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

Steven Tadelis
UC Berkeley and NBER

Florian Zettelmeyer
Northwestern University and NBER

November 12, 2014

Abstract

Market outcomes depend on the quality of information available to the market's participants. We measure the effect of information disclosure on market outcomes using a large-scale field experiment that randomly discloses information about quality in wholesale automobile auctions. We argue that buyers in this market are horizontally differentiated across cars that are vertically ranked by quality. This implies that information disclosure helps match heterogeneous buyers to cars of varying quality, causing both good and bad news to increase competition and revenues. Our empirical analysis confirms these hypotheses. These findings have implications for the design of other markets, including e-commerce, procurement auctions, and labor markets.

JEL classifications C93, D44, D82, L15

*We are indebted to the management and employees of the firm that provided the data and worked cooperatively with us to implement the experiment. We thank Meghan Busse and Igal Hendel for helpful discussions, and many seminar participants for comments. We are grateful to Karin Mauge from ebay Research Labs for pointing us to supporting evidence from ebay. Tadelis thanks the National Science Foundation for financial support.

1 Introduction

Information asymmetries often cause markets to be inefficient because they introduce adverse selection between buyers and sellers. These asymmetries may be reduced when sellers disclose information, either voluntarily or by regulatory mandate. As a result, the topic of information disclosure had indeed received much attention in the economics, finance, accounting, law and marketing literatures.[1] The ways in which information disclosure will affect market outcomes will, however, depend on the institutional features of the market.

We study the effects of information disclosure on market outcomes in a market setting in which bidders have to choose which of the many simultaneous, and mutually exclusive auctions they wish to participate in. Many large markets such as automobile sales, e-commerce sites, and government procurement share these features. We investigate the wholesale market for used automobiles where trade between car dealers is facilitated through auctions and where sellers can control the amount of information that they choose to disclose. Using a randomized field experiment, we are able to precisely document how more information affects auction outcomes.

We hypothesize that in these settings, disclosing information about the quality of objects will help bidders choose to participate in the auctions that they value the most. Perhaps surprisingly, revenues will increase even if the information disclosed about the object is negative. Even though cars are vertically ranked by quality, participants have horizontally differentiated preferences over quality rankings. As a result, if the information disclosed about quality is surprising, whether good or bad, the bidders who value that level of quality the most will choose to participate in the auction. This will result in more intense competition and higher revenues for auctions in which the news was a surprise, but will have no effect if the information disclosed was anticipated.

We test our hypotheses using a unique randomized field experiment. In a market where thousands of vehicles are sold each week with an average value of $8,500, we manipulate information disclosure while keeping all other aspects of the auction fixed. The disclosed information has a clear ranking of quality, which allows us to test how revenues change as a function of the informational content, and in particular, whether the information disclosed is either *better or worse* than expected. The empirical results are striking: When the information disclosed is either a positive or negative surprise relative to expectations, the seller benefits from increased revenues. When the information is consistent with prior expectations, revenues remain the same.

We can use the rather trivial example of horizontal differentiation to explain our findings more simply. Imagine that some people like blue cars, and some people like red cars, but people do not know at which of several auctions the different colored cars are sold. In this instance, sometimes blue car

[1] See, e.g., Grossman (1981), Diamond and Verrecchia (1991), Healy and Palepu (2001), Gilson and Kraakman (1984) and Day (1976), and the references therein.

lovers will appear at red car auctions and vice versa, resulting in a low willingness to pay for each car. If sellers would just announce in advance where each color car will be sold, buyers would sort which auctions to attend according to their preferences. Competition would then be fiercer at each auction, resulting in higher revenues. This would not surprise anyone.

In our setting, however, information discloses quality along a vertical dimension over which all bidders agree on the ranking: other things equal, everyone prefers higher quality. Hence, the analogy to red and blue cars is less obvious. We observe that if one bidder has a higher *marginal* value for a quality than another, he may still have a lower *absolute* value for lower quality levels, making the total willingness-to-pay vary across the quality range. Consider consumer A who just needs a car to get around. He has a moderate willingness to pay for any quality car, but his marginal benefit from higher quality is rather small. Consumer B, who cares a lot about quality, may be willing to pay almost nothing for the lowest-quality cars, but has a high marginal value for extra quality. Hence, for low-quality cars consumer A has a higher value than consumer B, while for high-quality cars their values switch order. Despite the fact that both consumers agree on the vertical dimension of quality (more is better), they are nonetheless horizontally differentiated along the quality dimension.[2] We discuss several other markets where our insights apply, arguing that the matching value of information has widespread applications.

This observation, which to the best of our knowledge has not been made before, contributes to the existing literature on the disclosure of vertically ranked information by moving beyond the more standard situation containing a fixed set of symmetric bidders at a single auction. In the standard setup, as Milgrom and Weber (1982) demonstrate, more information will improve revenues on average, but upon the arrival of bad news, revenues will fall below those from the no-disclosure setting. More generally, Milgrom (1981) shows that bad news about a single object will cause valuations to drop for every consumer who is considering that object. We instead consider a setting where asymmetry across bidders, coupled with their need to choose which of many auctions to participate in, will cause both *good and bad* news to create a better match between bidders and the quality they value relatively more than others. The standard literature focused on how information disclosure increases competition *within* a given auction or market, while we focus on the increased competition *across* auctions or markets.

We also contribute to the growing empirical literature on the effects of information disclosure on market outcomes.[3] Due to the challenge of testing how variation in information disclosure affects auctions in the field, there have been few such studies, and most refer to the Linkage Principle as their testing ground.[4]

[2]As an executive in the company that provided the data commented, "One man's trash is another man's treasure."

[3]See, for example, Porter (1995); Jin and Leslie (2003); and Lewis (2010). Anand and Shachar (2011) investigate the matching properties of TV ads for programs that are horizontally differentiated.

[4]See De Silva et al. (2008) and Cho et al. (2014).

The paper proceeds as follows. Section 2 describes the industry, the auction details, the information provided to bidders, and the way in which bidders are differentiated. Section 3 presents a simple theoretical framework from which we derive testable hypotheses. Section 4 describes the data and the experimental design, while Section 5 presents the experimental results that are consistent with our hypotheses. Robustness tests are performed in Section 6, and Section 7 presents concluding remarks.

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 from trade-ins or the wholesale market for used automobiles.

Wholesale automobile auctions provide a prominent source of used cars. According to the National Automobile Dealers Association (NADA), 35 percent of all used cars sold by new car dealers in 2008 were sourced in auctions.[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

Buyers at our auctions are exclusively dealers, while sellers mainly belong to one of three categories: dealers who 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 has ended. Sellers bring their vehicles to the auction site one or more days before the auction. Each vehicle is assigned "lane" and "run" numbers. Several thousands of vehicles may be auctioned off during a sale day. The vehicles are lined up in several (up to twelve) lanes, according to the lane and run numbers.[7]

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, options, 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.

[5]See the National Independent Automotive Dealer's Association (NIADA) website (http://www.niada.com/) for its 2010 annual report.

[6]See NADA DATA (2009), 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. Every auction takes about thirty seconds, implying that this vehicle will be offered for sale about sixty-six minutes after the auctions starts.

Each lane has an auction block from which an auctioneer conducts the auction, one car at a time, so that up to twelve auctions can occur simultaneously. The vehicle that is next in line to be sold is 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 thirty seconds and ends when no bidder is willing to raise the price.[9] If the price exceeds the seller's reserve price, the sale is consummated. About half the vehicles do not sell on any given auction day because their reserve price is not met. In many of these cases the seller keeps the vehicle at the site for no extra charge, to be auctioned later in the week or during following weeks.[10]

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 bring 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 them secretly to the auction house.[11] The auction house then informs the high bidder if the sale is accepted after the auction ends.

There are two distinct classes of bidders at the auction. "Lane" bidders are physically present at the auction and can visually inspect the car up close. Prior to the bidding, vehicles are parked outside so that potential bidders who arrive early enough can examine their exterior condition. "Online" bidders 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 basic information and, for the lane bidders, the potential to visually inspect the car and listen to the engine of those cars that can be driven. Because potential buyers cannot perform a serious inspection of the vehicles (not to mention the disadvantage of the online bidders, who cannot see

[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 place for decades (see Genesove, 1995, p. 26), and we have been told that it is also common in livestock auctions.

[10] The seller can also return the car to his lot. For the vehicles in our sample, 70 percent are consigned only once, 17 percent twice, 7 percent three times, 3 percent four times, 2 percent five times, and 1 percent more than five times.

[11] Dealers may change the reserve price after the car has been checked into the auction lot. However, from our discussions with the auction house personnel, we understood that the random assignment of the experiment described below did not cause sellers to change the reserve prices for their cars. This supports our assumption in Section 3 that ex ante, sellers know that ex post, bidders will almost perfectly observe the condition of the vehicle and bid accordingly.

the vehicles in any detail), there is residual uncertainty about a vehicle's quality. As a response, many auction houses offer condition reports that describe the vehicle's condition in more detail. Historically, fleet-sellers have requested tailor-made condition reports for the cars they sell, but dealer-sellers have not followed suit. The output from these reports was not standard, and buyers were not always pleased with the way in which information was presented.

In response, the auction house from which this paper's data originates developed a Standard Condition Report (SCR) that offers a standard set of inspections, and presentation format. The SCR is based on a detailed inspection that takes about twenty minutes per car. The inspections cover the vehicle's exterior condition, documenting all imperfections (including any additional layers of paint that imply 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 the information through a computerized hand-held device that registers the information on a central computer, and creates a standardized report.

The SCR is then posted online in a one-page format that documents a detailed summary of the inspection and two other summary statistics. First, a "condition score" (CS) is calculated based on the input of the inspection.[12] The grading system runs from 1 through 5, with increments of 0.1, where $CS = 1.0$ is considered "rough," and $CS = 5.0$ 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 condition score and the estimated costs offer standardized measures of vehicle quality.

We learned that auction bidders are quite experienced in assessing a vehicle's condition. A relatively quick up-close visual inspection can identify to a large degree whether the condition score ought to be low, high, or somewhere in between. Once bidders show up at a lane and see a vehicle, they have a pretty good idea of its condition as measured by the condition score. That is, *conditional* on a bidder showing up at a vehicle's auction, the information revealed by the SCR is not very useful.

2.3 Dealer Heterogeneity

Discussions with industry participants reveal that used-car dealers are heterogeneous. Dealers sell to customers in their geographical vicinity, implying that local tastes will shape their values for different vehicles. For instance, high-income consumers are not interested in beaten-up, low SCR vehicles, while low-income consumers cannot afford to be as picky. Dealers from low-income neighborhoods will then outbid their counterparts from high-income neighborhoods on low-SCR cars. High-income consumers

[12] Genesove (1993) and Overby and Jap (2009) also investigated the role of condition reports.

will pay more for cars in better condition than low-income consumers because their marginal value for appearance is greater. This implies that a car's quality, in the form of its condition score, has elements of horizontal differentiation despite the fact that it is usually considered as a vertical dimension.

To confirm that bidders are horizontally differentiated with respect to condition scores, we split each dealer's purchases into "early" and "late" car purchases: early purchases comprise the first 50 percent of cars purchased by the dealer during our sample period; late purchases comprise the remaining cars that the dealer bought. For each dealer we calculate the average condition score of cars purchased early and late. If dealers specialize in specific condition scores, we would expect the average condition scores between the two periods to be positively correlated. Indeed, for dealers who purchased more than two cars during the sample period, we found that the correlation coefficient is 0.45 (p-value < 0.01).

Another way of establishing specialization is to calculate a transition matrix between the condition scores chosen for early and late purchases. We split the early and late average condition scores into quintiles and calculated the percentage of dealers who were in a specific quintile for early purchases and in the same quintile for late purchases. Using 406 dealers who purchased more than two cars during the sample period, Table 1 illustrates the transition matrix and confirms that buyers who chose cars of particular condition scores during early car purchases tended to choose cars of similar condition scores during late car purchases as well.

Table 1: Early vs. late purchase transition matrix

	Condition Score Quintile	Late purchases 1	2	3	4	5
	1	51.56%	26.56%	9.38%	10.94%	1.56%
Early	2	28.74%	27.59%	21.84%	11.49%	10.34%
purchases	3	16.09%	24.14%	27.59%	21.84%	10.34%
	4	17.65%	8.24%	16.47%	23.53%	34.12%
	5	14.46%	9.64%	20.48%	18.07%	37.35%

3 Theoretical Framework

The most famous theoretical result on the effect of information disclosure on auction outcomes is the Linkage Principle, which was derived in the seminal work of Milgrom and Weber (1982). It shows that in a symmetric affiliated values auction setting, a seller who commits to disclose all information ex ante can increase expected revenues. Intuitively, information disclosure causes the assessments of the bidders to be more congruent, which results in lower "information rents."

The result is derived assuming that one auction is being conducted at any given time and that the set of bidders at the auction is fixed. Our environment violates both assumptions because multiple auctions are conducted simultaneously and bidders have to choose only one in which to participate. In fact, the observations we made in Section 2 suggest that our environment can be described by three basic features. First, bidders are heterogeneous and horizontally differentiated with respect to condition scores. Second, several goods are selling at several *mutually exclusive,* simultaneous auctions. Third, the disclosure of SCRs may help bidders find the vehicles they are interested in, but once they see a vehicle, the information content in the SCR is small.

Based on these observations, we outline below a simple example to guide our empirical analyses.[13] Imagine that cars have a quality q uniformly distributed over the interval $[0, 1]$, and there are two different types of bidders. The high-quality bidder (H) has a valuation equal to $v_H(q) = q$ and the low-quality bidder (L) has valuation $v_L(q) = 0.25 + 0.5q$. Hence, the H bidder values relatively high quality ($q > 0.5$) more than an L bidder, while the reverse is true for relatively low quality ($q < 0.5$).[14]

Now imagine that two open ascending auctions on two lanes offer vehicles simultaneously, that each bidder can be present at only one lane at a time, and that quality is independent across vehicles and lanes. The auction house can either disclose nothing, or it can disclose perfect, verifiable information about quality $q \in [0, 1]$. Once bidders arrive at a lane, they perfectly observe the quality q, but before choosing which lane to attend, bidders know only what the auction house chooses to disclose. To make things simple, assume that there are two distinct vehicles, one with quality $q_L < 0.5$ and the other with quality $q_H > 0.5$, and that they are randomly assigned to two different lanes. It clearly would be in the best interest of each bidder to go to the lane with his preferred vehicle, because then, given a fixed level of competition, he would face a car that offers him the highest value.[15]

If the auction house does not disclose information before the auctions take place, each bidder can do no better than to choose a lane randomly. However, if the auction house does disclose information about the vehicles' quality before the auction, then each type of bidder will flock to the lane with his preferred vehicle: the L type to the vehicle with $q_L < 0.5$ and the H type to the vehicle with $q_H > 0.5$, a consequence of optimal sorting. Each type will select into the lane where a comparative advantage exists, and information disclosure acts as a matching mechanism.

This observation is related to Board (2009), who observed that when bidders are asymmetric, information disclosure may affect who wins the good, which he labeled the "Allocation Effect." Simply put, asymmetry implies horizontal differentiation for vertically different goods. In our setting, unlike in the fixed set of participating bidders in Board (2009), the ex ante arrival of information on quality

[13]Developing and analyzing a more general formal model is beyond the scope of this paper.

[14]It is important that the the low type's value has a greater intercept and a smaller slope, so the value functions cross at an interior quality level. This requires more than the single-crossing condition to generate horizontal differentiation.

[15]To be precise, a fixed level of competition would be determined by the expectation over the number and types of other bidders that would attend the lane-auction. Given the simplicity of our example, the presence of at least one more competitor of the same type would reduce a bidder's profits to zero because the price would be bid up to its valuation. A more realistic example would have an additive idiosyncratic random value ϵ_i as part of the bidders' valuations, which in turn would give any bidder a positive expected profit from participating.

causes bidders to *endogenously choose lanes* where they can win.[16] Better matching of bidders with cars implies that bidders will be able to better target cars across lanes during the auction.

Hypothesis 1: Information disclosure helps bidders target the vehicles they choose to bid on.

Because information disclosure results in better matching between bidders and vehicles (L types with q_L vehicles and vice versa), it follows that information disclosure intensifies competition for *any* quality level. This conclusion may seem surprising. In any rational expectations model in which bidders have increasing values in quality, disclosing "bad news" (low quality) results in lower values, and hence in lower bids than those with no information disclosure (Milgrom, 1981). However, this logic applies to auctions with a fixed set of bidders. In our setting the disclosure of information matches bidders to auctions, resulting in fiercer competition even for the lower-quality scores.

Better matching will differentially affect different quality vehicles. In our example, the difference in valuations between L and H types is larger for vehicles at the extreme quality values, and smaller for vehicles in the middle of the quality range. Hence, the effects of information disclosure on improving sorting will be more pronounced for vehicles that are at the extremes of the distribution, where the two types are not close in valuation. This in turn means that revenues should increase more as the news becomes more extreme.

Before exploring how revenues will be affected, recall from Section 2 that about half the vehicles do not sell on any given auction day because their reserve price is not met, and in many of these cases the seller keeps the vehicle at the site (see footnote 10). This implies that sellers have an "outside option" that is surely not zero.[17] Reserve prices must therefore be considered to correctly predict the effect of information disclosure on auction outcomes.

Imagine that the seller has a small opportunity costs of keeping the car at the auction house until the next auction, and is very patient. In this case, a seller with a low (or high) quality vehicle would ideally want to wait until at least two low (or high) type bidders arrive at the auction for his car. In that case, the seller would set a reserve price close to the value of the type that most values the car. For example, if $q = 0.2$, then setting a price just below $v_L(0.2) = 0.25 + 0.5 \times 0.2 = 0.35$, guarantees that the car sells only when at least two L types appear at the auction. Recall that the H type only values this car at $v_H(0.2) = 0.2$. If sellers indeed have a low opportunity cost of leaving the car for future auctions, then this reserve price strategy implies that information disclosure will primarily affect the probability of a car selling, and less the price conditional on sale. That is, the reserve price will be set high enough so that the car sells only when the effective competition is strong. As a result, better

[16] Roberts and Sweeting (2011) show that selection effects may have subtle auction design implications even when one auction is considered.

[17] Dealers can also return 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.

sorting increases the likelihood that several of the right-type bidders will show up and bid above the reserve price. This results in the following testable hypothesis:

Hypothesis 2: Information disclosure increases the probability of sale for any given quality level, and the impact is larger for qualities at the extremes of the quality distribution.

Hypothesis 3: If reserve prices are adequately set by a patient seller, then the impact on prices of cars sold will be small across all quality levels.

4 Data

4.1 Experimental Design

The purpose of the experiment was to measure the treatment effect of SCRs on the probability of sale and final price for cars that were consigned to the auction by used-car dealers. A large number of dealer-consigned cars were inspected at one auction location over the course of nineteen weeks using the SCR inspection procedure, and 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; only the auction house knew that these cars had been inspected and their corresponding condition scores.

Due to a limited number of certified vehicle inspectors, not all dealer-consigned cars were inspected. The number of inspected cars depended on the number of available inspectors during that week (between three and twelve). For an auction conducted on Wednesday of a given week, all cars that were checked in starting Friday morning of the prior week were candidates for inspection. On days with many inspectors, all cars that were checked in until mid-day Tuesday were inspected, whereas on days with few inspectors, inspections were performed on cars that were checked in until some time on Monday. Specifically, out of approximately 1,500 dealer-consigned vehicles that were registered each week, between 150 and 600 cars were inspected per week (see Table A-1). In total, 8,098 cars were inspected, 3,980 of which were in the control group (SCR not reported) and 4,118 were in the treatment group (SCR reported).

During the check-in process, cars whose vehicle identification number (VIN) ended in an even digit were assigned to the treatment group, while those with an odd digit were assigned to the control group. The first digits of a VIN number designate manufacturer, country of origin, make, model, model-year, as well as some trim-level information, and 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 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 percent of cars to the control group and 50.85 percent to the treatment group.[18]

As we analyzed auction outcomes after the first nine weeks of the experiment, we found little evidence that cars with SCRs were more likely to sell or sold at higher prices (these findings are described in Section 4). One possibility was that the information contained in SCRs had little content. Alternatively, perhaps dealers did not know that SCRs were made available for many dealer-consigned cars. As discussed earlier, SCRs are available only online, not on the standard printout that dealers can obtain in advance on the website or on auction day at the facility. Hence, for the remainder of the experiment a weekly email was sent to all registered buyers informing them that they could find SCRs for some of the dealer-consigned cars on a particular website prior to the auction day. The emails stated that the company was ramping up its capabilities to offer SCRs, and technicians were assigned to inspect vehicles that were chosen randomly based on inspector availability. It was made clear that these were not solicited or affected by the sellers. As a result, our experiment covers two periods: weeks 21-30 (5,402 cars), during which dealers were not likely to have been aware of the existence of SCRs, and weeks 31-39 (2,696 cars), during which SCRs were publicized. This variation will prove useful in analyzing the data and shedding light on the impact of information disclosure.

4.2 Auction and Inspection Data

For each consigned car we observed the model, model-year, body type, engine and trim level (e.g., a Honda Accord, 1999, 4-door, V6, EX trim), as well as its mileage. More detailed information about the car's condition came from the SCRs as described in Section 2.2. We used two key measures. The first was the condition score, a number between 1 (rough) and 5 (clean). The second was the estimated cost to fix the damage detailed in the SCR. This included the auction house's estimates of both part and labor costs and is reported in dollars.

A unique seller ID allowed us to identify whether different cars were consigned by the same seller. The data reports whether a car was sold during the auction, the final auction price, and a unique buyer ID allowing us to identify whether different cars were purchased by the same buyer. Finally, we had the average auction price for cars of the same type that sold at any of the auction house's locations nationwide during the prior week (henceforth "National Auction Price" or NAP). This allowed us to construct a useful normalization of price that was independent of the type of car. Summary statistics are reported in Table A-2.

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

4.3 Randomization Check

We compared the treatment and control groups on a variety of observable characteristics. Specifically, if the randomization worked as intended, the distribution of condition scores, repair costs, mileage, vehicle age (model year), and national auction prices in the prior week should have been comparable across control and treatment groups. We used a Kolmogorov-Smirnov test for equality of distribution functions. The results are reported in Table 2.

Table 2: Kolmogorov-Smirnov test for equality of distribution functions

Variable	D	p-value
Condition score	0.0137	0.83
Repair costs	0.0301	0.05
Mileage	0.0172	0.58
Model year	0.0167	0.61
NAP	0.0246	0.17

For four of the five measures we failed to reject the hypothesis that the distribution functions were the same. However, the test statistic for repair costs was just at the critical level, indicating that repair cost may have had a different distribution between control and treatment groups. Repair costs for the control group were on average $1,382, while for the treatment group they were $1,316. We will account for this $66 (less than 5 percent) difference when interpreting our auction price results.

We will also explore our randomization when estimating treatment effects. We will analyze whether our estimates change as we add a large set of controls, namely fixed effects for seller ID, model year, vehicle segment, nameplate, and sale week, as well as measures of the car's condition. If our randomization procedure worked, then these controls should not substantially change our estimates. The results of this analysis are reported in Section 6, where we explore the robustness of our findings.

5 Results

We organize the results into four parts. First, in Section 5.1 we report the average effects of our experiment and show that information disclosure increased the likelihood that cars sold, and that, conditional on selling, they sold for a slightly higher price. After establishing these average effects, we turn to testing our hypotheses. In Section 5.2 we test Hypothesis 1 by analyzing whether information disclosure helps bidders target the vehicles they choose to bid on. In Section 5.3 we test Hypothesis 2 by investigating how information disclosure affects the probability of sale for cars of different quality levels. Finally, in Section 5.4 we test Hypothesis 3 by investigating how information disclosure affects auction prices for cars of different quality levels.

5.1 Average Effects of Information Disclosure

We considered first the probability of sale. Table 3 shows that during weeks 21-30, cars with and without a posted SCR were equally likely to sell; approximately 43 percent 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 email, cars with a posted SCR were 6.3 percentage points (or 16 percent) more likely to sell than cars without a posted SCR. This difference is highly statistically different from 0 (using a test of proportions with p-value < 0.01).[19] One concern in evaluating the statistical significance of this difference was that 30 percent of cars in our sample were offered for sale more than once during the sample period. As a result, the error terms for cars that were offered multiple times could have been correlated. To account for this potential correlation, we clustered the standard errors at the VIN level (for detailed results see section 6). The correction had a very small effect and did not alter our conclusions in Table 3 or any other table in Section 5.

Table 3: Sales probability by experimental condition

	No posted SCR	Posted SCR	Difference	% Difference	z-statistic	p-value
Weeks 21-30	0.43 2,605 cars	0.436 2,797 cars	0.006	1.39%	0.43	0.66
Weeks 31-39	0.392 1,375 cars	0.455 1,321 cars	0.063	16.1%	3.31	0.001

Next, we analyzed the effect of information disclosure on the auction price conditional on a sale. A problem we faced is the high price variance of sold cars, because sales include everything from 11-year-old small cars (e.g., Honda Civic) to current-year luxury cars (e.g., BMW 740). We 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 used the average auction price for cars of the same type that sold at any of the auction house's locations during the prior week. We used this measure to construct a normalized price for each car in the sample, specifically, the price of the car divided by the NAP. 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

[19] For a week-by-week graph of the probability of sale, see Figure A-1 in the online appendix.

After week 31, prices were higher by 1.9 percent for cars with a posted SCR relative to cars without a posted SCR. The difference was marginally significant (p-value of 0.11).

In interpreting these results, one concern was that bidders may have responded to something other than the informational content of SCRs. Of particular concern was that the emails from week 31 onward simply focused buyers' attention on cars with posted SCRs (i.e., a "salience" effect.) We will later show that bidder behavior was not consistent with such a salience effect. Instead, it seems that bidders reacted to the information contained in SCRs, as opposed to the mere existence of SCRs, as we later explain.

Overall, an analysis of the probability of sale and prices conditional on sale suggested that most of the effect of SCRs on expected auction revenues came from an increased probability of sale; transaction prices did increase, but only slightly.

5.2 How Information Disclosure Changes Bidder Behavior

We first analyzed evidence that information disclosure changes bidder behavior. Under Hypothesis 1 we should have observed that after information was disclosed, there was less variance in the condition score of vehicles that any given bidder chose to bid on. Unfortunately, we only observed the cars that bidders successfully won, and not the cars bidders chose to bid on. Using a variance test on the vehicles that a bidder wins was not informative, because given the endogenous choice of the reserve price, both with and without information disclosure, the right type of bidder should have won most of the time.[20]

Instead, we indirectly tested to see if bidders responded to the disclosed information. The auction registration process assigned vehicles to lanes before the SCRs were generated. During weeks 21-30 bidders knew where vehicles were but had less information about them. Hence, the benefit of switching from one lane to another in search of better matched vehicles was limited. After week 30, however, bidders had more information. Using this information, the benefit of switching lanes in pursuit of a better matched vehicle was higher. We expected, therefore, that for any given number of vehicles that a bidder bought, he would have visited more lanes after week 30.

We regressed the weekly number of vehicles purchased by each dealer on the number of lanes from which they were purchased. We allowed this relationship to differ for weeks 21-30 and 31-39, respectively. To ensure that the estimation was from within-dealer variation in the number of cars purchased over time, we estimated all specifications with buyer fixed effects. The results are in column 1 of Table 5.

As hypothesized, buyers on average used more lanes after week 30: up to week 30, for every additional car purchased, dealers purchased cars on 0.47 additional lanes. Starting in week 31, for every additional car purchased, dealers purchased these on 0.64 (0.47+0.17) additional lanes. In columns 2 and 3 of Table 5, we split our sample into cars with and without an SCR. We expected the relationship to hold

[20]The variance in quality of the cars bought by each bidder with and without SCRs was indeed the same.

Table 5: 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.

strongly for cars with a SCR. Indeed, the interaction between the dummy for weeks 31-39 and number of cars was significant for cars with an SCR. For non-SCR cars, the relationship might also have held, if buyers concluded that that some cars with an SCR were not the cars they were looking for. We found weak evidence for this. The point estimate of the interaction between the dummy for weeks 31-39 and number of cars was positive but not significantly different from zero.

5.3 The Effect of Information Disclosure on the Probability of Sale

We then tested Hypothesis 2: Information disclosure increases the probability of sale for any given quality level, and the impact is larger for qualities at the extremes of the quality distribution. To analyze the effect of posted SCRs by condition of the vehicle, one has to recognize that bidders have some ex ante information that can predict the condition score, namely mileage and age. As shown in Table 6, the average condition score varies substantially by vehicle age and mileage, as one would predict: cars that are older or that have higher mileage will, on average, have worse condition scores. This information allows buyers to predict the car's condition as a function of age and mileage.

As a result, it was necessary to perform an empirical test that explicitly allowed for condition score expectations that differed with vehicle age and mileage. We first estimated the predicted condition score of each car in our sample based on the vehicle age and vehicle mileage. We did this by regressing condition score on vehicle age dummies, vehicle mileage, and vehicle mileage deciles. To account for potential interaction effects, we interacted the vehicle age dummies with vehicle mileage. We took the difference between the *actual* condition score and the *predicted* condition score to construct a distance measure from the expected condition score. Finally, we split this distance measure into terciles, where

Table 6: Average condition score (CS) by mileage category and vehicle age

Mileage Category	Average CS	Vehicle Age	Average CS
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

the bottom tercile contained cars with worse-than-expected condition scores ("bad news"), the middle tercile contained cars with close-to-expected condition scores, and the top tercile contained cars with better-than-expected condition scores ("good news").[21]

Table 7: Sales probability by difference of expected condition score (CS), weeks 31-39

Tercile of Difference from Expected CS	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	z-statistic	p-value
Worse-than-expected	899	0.327	0.411	0.084	25.7%	2.61	0.009
Close-to-expected	899	0.429	0.418	-0.011	-2.6%	0.34	0.74
Better-than-expected	899	0.419	0.529	0.109	26.1%	3.28	0.001

As Table 7 shows, during weeks 31-39 there was no statistically significant effect of a posted SCR on the probability of sale for cars in the middle tercile, while in both terciles where condition scores had informational content, the effect on the probability of sale was positive and significant. This also alleviated the concern that SCRs created a salience effect. If they had, the effect should not have depended on the informational content, but rather on whether an SCR was posted.

Moreover, because early in the experiment the availability of SCRs was not publicized, we should not have found the hypothesized pattern during weeks 21-30. As Table 8 shows, there was no statistically significant effect of a posted SCR on the probability of sale for cars in any of the terciles.

One might be concerned that buyers form expectations not over condition scores but over the repair cost estimates contained in all SCRs. To address this concern, in section 6 we re-estimated the results in this subsection and the next subsection using terciles of "difference from expected repair cost." The

[21]The current definition of worse-than-expected, close-to-expected, and better-than-expected does not account for stock dynamics. In Section 7 of the online appendix we allow for expectations to be formed only over cars that were sold in recent weeks. None of our results changed.

Table 8: Sales probability by difference of expected condition score (CS), weeks 21-30

Tercile of Difference from Expected CS	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	z-statistic	p-value
Worse-than-expected	1,802	0.383	0.375	-0.08	-0.2%	-0.36	0.72
Close-to-expected	1,800	0.429	0.452	0.02	4.6%	0.99	0.32
Better-than-expected	1,800	0.477	0.483	0.005	1.3%	0.23	0.82

results were similar in that they preserved the "U-shape" of the impact of news on sales probability and showed no significant effect on prices.[22]

5.4 The Effect of Information Disclosure on Auction Prices

Next, we tested Hypothesis 3: If reserve prices are adequately set by patient sellers then the impact on prices of cars sold will be small across all quality levels. We investigated how information disclosure affected auction prices for cars of different quality levels. Table 9 shows that during weeks 31-39 the normalized average prices of cars with a posted SCR were slightly higher but statistically no different from those of cars without a posted SCR. This is consistent with Hypothesis 3. Since early in the experiment the wide availability of SCRs was not publicized, we should also not have found a price difference during weeks 21-30. Table 10 confirms this prediction.

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

Tercile of Difference from Expected CS	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	t-statistic	p-value
Worse-than-expected	324	0.978	0.999	0.022	2.2%	1.05	0.30
Close-to-expected	375	1.04	1.08	0.035	3.3%	1.58	0.11
Better-than-expected	422	1.07	1.08	0.006	0.6%	0.31	0.75

Table 10: Price/NAP by difference of expected condition score (CS), weeks 21-30

Tercile of Difference from Expected CS	# 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

[22] The only change is that the SCR effect for cars in the "worse-than-expected" estimated repair cost tercile is marginally significant (p-value 0.09).

In summary, we found support for all three hypotheses. First, in lieu of actually observing which cars bidders chose to bid on, we concluded that they used the disclosed information to switch lanes more often and target vehicles that were a better match for their potential customers. Second, we observed that both good news and bad news caused more sales. Third, we observed that both good and bad news did not lead to a significant increase in auction prices.

6 Robustness

This section presents further evidence about the reliability and interpretation of our results. We began with the concern that buyers form expectations not over condition scores but over the repair cost estimates contained in all SCRs. To address this concern, we re-estimated the main sales probability and price results using terciles of "difference from expected repair cost," instead of "difference from expected condition score." The estimation and results are in Section 1 of the online appendix. As in Section 5, during weeks 31-39 there was no statistically significant effect of a posted SCR on the probability of sale for cars in the middle tercile (close-to-expected estimated repair costs). However, in both terciles where estimated repair costs had informational content, the effect on the probability of sale was positive. The effect was strongly significant for the better-than-expected tercile (p-value $<$ 0.01) and marginally significant for the worse-than-expected tercile (p-value = 0.09). The price results were also similar to the condition score findings in that there was no statistically significant effect of a posted SCR on the probability of sale for cars in any of the "difference from expected repair cost" terciles.

We were also concerned that 30 percent of cars in our sample were offered for sale more than once during the sample period. We wanted to account for the potential correlation in the error terms for cars that were offered multiple times. We did so by estimating linear probability models with robust clustered standard errors at the VIN level. This accommodated arbitrary correlation between the errors of observations of the same car. We present the results in Section 2 of the online appendix. We found that clustering standard errors did not alter our conclusion that average prices seem not to have significantly increased due to SCRs.

We also revisited whether our randomization procedure yielded a random assignment to treatment and control groups. We analyzed whether our basic results changed after adding a large set of controls, namely seller fixed effects (267), model year fixed effects (13), vehicle segment fixed effects (21), nameplate fixed effects (38), sale week fixed effects (9), condition score tercile (3), and some (non-SCR) measures that represented the condition of the car. If our randomization procedure worked, these controls should not have substantially changed our estimates, because the randomization should have ensured that cars of the same make, model year, segment, and approximate condition were randomly assigned between treatment and control groups. The results of this analysis are in Section 2 of the online appendix. We can't reject the hypothesis that the treatment effect in probability of sale and price is unchanged by the inclusion of the extensive set of fixed effects, consistent with random assignment.

Next, we further explored our argument that SCRs had no effect on auction outcomes during weeks 21-30, because dealers were not aware that they had been posted. We confirmed this by exploring the behavior of dealers who must have been aware that SCRs were posted during weeks 21-30: Those who made use of the auction house's online bidding feature ("online dealers" for short). To access online bidding, dealers must use the web portal that posts SCRs. Moreover, this was the only source of information that puts online dealers on some equal footing with the on-site lane bidders. Showing that online dealers behaved no differently before and after week 31 would present evidence in support of our argument. In Section 2 of the online appendix we found that the effect of a posted SCR on the behavior of online dealers was similar between weeks 21-30 and 31-39. We concluded that the effect of SCRs we observed offline during weeks 31-39 was most likely tied to dealers learning about SCRs.

Next, we addressed a concern about the way we classified cars as either "worse-than-expected," "close-to-expected," or "better-than-expected." Recall that we constructed an expected condition score using a regression of condition score on vehicle age and vehicle mileage, and used the residuals of that regression to classify cars into terciles. A concern was that the variance of the residuals may have differed for different types of cars. If this were the case, by placing cars into terciles relative to all other cars, we may have categorized a car as "worse-than-expected," when it was "close-to-expected" relative to other cars of the same type. To address this concern, we categorized cars into terciles only in comparison to cars of the same standard car category, namely "Compact Car," "Fullsize Car," "Luxury Car," "Midsize Car," "Pickup," "Sports Car," "SUV," and "Van." For example, a Honda Civic would be placed into the "worse-than-expected" tercile if its residual was low relative to those of other compact cars. We then re-estimated Tables 7 and 9 with the car category specific terciles. As we show in Section 4 of the online appendix, there were no meaningful changes relative to Tables 7 through 10.

Next, we wanted to see whether using (coarse) terciles to categorize news masked findings for cars that were "much better-than-expected" or "much worse-than-expected." In Section 5 of the online appendix, we replicated our main analysis using quintiles instead of terciles. The results were largely consistent with our earlier findings preserving the U-shape of the impact of news (as seen in Table 7). However, the smaller cell sizes made the differences at worse-than-expected quality statistically insignificant. We found one new insight, namely that the sales probability effect of the "better-than-expected" tercile was driven by cars that were "much better-than-expected."

For our last robustness check, we felt that it was important to rule out that bidders were responding to something other than the informational content of SCRs. Of particular concern was that emails from week 31 onward focused buyers' attention on cars with posted SCRs (a "salience" effect.)

We addressed this concern earlier, by considering the results in Table 7. If SCRs caused a pure salience effect, bidders should have responded to SCRs *regardless* of their informational content. The table offers strong evidence that bidders responded only to the informational content of the SCRs and not to their mere presence. In particular, for the middle tercile, where SCRs had little informational content, there was no significant effect of SCRs on sale probabilities. We concluded that salience was not driving our results.

We are left to explain, however, why the probability of sale for cars without a posted SCR dropped in weeks 31-39. Recall that the experiment was conducted between May and September 2008, a period of significant decline in the stock market and the housing market. Arguably, these events may have affected sales probabilities. We therefore tested 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. To estimate a secular trend, we needed data that were not part of the market in which the experiment was conducted, so we chose cars that were offered for sale by fleet-sellers. For these cars there was no change in available information due to the experiment.[23] As we show in Section 8 of the online appendix, we found no evidence that the emails sent out starting in week 31 led dealers to substitute cars without posted SCRs with 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.

7 Concluding Remarks

It is well established that disclosing vertically differentiated information can help market participants better evaluate the value of goods and services in which they are interested, often resulting in more efficient outcomes and less distortionary information rents. For example, Lewis (2010) showed that by voluntarily disclosing information on ebay Motors, sellers could effectively offer protection to buyers from adverse selection. This insight helps explain the prevalence of many online transactions that otherwise seem puzzling due to concerns over potential “lemons.”

We have demonstrated that in addition to mitigating adverse selection, information disclosure can play an important role by helping buyers choose *which market* to participate in.[24] This simple, yet novel insight has broader applications beyond our market for used automobiles. If heterogeneous participants

[23] More than 98 percent of fleet-seller-consigned cars receive some type of inspection by the auction house. 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.

[24] Johnson and Myatt (2006) show that information disclosed through advertising can cause consumers to better learn their true value for the product. As a result, some may be discouraged from purchasing it while others will choose to buy it. In their setup this causes a change in the dispersion of valuations, and a rotation of the demand curve. In our setup, all buyers are better matched, causing demand curves to effectively shift out for every auction.

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.

In ebay's huge marketplace, sellers who reveal more information give buyers the opportunity to self-select into the auctions in which they are most interested, even if the information disclosed is considered bad news. We were able to obtain further insights from ebay auction data for laptops, cameras, and cell phones during 2011. We identify bad news for products that have descriptions with words like "broken," "not working," "cracked," etc. The data show that items with bad news are more likely to sell than items without bad news. We conjecture that our approach can explain this easily: certain buyers are looking for functional equipment, and have no value for a broken item. Others, like repair shops, have some positive value for broken items. At the same time, these repair shops have less value for fully functioning products than end users. Here again, one man's trash is another man's treasure.

Online auction sites such as ebay also offer exclusive, simultaneous auctions for vertically differentiated goods. Yet other markets have different institutional details that present matching benefits from disclosing vertically ranked information. For example, a firm looking to hire people for similar yet distinct positions may gain from providing information about generally undesirable aspects of a job. This will attract potential employees who have a lower marginal aversion to these aspects, providing a better match that can result in a higher likelihood of filling the position, and more importantly, a higher likelihood of retaining the employee.

Our insights may also apply to government procurement. Typically, governments engage in both parallel and sequential procurement of many similar yet distinct projects. For example, several construction, road, or defense acquisition projects may be let out to bid simultaneously, and yet many more are anticipated to materialize within weeks or months. Contractors with capacity constraints may not be able to bid on later auctions if they win earlier ones. If the procurement authority releases information about unattractive aspects of a job not only on current but also on future tenders, heterogeneous contractors may be able to better select in which of the coming auctions to participate.[25] This can increase effective competition both within and between projects, a desirable outcome of information disclosure as a matching mechanism.

[25] Budish and Zeithammer (2008) have shown that providing information for sequential auctions can increase expected revenues by having bidders use early auctions to learn about their competitors' willingness to pay, and then decide whether and how to bid in future auctions.

References

Anand, Bharat N. andRon Shachar (2011) "Advertising, the Matchmaker," *RAND Journal of Economics*, **42(2)**:205-245

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

Budish, Eric and Robert Zeithammer (2014) "An Efficiency Ranking of Markets Aggregated from Single-Object Auction," mimeo.

Cho, Sung-Jin, Harry J. Paarsch and John Rust (2014) "Is the 'Linkage Principle' Valid?: Evidence from the Field," *The Journal of Industrial Economics,* **62:** 346-375.

George S. Day (1976) "Assessing the Effects of Information Disclosure Requirements," *Journal of Marketing,* **40(2)**:42-52.

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.

Diamond, Douglas W. and Robert E. Verrecchia (1991) "Disclosure, Liquidity, and the Cost of Capital," *The Journal of Finance*, **46(4)**:1325-1359

Genesove, David (1993) "Adverse Selection in the Wholesale Used Car Market," *Journal of Political Economy* **101(4)**:644-665.

Gilson Ronald J. and Reinier H. Kraakman (1984) "The Mechanisms of Market Efficiency," *Virginia Law Review,* **70(4)**:549-644.

Grossman, Sanford (1981) "The informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics,* **24(2)**:461-489.

Healy, Paul M. and Krishna G. Palepu (2001) "Information asymmetry, corporate disclosure, and the capital markets: A review of the empirical disclosure literature," *Journal of Accounting and Economics,* **31(1-3)**:405?440

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.

Johnson, Justin P., and David P. Myatt. 2006. "On the Simple Economics of Advertising, Marketing, and Product Design." *American Economic Review,* **96(3)**:756-784.

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.

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

Milgrom, Paul R. and Robert J. Weber (1982) "A Theory of Auctions and Competitive Bidding," *Econometrica* **50 (5)**:1089–1122.

Overby, Eric and Sandy Jap (2009) "Electronic and Physical Market Channels: A Multi-Year Investigation in a Market for Products of Uncertain Quality," *Management Science*, **55(6)**:940-957.

Porter, Robert H. (1995) "The Role of Information in US Offshore Oil and Gas Lease Auctions," *Econometrica* **63(1)**:1–27.

Roberts, James W. and Andrew Sweeting (2014) "Bailouts and the Preservation of Competition," mimeo.

ONLINE APPENDIX

Table A-1: Dealer-consigned and inspected cars by week[†]

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

Weeks are of 2008.

Table A-2: 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 score	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.

Figure A-1: Sales probability by week

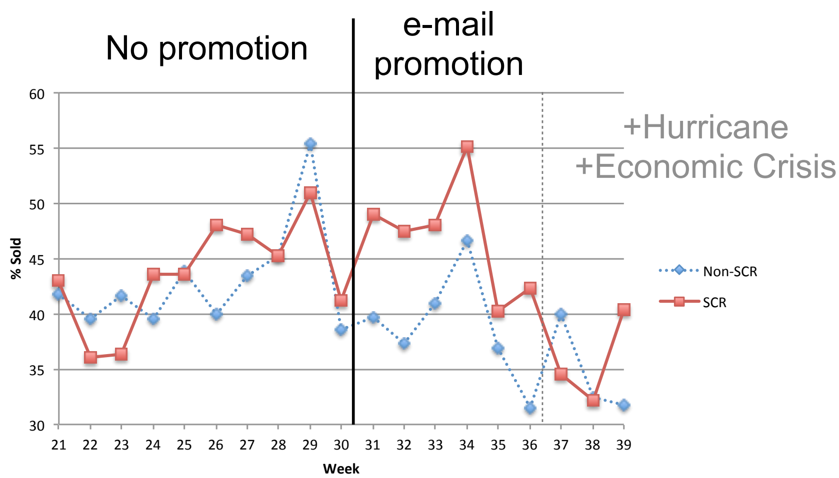

Robustness Checks

1 Expectations about Repair Cost Estimates

To obtain terciles of "difference from expected repair cost," we first estimated the predicted repair cost estimate of each car in our sample based on the vehicle age and vehicle mileage. We made this prediction by regressing estimated repair cost on vehicle age dummies, vehicle mileage, and vehicle mileage deciles, and an interaction between the vehicle age dummies with vehicle mileage. We took the difference between the actual estimated repair cost and the predicted estimates repair cost to construct a distance measure from the expected condition score. Finally, we split this distance measure into terciles, where the bottom tercile contains cars with worse-than-expected estimates repair cost, the middle tercile contains cars with close-to-expected estimates repair cost, and the top tercile contains cars with better-than-expected estimates repair cost.

During weeks 31-39 there was no statistically significant effect of a posted SCR on the probability of sale for cars in the middle tercile, where actual estimated repair costs are close to expected estimated

Table A-3: Sales probability by difference of expected estimated repair cost (ERC), weeks 31-39

Tercile of Difference from Expected ERC	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	z-statistic	p-value
Worse-than-expected	901	0.343	0.398	0.055	16.0%	1.69	0.09
Close-to-expected	898	0.429	0.442	0.012	2.8%	0.37	0.71
Better-than-expected	897	0.408	0.522	0.11	27.0%	3.45	0.001

repair costs. However, in both terciles where estimated repair costs have informational content, the effect on the probability of sale was positive. The effect was strongly significant for the better-than-expected tercile (p-value < 0.01) and marginally significant for the worse-than-expected tercile (p-value $= 0.09$).

The price results were also similar to the condition score findings in that there was no statistically significant effect of a posted SCR on the auction price of cars for any of the "difference from expected repair cost" terciles.

Table A-4: Price/NAP by difference of expected estimated repair cost (ERC), weeks 31-39

Tercile of Difference from Expected CS	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	t-statistic	p-value
Worse-than-expected	322	0.992	1.01	0.018	1.8%	0.76	0.45
Close-to-expected	386	1.03	1.04	0.012	1.12%	0.67	0.50
Better-than-expected	413	1.08	1.1	0.018	1.2%	0.84	0.40

Finally, as we found in the condition score results, when we use terciles of "difference from expected repair cost," we found no effect of information disclosure on probabilily of sale or auction prices during weeks 21-30 (not reported).

2 Standard Error Correction and Randomization Check

As described in Section 6, we estimated linear probability models with robust clustered standard errors at the VIN level. This accommodates arbitrary correlation between the errors of observations of the same car. We also added a large set of controls, namely seller fixed effects (267), model year fixed effects (13), vehicle segment fixed effects (21), nameplate fixed effects (38), sale week fixed effects (9), condition score tercile (3), and some (non-SCR) measures that represented the car's condition, namely its mileage and whether it was offered under a green, yellow, or red light, as well as a blue light.[26]

First, consider the aggregate finding that during weeks 31-39, cars with posted SCRs had a significantly higher probability of sale than cars without a posted SCR (second row of Table 3). Column 1 of Table A-5 shows that clustered standard errors don't change this inference.

Column 2 contains the treatment effect on the probability of sale controlling for the large set of

[26]The seller of every car sold at the 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 listed (e.g., "rough engine"). A red light means that the seller sells the car "as is," with no assurance of 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 was not at the auction site.

Table A-5: 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** (.019)	.046* (.021)	.02 (.012)	.0097 (.013)
CS close to expected		.035 (.028)		.08** (.018)
CS better than expected		.11** (.033)		.085** (.021)
Mileage on Car		6.8e-07 (4.4e-07)		3.1e-07 (3.3e-07)
Green light		.087+ (.047)		.17** (.046)
Yellow light		-.039 (.033)		-.033 (.027)
Blue light		-.12+ (.069)		-.0093 (.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	0.433

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

controls we listed above. The point estimate of the treatment effect drops from 6.3 percentage points to 4.6 percentage points. However, we can't reject the hypothesis that the treatment effect was unchanged by the inclusion of the extensive set of fixed effects. Columns 3 and 4 of Table A-5 show that our inference about the effect of a posted SCR on prices during weeks 31-39 (second row of Table 4) also remains unchanged. Clustering standard errors and controlling for fixed effects did not alter our conclusion that average prices seem not to have significantly increased due to SCRs.

We repeated these tests for the results that that were decomposed by the difference from the expected conditions score (Tables 7 and 9). Columns 1 and 2 of Table A-6 contain the effect of posted SCRs on the probability of sale by the difference from the expected conditions score. The relevant comparisons to the effects listed in Table 7 under the "Difference column" are the first three coefficients in the table. Clustering standard errors (column 1) did not change our inference. Adding controls changed the estimated coefficients very little. Similarly, columns 3 and 4 of Table A-6 didn't affect the interpretation of our price results.

In summary, the conclusions of the the paper's key specifications were unaffected by clustering standard errors and by adding a large set of controls—there was no evidence that our procedure yielded a non-random assignment to treatment and control groups.

3 Online Transactions

In this section we explore our argument that SCRs had no effect on auction outcomes during weeks 21-30 because dealers were not aware that they had been posted. We confirmed this by exploring the behavior of dealers who made use of the auction house's online bidding feature ("online dealers" for

Table A-6: Randomization check on results by expected CS: Sales probability and Transaction Prices for weeks 31-39

	Sales Probability		**Transaction Prices**	
	Base Result	Fixed Effects	Base Result	Fixed Effects
Posted SCR * CS Worse than expected	.084* (.036)	.084* (.036)	.022 (.021)	.029 (.02)
Posted SCR * CS Close to expected	-.011 (.036)	.00045 (.036)	.035 (.022)	.019 (.022)
Posted SCR * CS Better than expected	.11** (.038)	.11** (.037)	.0065 (.02)	.0022 (.018)
CS Close to expected	.1** (.036)	.1** (.036)	.067** (.021)	.061** (.021)
CS Better than expected	.092* (.036)	.11** (.037)	.094** (.02)	.1** (.02)
Mileage on Car		-5.3e-07 (4.2e-07)		1.9e-07 (2.9e-07)
Green light		.11* (.042)		.19** (.036)
Yellow light		-.033 (.031)		-.042+ (.022)
Blue light		-.11 (.067)		.012 (.037)
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
Constant	.33** (.025)	.25 (.21)	.98** (.015)	.87** (.13)
Observations	2696	2696	1121	1121
R-squared	0.015	0.076	0.034	0.218

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

short). These dealers must have been aware that SCRs were posted during weeks 21-30.

We considered three measures of online behavior as a function of whether an SCR was posted or not: (1) the percentage of vehicles that received an online bid, (2) the average number of online bidders, and (3) the percentage of sold vehicles bought by an online bidder.

As Table A-7 shows, over all weeks (21-39), 3.45 percent of cars with a posted SCR received an online bid, compared to 2.54 percent without a posted SCR. This 36 percent difference 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 driven by dealer behavior in weeks 31-39. We found that the effect of an SCR on the percentage of vehicles that received an online bid was statistically no different in weeks 21-30 as it was in weeks 31-39 (p-value 0.75).

Next, as Table A-8 shows, we found that over all weeks, 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 measure, the SCR effect was statistically no different in weeks 21-30 compared to weeks 31-39 (p-value 0.71).

Finally, the winning bids of 4.7 percent of cars with a posted SCR were placed online, compared to 3.07 percent without a posted SCR. This 53 percent difference is statistically significant (p-value 0.01). The effect of an SCR on the percentage of vehicles with an online winning bid was statistically

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

	No posted SCR	Posted SCR	Difference	% Difference	z-statistic	p-value
All weeks	2.54% 3,980 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

Table A-8: 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

no different in weeks 21-30 from weeks 31-39 (p-value 0.44). It is worth noting that it is ambiguous whether or not online bidders should win more or less auctions after the existence of SCRs was made known to the lane-bidders. Online bidders lose their information advantage in weeks 31-39 relative to lane-bidders. Hence, this may decrease the likelihood that a winning bid is placed online. Regrettably, there are few winning bids placed online and we cannot test for this effect. Interestingly, however, Table A-9 shows point estimates that suggest that online bidders were more successful in winning cars with a posted SCR during weeks 31-39 (they won 5.15% of auctions) than they were during weeks 21-30 (they won 4.51%), contrary to this conjecture. The difference is not statistically significant (not reported.) The ranking of these point estimates, however, is not implausible. Recall that because almost half the cars did not historically sell, the better targeting may have led to an increase in lane sales without decreasing online sales because online bidders, if they lose an auction, can go to the next car they want to buy before their demand is met. One can argue that this implies a clear welfare improvement, but given the size of the sample we cannot confirm this point.

Table A-9: 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

Given that online dealers knew about SCRs from the beginning of the experiment (week 21), and given that the effect of a posted SCR on their behavior was similar between weeks 21-30 and 31-39, we concluded that the effect of SCRs we observed offline during weeks 31-39 was most likely tied to dealers learning about SCRs.

4 Definition of "Worse-than-Expected," "Close-to-Expected," and "Better-than-Expected."

Recall that we constructed an expected condition score using a regression of condition score on vehicle age and vehicle mileage. We then used the residuals of that regression to classify cars into terciles. In this section we addressed the concern that the variance of the residuals may have been different for different types of cars. If this were the case, by placing cars into terciles relative to all other cars, we may have categorized a car as worse-than-expected, when it was close-to-expected relative to other cars of the same type.

We looked at the variance of the residuals across different types of cars. We used the standard car categories "Compact Car," "Fullsize Car," "Luxury Car," "Midsize Car," "Pickup," "Sports Car," "SUV," and "Van." Table A-10 shows the 10th and 90th percentile and the variance of the residuals by car category.

Table A-10: Percentiles and variance of difference between CS and expected CS by car category

Car Types	10th	90th	Variance
Compact Car	-1.86	1.10	1.34
Fullsize Car	-1.26	1.57	1.18
Luxury Car	-1.12	1.74	1.25
Midsize Car	-1.64	1.10	1.15
Pickup	-1.67	1.38	1.30
Sports Car	-1.52	1.59	1.58
SUV	-1.23	1.75	1.27
Van	-1.53	1.25	1.27

There were clear differences between car types. Compare, for example, the 10th and 90th percentile of midsize and luxury cars. Also, a test of equality of variances rejected the hypothesis that the variances of all of these groups were the same (not reported). The consequence is that by placing cars into terciles relative to all other cars, we may indeed have categorized a car as "worse-than-expected," when it was not "worse-than-expected" relative to other cars of the same category.

To address this concern we categorized cars into terciles only in comparison to cars of the same category. For example, we placed a Honda Civic into the "worse-than-expected" tercile if its residual was low relative to other compact cars. We then re-estimated Tables 7 and 9 with the car category specific terciles. The results are in Tables A-11 and A-12.

Table A-11: Sales probability by difference of expected condition score (CS) by car category, weeks 31-39

Car Category Specific Tercile	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	z-statistic	p-value
Worse-than-expected	902	0.347	0.427	0.080	23.1%	2.48	0.013
Close-to-expected	902	0.397	0.394	-0.003	-0.8%	0.09	0.927
Better-than-expected	892	0.433	0.537	0.104	24.0%	3.1	0.002

Table A-12: Price/NAP by difference of expected condition score (CS) by car category, weeks 31-39

Car Category Specific Tercile	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	t-statistic	p-value
Worse-than-expected	341	.98	1.00	0.02	2.4%	-1.15	0.25
Close-to-expected	352	1.05	1.08	0.029	2.8%	1.31	0.19
Better-than-expected	428	1.07	1.08	0.014	1.3%	0.65	0.51

Clearly, there were no meaningful differences between the two sets of tables. In addition, during weeks 21-30 there were no significant differences between cars with and without SCRs (not reported). This mirrors the finding in Tables 8 and 10.

5 Quintiles of Difference between Actual and Expected CS

To investigate whether using (coarse) terciles to categorize news' masked findings for cars that were "much better-than-expected" or "much worse-than-expected," we replicated our main analysis using quintiles instead of terciles. For an easy comparison with section 4 in this online appendix, we followed the procedure of categorizing cars relative to other cars in their category (The conclusion were not different if we categorize cars relative to all other cars). The results are in tables A-13 and A-14.

Table A-13: Sales probability by difference of expected condition score (CS) by car category, weeks 31-39

Car Category Specific Quintiles	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	z-statistic	p-value
Much worse-than-expected	540	0.343	0.397	0.054	16%	1.29	0.20
Worse-than-expected	540	0.362	0.413	0.051	14%	1.21	0.23
Close-to-expected	538	0.423	0.438	0.015	3%	-0.34	0.73
Better-than-expected	539	0.417	0.441	0.025	6%	-0.58	0.56
Much better-than-expected	539	0.415	0.588	0.173	42%	-4.01	0.0001

Table A-14: Price/NAP by difference of expected condition score (CS) by car category, weeks 31-39

Car Category Specific Quintiles	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	t-statistic	p-value
Much worse-than-expected	198	0.930	0.966	0.036	3.8%	1.49	0.14
Worse-than-expected	201	1.019	1.065	0.046	4.5%	1.52	0.13
Close-to-expected	227	1.063	1.078	0.015	1.5%	0.55	0.58
Better-than-expected	229	1.065	1.071	0.005	0.5%	0.22	0.82
Much better-than-expected	266	1.071	1.087	0.017	1.6%	0.59	0.56

Our interpretation of these tables was that the results were largely consistent with our earlier findings preserving the U-shape of the impact of news on the probability of sale (see Table A-11 in this appendix, and Table 7 in the paper). However, the smaller cell sizes made some of the differences statistically insignificant. We fond one new insight, namely that the sales probability effect of the "Better-than-expected" tercile, seems to be driven by the "Much better-than-expected" quintile.

6 Expectations over Estimated Repair Costs

We re-estimate Tables 7 and 9 in the paper, using the surprise in estimated repair costs instead of condition scores. First, we predicted estimated repair costs using the same specification we used to predict condition scores. Next, we took the difference between predicted and actual estimated repair costs to form terciles, analogous to our procedure in section 5.3. The results are reported in tables A-15 and A-16.

Table A-15: Sales probability by difference of expected repair cost (RC), weeks 31-39

Tercile of Difference from Expected RC	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	z-statistic	p-value
Worse-than-expected	901	0.343	0.398	0.055	16%	-1.69	0.09
Close-to-expected	898	0.429	0.442	0.013	3%	0.37	0.71
Better-than-expected	897	0.408	0.522	0.114	28%	3.45	0.001

Table A-16: Price/NAP by difference of expected repair costs (RC), weeks 31-39

Tercile of Difference from Expected RC	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	t-statistic	p-value
Worse-than-expected	322	0.992	1.01	0.018	2%	0.76	0.45
Close-to-expected	386	1.03	1.04	0.01	1%	0.67	0.50
Better-than-expected	413	1.08	1.1	0.02	2%	0.84	0.40

Our interpretation of these tables was that the results were largely consistent with our earlier findings, preserving the U-shape of the impact of news on sales probability and showing no significant effect on prices. The only change was that the effect of an SCR for cars in the "worse-than-expected" estimated repair cost tercile became marginally significant (p-value 0.09).

7 Stock Dynamics in the Formation of Expectations

The current definitions of "worse-than-expected," and "better-than-expected" do not account for stock dynamics. However, we can easily test whether our results continue to hold if we account for changing beliefs over time. Instead of using the whole set of cars over the entire period to estimate a predictor of condition scores, we can predict condition scores using only cars that were consigned recently. Specifically, for any car sold in a focal week we used only cars that were consigned during that week and the prior week to form a prediction. We then define the "worse-than-expected," "close-to-expected," and "better-than-expected" terciles based on this rolling-prediction and relative to cars in the focal week and the prior week. We then re-estimated Tables 7 and 9, accounting for this form of stock dynamics. The results are reported in tables A-17 and A-18.

All of our results were preserved accounting for this form of stock dynamics. The results were also robust to extending the rolling window from 2 to 3, 4, or 8 weeks (not reported).

8 Salience and Substitution

In this section we explain why the probability of sale for cars without a posted SCR dropped in weeks 31-39. As we have explained in Section 6, we used fleet-seller-consigned cars to estimate a secular trend

Table A-17: Sales probability by difference of expected condition score (CS) using rolling prediction, weeks 31-39

"Rolling" tercile	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	z-statistic	p-value
Worse-than-expected	916	0.351	0.416	0.065	19%	2.04	0.041
Close-to-expected	910	0.413	0.444	0.031	8%	0.93	0.354
Better-than-expected	870	0.412	0.504	0.092	22%	2.72	0.006

Table A-18: Price/NAP by difference of expected condition score (CS) using rolling prediction, weeks 31-39

"Rolling" tercile	# of Cars	No posted SCR	Posted SCR	Difference	% Difference	t-statistic	p-value
Worse-than-expected	343	0.967	0.996	0.029	3%	1.4298	0.154
Close-to-expected	384	1.06	1.09	0.03	3%	1.3568	0.176
Better-than-expected	394	1.07	1.08	0.01	1%	0.2682	0.789

in probability of sale and auction prices over the sample period. The probability of sale for fleet-seller-consigned cars was 67.25 percent in weeks 21-30 (13,491 cars) and 59.83 percent in weeks 31-39 (12,864 cars), a drop of more than 7 percentage points. This suggests that demand for cars at the auction site decreased over the period of the experiment. Adding fleet-seller-consigned cars to our sample allowed us to use a difference-in-differences linear probability regression to estimate the change over time in the probability of sale for cars with and without a posted SCR relative to fleet-seller-consigned cars.[27]

The results are in column 1 of Table A-19. The constant is the probability of sale for fleet-seller-consigned cars during weeks 21-30. The coefficient of "Week 31-39" is the change in the probability of sale for fleet-seller-consigned cars relative to weeks 21-30 and measures 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 was offered for auction more than once during the sample period, we cluster the standard errors at the VIN level.

The coefficient on "Week 31-39 * Dealer-consigned car, no posted SCR" is 0.031 (p-value 0.19). We cannot therefore 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" is 0.089 and is significantly different from 0 (p-value < 0.01). We made the following interpretation: under the maintained assumption that the demand conditions of fleet-seller-consigned cars changed similarly to the those for dealer-consigned cars, we found no evidence that the emails sent out starting in week 31 led dealers to substitute cars without posted SCRs with cars with posted SCRs. Instead, it seems that the probability of sale for cars without posted SCRs was unchanged (relative to fleet-seller-consigned

[27]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 was 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 (nineteen weeks), we used a linear probability model that estimated a linear time trend in the probability of sale for cars, separately for fleet-seller- and dealer-consigned cars. The results are in Table A-20. We cannot reject the hypothesis that the secular trend in probability of sale was the same for fleet-seller- and dealer-consigned cars.

Table A-19: 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.

cars), while the probability of sale for cars with posted SCRs increased.

A concern may be that the types of cars sold by fleet-sellers were not comparable to cars sold by dealers, making fleet-seller-consigned cars unsuitable for estimating the secular trend. We can (partially) address this concern by re-estimating the specification in column 1 of Table A-19 with model-year, vehicle segment, nameplate and sale-week fixed effects, and some (non-SCR) measures that represented the car's condition, namely, mileage and whether it was offered under a green, yellow, or red light and a blue light. This identifies 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 A-19, the estimates changed very little.

A remaining concern may be that there was 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. This category comprised 616 dealers, a large fraction of the 1,670 dealers who purchased at least one car (fleet-seller- or dealer-consigned) during our experimental period. If there was substitution between fleet-seller- and dealer-consigned cars with a posted SCR, we should find that these 616 dealers—if they purchased *any* dealer consigned cars during weeks 31-39—were more likely to buy cars with a posted SCR than without a posted SCR. We found no evidence of such behavior: dealers who purchased only fleet-seller-consigned cars during weeks 21-30 purchased forty-eight dealer-consigned cars with a posted SCR and fifty-three dealer-consigned cars without a posted SCR during weeks 31-39. We concluded that substitution was unlikely to explain why SCRs increased expected auction revenues.

Table A-20: Pre-treatment 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.