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

ONLINE APPENDIX

Steven Tadelis
UC Berkeley and NBER

Florian Zettelmeyer
Northwestern University and NBER

December 4, 2014

Abstract

The following Appendix provides additional details and estimations referred to in the published paper "Information Disclosure as a Matching Mechanism: Theory and Evidence from a Field Experiment," published in the *American Economic Review*

1 Additional Tables

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

Sale Week	Dealer-Consigned	With SCR	
	Total	Not reported	Reported
21	1,442	237	222
22	1,709	195	186
23	1,438	324	330
24	1,606	281	365
25	1,249	303	344
26	1,408	230	250
27	1,170	291	305
28	1,462	246	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	3,980	4,118

Weeks are of 2008.

Table A-2: Summary Statistics

Variable	N	mean	p50	sd	min	max
Model Year	8098	2004	2004	2.65	1997	2009
Mileage	8098	75959	71316	44359	0	508112
Condition score	8098	2.79	2.8	1.37	1	5
Repair Costs	8098	1348	1024	1237	0	16111
Sold	8098	0.43	0	0.495	0	1
Sales Price	3429	8606	7300	5862	500	59000
National Auction Price	3429	8397	6975	5811	200	62000
Sales Price/National Auction Price	3429	1.06	1.03	0.243	0.236	5.6

Figure A-1: Sales probability by week

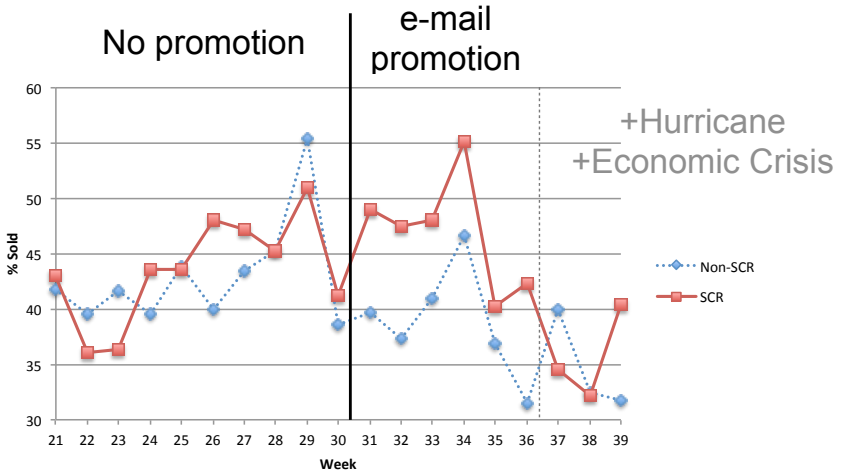

2 Robustness Checks

2.1 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.

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

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 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).

2.2 Standard Error Correction and Randomization Check

As described in Section V of the paper, 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.[1]

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 in the paper). Column 1 of Table A-5 shows that clustered standard errors don't change this inference.

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**	.044*	.02	.0078
	(.021)	(.021)	(.012)	(.013)
CS close to expected		.04		.049**
		(.027)		(.017)
CS better than expected		.095**		.067**
		(.029)		(.016)
Mileage on Car		4.6e-07		8.4e-08
		(4.4e-07)		(3.2e-07)
Green light		.089+		.16**
		(.047)		(.046)
Yellow light		-.04		-.031
		(.033)		(.028)
Blue light		-.13+		-.0082
		(.069)		(.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	0.272	0.002	0.426

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

Column 2 contains the treatment effect on the probability of sale controlling for the large set of controls we listed above. The point estimate of the treatment effect drops from 6.3 percentage points to 4.4 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 in the paper) 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.

[1]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.

We repeated these tests for the results that that were decomposed by the difference from the expected conditions score (Tables 7 and 9 in the paper). 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 in the paper 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.

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*	.084*	.022	.029
	(.036)	(.036)	(.021)	(.02)
Posted SCR * CS Close to expected	-.011	.00045	.035	.019
	(.036)	(.036)	(.022)	(.022)
Posted SCR * CS Better than expected	.11**	.11**	.0065	.0022
	(.038)	(.037)	(.02)	(.018)
CS Close to expected	.1**	.1**	.067**	.061**
	(.036)	(.036)	(.021)	(.021)
CS Better than expected	.092*	.11**	.094**	.1**
	(.036)	(.037)	(.02)	(.02)
Mileage on Car		-5.3e-07		1.9e-07
		(4.2e-07)		(2.9e-07)
Green light		.11*		.19**
		(.042)		(.036)
Yellow light		-.033		-.042+
		(.031)		(.022)
Blue light		-.11		.012
		(.067)		(.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**	.25	.98**	.87**
	(.025)	(.21)	(.015)	(.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 .

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.

2.3 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 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).

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%	3.45 %	0.91%	35.8%	2.40	0.016
	3,980 cars	4,118 cars				
Weeks 21-30	2.69%	3.50%	0.81%	30.2%	1.73	0.084
	2,605 cars	2,797 cars				
Weeks 31-39	2.25%	3.33%	1.08%	47.7%	1.70	0.089
	1,375 cars	1,321 cars				

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).

Table A-8: Expected number of online bidders per 100 auctions

	No posted SCR	Posted SCR	Difference	% Difference	t-statistic	p-value
All weeks	3.64	4.74	1.09	29.8%	2.22	0.026
	3,980 cars	4,118 cars				
Weeks 21-30	3.76	4.72	0.96	25.5%	1.58	0.11
	2,605 cars	2,797 cars				
Weeks 31-39	3.42	4.77	1.35	39.5%	1.60	0.11
	1,375 cars	1,321 cars				

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 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	1.3%	40.5%	1.62	0.10
Weeks 31-39	2.78% 539 cars	5.16% 601 cars	2.38%	85.6%	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.

2.4 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

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 7 and 9.

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

2.5 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 2.4 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.

2.6 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 V of the paper, we used fleet-seller-consigned cars to estimate a secular trend in probability of sale and auction prices over the sample period. The probability of sale for fleet-seller-consigned cars was 66.63 percent in weeks 21-30 (13,942 cars) and 59.65 percent in weeks 31-39 (13,247 cars), a drop of 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.[2]

The results are in column 1 of Table A-15. 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 cars), while the probability of sale for cars with posted SCRs increased.

[2]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-16. 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-15: 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.

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-15 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-15, 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 600 dealers, a large fraction of the 1,650 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 600 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-16: 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.