H-types. If he switches, then there is a positive probability that his idiosyncratic noise ε_i is greater than that of the other L-type in lane q, in which case the switching bidder would win and obtain a small rent.[32] (A symmetric argument holds

[32] If there is no idiosyncratic noise then this bidder would be indifferent between losing to the H-types and switching lanes only to see the price of the g vehicle rise to v_L^q, leaving him with no rents.

for a sole H-type in the q' lane.) Finally, assume that each lane has two bidders, one of each type. In this case the H-type in lane q and the L-type in lane q' both lose, while if they switched then there is a positive probability that each of their idiosyncratic noise ε_i is enough to make them win and obtain a small rent. To complete the analysis, observe that if the two types perfectly sort as stated in Claim 2 above then no one has an incentive to switch. Each has a positive probability of winning and obtaining a small rent, while by switching each is guaranteed to lose. ■

Table 17: Dealer-consigned and inspected cars by week†

Sale Week	Dealer-Consigned Total	With SCR Not reported	With SCR 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 18: Summary Statistics

Variable	N	mean	p50	sd	min	max
Model Year	8098	2003.5	2004	2.7	1997	2009
Mileage	8098	75958.6	71315.5	44359.1	0	508112
Condition Grade	8098	2.42	2	1.31	1	5
Repair Costs	8098	1347.9	1024	1236.7	0	16110.8
Sold	8098	0.43	0	0.50	0	1
Sales Price	3481	8660.8	7300	5929.9	500	59000
National Auction Price	3429	8397.2	6975	5810.8	200	62000
Sales Price/National Auction Price	3429	1.06	1.03	0.24	0.24	5.6

* The number of observations for the "National Auction Price" and "Sales Price/National Auction Price" is lower than for "Sales Price" because the "National Auction Price" is missing for a few cars in our data.

Table 19: Pre-promotion trends: Sales probability during weeks 1-19

	Sold
Time Trend	-.0045**
	(.0005)
Fleet-Seller	.33**
	(.0084)
Fleet-Seller*Time Trend	-.00096
	(.00073)
Constant	.48**
	(.0057)
Observations	57513
R-squared	0.105

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

Table 20: Randomization check on aggregate results: Sales probability and Transaction Prices for weeks 31-39

	Sales Probability		**Transaction Prices**	
	Base Result	Fixed Effects	Base Result	Fixed Effects
Posted SCR	.063**	.048*	.02	.0077
	(.019)	(.021)	(.012)	(.013)
Condition grade=2		.044		.038*
		(.031)		(.019)
Condition grade=3		.053+		.064**
		(.031)		(.023)
Condition grade=4		.15**		.078**
		(.042)		(.02)
Condition grade=5		.067		.069**
		(.046)		(.024)
Mileage on Car		6.2e-07		2.7e-07
		(4.4e-07)		(3.3e-07)
Green light		.089+		.17**
		(.047)		(.046)
Yellow light		-.041		-.033
		(.033)		(.028)
Blue light		-.12+		-.0056
		(.07)		(.037)
Seller Fixed Effects	no	yes	no	yes
Model Year Fixed Effects	no	yes	no	yes
Vehicle Segment Fixed Effects	no	yes	no	yes
Nameplate Fixed Effects	no	yes	no	yes
Sale Week Fixed Effects	no	yes	no	yes
Observations	2696	2696	1121	1121
R-squared	0.004	0.273	0.002	0.426

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

Table 21: Randomization check on results by CG: Sales probability and Transaction Prices for weeks 31-39

	Sales Probability	Transaction Prices
Posted SCR * CG = 1	.067+ (.034)	.031 (.022)
Posted SCR * CG = 2	.074 (.05)	.0047 (.024)
Posted SCR * CG = 3	.015 (.041)	.036 (.023)
Posted SCR * CG = 4	.1 (.066)	-.029 (.023)
Posted SCR * CG = 5	.16* (.067)	-.02 (.027)
Condition Grade = 2	.061 (.044)	.08** (.023)
Condition Grade = 3	.12** (.04)	.087** (.022)
Condition Grade = 4	.19** (.055)	.13** (.026)
Condition Grade = 5	.036 (.054)	.12** (.026)
Mileage on Car	-3.2e-07 (4.2e-07)	4.4e-07 (3.0e-07)
Green light	.1* (.042)	.19** (.036)
Yellow light	-.029 (.032)	-.044* (.022)
Blue light	-.1 (.066)	.015 (.037)
Model Year Fixed Effects	yes	yes
Vehicle Segment Fixed Effects	yes	yes
Nameplate Fixed Effects	yes	yes
Sale Week Fixed Effects	yes	yes
Constant	.26 (.2)	.85** (.13)
Observations	2696	1121
R-squared	0.079	0.224

* significant at 5%; ** significant at 1%; + significant at 10% level. Robust and clustered (by VIN) SEs in parentheses .

Table 22: Randomization check on results by expected CG: Sales probability and Transaction Prices for weeks 31-39

	Sales Probability	Transaction Prices
Posted SCR * CG Worse than Expected	.07+ (.036)	.025 (.021)
Posted SCR * CG Close to Expected	.018 (.037)	.017 (.021)
Posted SCR * CG Better than Expected	.1** (.037)	.0045 (.018)
CG Close to Expected	.088* (.037)	.059** (.021)
CG Better than Expected	.1** (.037)	.099** (.02)
Mileage on Car	-5.9e-07 (4.2e-07)	1.9e-07 (2.9e-07)
Green light	.11* (.042)	.19** (.036)
Yellow light	-.029 (.031)	-.041+ (.022)
Blue light	-.11 (.067)	.012 (.037)
Model Year Fixed Effects	yes	yes
Vehicle Segment Fixed Effects	yes	yes
Nameplate Fixed Effects	yes	yes
Sale Week Fixed Effects	yes	yes
Constant	.27 (.21)	.87** (.14)
Observations	2696	1121
R-squared	0.075	0.218

* significant at 5%; ** significant at 1%; + significant at 10% level. Robust and clustered (by VIN) SEs in parentheses .