Asymmetric Information, Adverse Selection and Online Disclosure: The Case of eBay Motors[*]

Gregory Lewis [†]

January 28, 2011

Abstract

Since George A. Akerlof (1970), economists have understood the adverse selection problem that information asymmetries can create in used goods markets. The remarkable growth in online used goods auctions thus poses a puzzle. Part of the solution is that sellers voluntarily disclose their private information on the auction webpage. This defines a precise contract — to deliver the car shown for the closing price — which helps protect the buyer from adverse selection. I test this theory using data from eBay Motors, finding that online disclosures are important price determinants; and that disclosure costs impact both the level of disclosure and prices.

[†]Department of Economics, Harvard University, 125 Littauer Center, 1805 Cambridge St, Cambridge MA 02138. Email: glewis@fas.harvard.edu.

1 Introduction

The rise of the internet has seen a huge rise in the volume of used goods traded online. Online auction sites such as eBay and Yahoo! Auctions compete worldwide with specialized listing sites such as usedcomputer.com and cars.com in the retail trade of consumer goods. Meanwhile, business to business transactions totaling billions of dollars take place through online auctions in industries as diverse as aviation and mining. At first glance, this growth is somewhat surprising. Since Akerlof's classic paper, economists have been aware of the potential for adverse selection in markets with information asymmetries, such as used good markets. Information asymmetries are exacerbated in online transactions, where the buyer typically does not view the good in person. Why then has the volume of trade in these markets proved so robust to adverse selection?

In this paper, I argue that it is fundamentally because sellers are able to partially contract on the quality of their goods. By disclosing their private information on the auction webpage in text and photos, the seller offers a contract to potential buyers to deliver the item described in the listing. If the disclosures define sufficiently detailed and enforceable contracts, the initial information asymmetry should play no role in determining the performance of the market. Both of these caveats are important: if contracts are not enforceable, they are meaningless; while if disclosures are costly, the resulting contracts may be coarse and market efficiency may suffer.

To test this hypothesis, I examine the role of disclosure and disclosure costs on eBay Motors, the largest used car marketplace in the United States. In this market, despite high stakes for both sides and substantial information asymmetries, there is a high volume of trade with approximately 36000 cars sold each month.[1] By analyzing how the market works, I first argue that the institutional framework makes certain kinds of claims relatively easy to enforce. Then, using a large new dataset of over 80000 car auctions, I show that photos and text posted by the seller on the auction webpage strongly influence prices, suggesting that online disclosures are important empirically. Finally, I show that disclosure costs affect how much information the seller decides to post, and therefore the prices they obtain. Taken together, I conclude that disclosure costs — whether caused by technology, bandwidth, or time costs — are an important determinant of the extent to which parties can create well-defined contracts online, and therefore of the success of online goods marketplaces.

The theoretical foundations of the paper lie in the work of Sanford J. Grossman & Oliver D. Hart (1980), Sanford J. Grossman (1981) and Paul Milgrom (1981), who argued that verifiable disclosure might mitigate the adverse selection problems of Akerlof (1970). Boyan Jovanovic (1982) investigated the welfare implications of disclosure costs. On the empirical side, there is a diverse literature on both mandatory

---

[1]Source: eBay Press Releases, (ReleaseID = 206868, 306677).

and voluntary disclosure (see e.g. Alan D. Mathios (2000), Ginger Zhe Jin & Phillip Leslie (2003), Ginger Zhe Jin (2005) Ginger Zhe Jin & Andrew Kato (2006)). There is also a literature on adverse selection in used vehicle markets (e.g. Eric W. Bond (1982), David Genesove (1993)). The contribution of the paper to the online auctions literature (e.g. Paul Resnick & Richard Zeckhauser (2005), Daniel Houser & John Wooders (2006)) is to provide a different perspective, shifting the focus from the seller feedback mechanism, and the implicit contract it helps enforce, to the role of disclosure and the explicit contracts thus defined. In a similar vein, Pai-Ling Yin (2008) shows that when the webpage leaves bidders uncertain as to the value of the object sold, prices are lower. The latter part of the paper examines how software availability changes what information is disclosed. George P. Baker & Thomas N. Hubbard (2004) also examine the impact of technology on contracting, looking at how the introduction of on-board computers impacted asset ownership in the trucking industry. The paper proceeds as follows. Section 2 describes the market; section 3 contains the empirical analysis; and section 4 concludes.

2   eBay Motors

eBay Motors, the automobile arm of online auctions giant eBay, is thriving. It is the largest automotive site on the Internet, and every month, approximately 36000 vehicles are sold, a rate just slightly slower than a car a minute. This trading volume dwarfs those of its online competitors, the classified services cars.com, autobytel.com and Autotrader.com. In contrast to these sites, most of the sellers on eBay Motors are private individuals, although dealers still account for around 30% of the listings. Another big difference is that a large proportion (75%) of vehicles are sold to out-of-state buyers. Because of this, bidders can typically neither examine the car in person nor rely entirely on the seller's reputation; they must rely on the information on the auction webpage to evaluate potential purchases.[2]

Some information is standardized and mandatory, such as car make, model, mileage etc. But most of the details are voluntarily disclosed by the seller in the item description, which can include text, photos, graphics and video. eBay charges little for posting information, as text and graphics are free, while each additional photo costs $0.15. Yet the opportunity costs are higher, as it is time-consuming to take, select and upload photos, write the description, generate graphics etc. While these opportunity costs may seem small, the fact that professional car dealers typically invest in advanced listing management software to limit these costs suggests that they are not insignificant. Such software allows easier photo uploading and maintenance, graphics production and listing management, and is offered by companies such as CARad, eBizAutos and Auction123 at costs ranging from $10 a listing to a flat $300 a month

---

[2]Source: Auction123 (http://www.auction123.com/ebayadvantages.aspx).

fee. It is typical in most eBay car auctions for sellers to post many photos, a full text description of the car's history and features, and sometimes graphics and videos showing the car's condition.

The webpage created by the seller defines the contract between the buyer and the seller, in accordance with whose terms the buyer agrees to purchase the vehicle described by the seller at the final closing price of the auction.[3] Rich media such as photos and videos may define the contract terms more precisely than text. Two features of the market make these contracts practically enforceable. First, "most buyers opt to pick up the vehicle in person."[4] Even when the seller ships the vehicle to the buyer, payment is often held in escrow (e.g. through PayPal) until the buyer has had a chance to examine the vehicle. The result is that much of the information provided by the seller is often verifiable before payment is made. Second, material misrepresentations by the seller constitute fraud. In contrast to private car sales offline where it may be difficult to establish exactly what the seller did or did not promise, the webpage is stored by eBay for at least 20 days after the sale. As a result, these online transactions have a clearly defined contract in the event of a dispute. Because of these institutional features, sellers have little incentive to lie, and buyers can take much of the webpage information at face value. Conversely, sellers have the incentive to create detailed webpages, knowing that buyers will rely on the information presented. We examine how buyers respond to disclosures below.

3   Analysis

To fix ideas, it is useful to have the disclosure model of Grossman & Hart (1980), Grossman (1981) and Milgrom (1981) in mind. A seller knows a number of pieces of information about the car he is selling. On some dimensions, the information known may not be verifiable ex-post or enforced as a contractual claim by the buyer. Then the buyer should not update based on seller statements about these dimensions. For example, the buyer may not be able to judge the mechanical condition of the car upon pick-up. In this case, statements along the lines of "this car has no mechanical problems whatsoever," should be treated as cheap talk.

On other dimensions, though, the information can be directly exhibited on the auction webpage, and verified ex-post. For example, the seller can post photos showing the condition of the car exterior. Since the seller has strong incentives not to misrepresent this information, buyers should update their priors based on the webpage content. In addition, the disclosure model tells us that buyers should be skeptical, and interpret the absence of information along these dimensions as a bad signal.

---

[3] *Caveat emptor* applies: it is the buyer's responsibility to ask questions about undisclosed details before bidding. Experienced sellers often explicitly include a boilerplate disclaimer of this form.

[4] Source: eBay Motors Seller's Guide, http://pages.motors.ebay.com/howto/selling/closeB.html

Disclosure costs play a role because they determine the marginal cost to a seller of posting a piece of information. The marginal benefit is endogenous and depends upon buyer expectations: the value of posting photos showing that the exterior has no dents is highest for older cars, because buyers expect older cars to be dented, and thus the change in their willingness to pay on seeing no-dent photos is higher.

In a market with infinite disclosure costs, no-one ever reveals information and the potential for adverse selection is high. At the other extreme, when disclosure is costless, the information asymmetry "unravels" on every dimension that can be ex-post verified, as every type has an incentive to reveal their information rather than get pooled with worse types. Adverse selection can thus occur only on the other dimensions. In the supplementary appendix, I show more generally that adverse selection is increasing in disclosure costs. Consequently, the level of disclosure costs has implications for the efficiency of the market.

For the remainder of the paper, I take this theory to the data. First I test whether bidder behavior is causally influenced by information on the auction webpage, as the theory would predict. Second, I look for a relationship between disclosure costs and the level of disclosure.

3.1   The Data

The main data source is a collection of auction webpages from completed used car auctions on eBay Motors. This data was obtained by downloading the auction web-pages for selected car models over an 8 month period (March-October 2006), and then implementing a pattern matching algorithm to pull variables of interest from the webpage html code. I drop observations with nonstandard or missing data; and those pertaining to new or certified pre-owned cars or cars under salvage title.[5] I also drop auctions in which the webpage was not created using either the basic eBay listing tools, or one of the most commonly used proprietary listing platforms, CARad, Auction123 or eBizAutos (11% of the remaining listings). The resulting dataset consists of 82538 observations of 18 models of vehicle. The models of vehicle are grouped into three main types: those which are high volume Japanese cars (e.g. Honda Accord, Toyota Corolla), a group of vintage and newer "muscle" cars (e.g. Corvette, Mustang), and most major models of pickup truck (e.g. Ford F-series, Dodge Ram). I call these groups "reliable", "classic" and "pickups" respectively. I also split out classic cars of model-year less than 1980, and call these "collectible".

Table 1 summarizes the variables in the dataset. For each auction, I observe a number of item characteristics including model, year, mileage and transmission and the number of options/accessories such as car radio etc listed by the seller. I also observe

---

[5]I drop cars under salvage title because they attract a completely different set of buyers, and are arguably in a different market.

whether the vehicle sold is currently under manufacturer warranty. As a measure of reputation, I have the seller's eBay feedback. All of this information is standardized and mandatory, in that the seller must provide it when listing the vehicle.

But my focus here is on the information *voluntarily* disclosed by the seller in the item description. I have two simple measures of this content. First, the number of photos posted on the auction webpage (my primary measure). Second, I have dummies for whether key text phrases — "rust", "scratch" and "dent" — are used in the item description, and modifiers for how they are used. For example, a negation is a phrase like "rust-free", or "never seen any rust".[6] As is clear from the summary statistics in Table 1 these webpages exhibit substantial variation in information content.

On average, the cars are old (nearly 16 years on average) and well-travelled (about 90000 miles on the odometer). This is because of the large fraction of collectible cars sold on eBay Motors. There are 17 photos on a typical webpage. Sellers are typically experienced, with average feedback scores of 148. The minimum bid is usually set well below the book value of the vehicle, and thus most (85%) auctions receive at least 1 bid, with the highest bid averaging just over $11 000. But only 28% of the cars actually sell, because of the widespread use of secret reserves.

In the last two columns, I distinguish between "dealers" and "non-dealers", where I define a dealer to be any seller who lists more than one vehicle on eBay. Dealers and non-dealers differ quite markedly. Dealers list newer cars (3.5 years newer), with lower mileage (17000 miles less) and these cars are more than twice as likely to be under warranty. They also behave quite differently, using professional listing software for 47.5% of listings, versus 3.7% for non-dealers; and they put up many more photos (21.4 versus 12.7). They use lower minimum bids, but higher secret reserves, so that average dealer sales rates are around 6% lower.

I supplement this main data source with data on private party book values publicly available at edmunds.com.[7] For model-years dated 1990 or later, I obtained the typical dealer retail value for each model-year of the models in my data set, and then matched this with each observation in the main data set, matching on trim where possible. This gives me book value data for nearly 55000 observations.

---

[6]I give more details on the construction of these variables, and choice of phrases in the supplementary appendix, where I detail the content analysis methodology.

[7]I used the "used car appraiser" at http://www.edmunds.com/tmv/used/index.html, which generates a book value estimate based on recent average dealer sales prices for that model-year, adjusted via a proprietary formula for factors like current vehicle inventory levels, economic trends and unpublished incentives.

Table 1: Summary Statistics

	Full Sample		Non-Dealers	Dealers
	Mean	Std. Dev[a]	Mean	Mean
Car Characteristics				
Miles	90181	90663	98320	81217
Age (in years)	15.8	13.6	17.5	14.0
% Manual Transmission	30.4	—	33.6	26.8
% Warranty	18.6	—	12.2	25.5
# of Options	5.2	5.2	5.4	5.0
# of Photos	17.0	10.8	12.7	21.4
% "Classic" Cars	51.0	—	54.6	47.0
% "Reliable" Cars	18.7	—	16.5	21.1
% Pickups	30.3	—	28.9	31.9
% "Collectible" Cars	19.1	—	19.4	18.6
% CARad	14.7	—	2.9	27.7
% Auction123	3.8	—	0.5	7.4
% eBizAutos	6.1	—	0.3	12.4
% "Rust" Phrase	19.3	—	21.7	16.7
% "Rust" Negation	6.8	—	6.8	6.8
% "Scratch" Phrase	16.4	—	13.3	19.8
% "Scratch" Negation	2.1	—	1.9	2.3
% "Dent" Phrase	12.1	—	12.6	11.5
% "Dent" Negation	3.2	—	3.0	3.3
Seller Characteristics				
Seller Feedback Score	148.0	556.7	115.0	184.5
% negative feedback	1.60	6.00	1.42	1.78
Auction Characteristics and Outcomes				
Minimum Bid (% of book value)[b]	52.4	78.5	62.6	43.0
% auctions with ≥ 1 bid	85.2	—	82.5	87.9
% sold	28.4	—	31.4	25.2
Highest Bid	11110	13018	9173	13113

[a]Standard deviations for categorical variables are not reported.
[b]Statistics calculated from the book value subsample.

3.2 Prices and Information

In the first part of the analysis, I examine the relationship between price and information measures such as photos and text. I run log-linear hedonic regressions of the following form:

$$\log p_t = x_t \beta + \varepsilon_t \tag{1}$$

where p_t is the price in auction t, x_t is a vector of item and webpage characteristics in auction t, and ε_t is an error term capturing the idiosyncratic taste of the winning bidder in this auction. For the moment, I assume that x_t and ε_t are uncorrelated; later I examine potential sources of endogeneity.

This equation is the workhorse of the empirical analysis. It is motivated by the idea that under the theory, the "quality" of the car portrayed in the webpage should be positively correlated with the bid of the second highest bidder (and thus the price). By "quality" I mean an index that captures the difference between how an average bidder perceives the value of this car, relative to an average car of the same base characteristics, after updating on the webpage content. For example, a 1960 Honda where the photos show only a few small dents might be perceived as "high-quality", whereas a 2007 Honda with the same dents may be "low-quality". In the first round of regressions, the only webpage characteristic included is the number of photos. This is a good proxy for car quality, since if the seller has a high quality vehicle, he should include many photos; but if not, he should put up very few.

I report the results of a wide variety of specifications in table 2. In the "base specification" (1), the vector of covariates includes car characteristics (mileage, number of options, model, year and transmission fixed effects), the number of photos and it squared, a fixed effect for the week of listing (to control for seasonal demand fluctuations), and a pair of seller characteristics (log feedback and percentage negative feedback). The coefficients generally have the expected sign and all are highly significant. Of particular interest is the sheer magnitude of the positive coefficients on the number of photos. A change from 9 to 10 photos is associated with a selling price that is approximately 1.54% higher, which for the average car in the dataset is around $171 more. To put this in context, the value of a used car of a given model-year and mileage can vary by thousands of dollars depending on factors such as vehicle condition, maintenance history and documentation, all of which can be shown in photos. What this result suggests then is that bidders do rely heavily on photos to form perceptions of quality, and that the market is operating as expected. Sellers of high quality cars contract to provide high quality cars by carefully describing them on the webpage; those selling low quality cars provide weakly specified contracts through minimally descriptive webpages, and duly receive lower bids.[8]

---

[8]Notice that the effects of negative feedback are quite small and the coefficient on total log feedback is actually negative. This may be because total feedback conflates transactions across product categories.

In specification (2) I interact photos with age and warranty status, expecting that photos have a greater impact on prices for older cars (due to greater heterogeneity) and a lower impact for cars under warranty (since the buyer is partially insured by the warranty). The sign is as expected in both cases, though only significantly so for the interaction with age. In the final four columns, I consider specific subsamples. Column (3) is non-dealers, column (4) is dealers, and column (5) is collectible cars. Photos are more strongly correlated with price for non-dealers, possibly because buyers cannot rely on reputation as an alternate source of information about quality. They are also particularly important for collectibles. Finally, in column (6) I look at the subsample of newer cars for which I have book value data, and include log book value as a control. As one would expect, the estimated relationship is weaker since these are newer cars with less underlying heterogeneity, but is still significant and large in magnitude.

Endogeneity: One might naturally be concerned that the correlation between price and photos is not causal, and is instead driven by an omitted variable or selection. In table 3 I examine how robust the relationship is. In column (1), I include the observed number of bidders as a control variable. The idea is to test if there is a partial correlation between photos and prices after controlling for observed participation. The estimated relationship remains strong and positive, which rules out a story in which prices are higher in auctions with many photos purely because of increased competition, and not because the photo content causes bidders to update their valuations. In column (2), I attempt to deal with the selection problem induced by only using auctions with non-zero bidders. I estimate a Tobit-like model in which the latent variable — the log intended bid of the bidder with the second highest valuation — is equal to the log price when observed, and censored below at the log minimum bid when there are no bids. The results are very similar to the baseline specification.

Finally, in column (3) I try to deal with the concern that the results are driven by seller heterogeneity. Frequent sellers like car dealerships may have lower disclosure costs and put up more photos. Then if buyers prefer to buy from professional car dealers, I may be picking up this preference rather than the effects of information disclosure. To analyze this, I restrict to the subsample of dealers and include a seller-specific fixed effect for each of them. The results show that even after controlling for seller identity, there is a large and significant relationship between price and the number of photos. This suggests that dealers vary the amount of photos for each individual listing (i.e. the information posted is vehicle specific), and that furthermore such information positively co-varies with prices. Such results are consistent with selective disclosure.

Text Analysis: Another measure of webpage content is the text of the car description. In table 4, I add dummies for the presence of certain phrases in the item description to the set of covariates in the base specification. The supplementary appendix describes in detail how these phrases were chosen and the variables constructed. For each noun (e.g. "rust"), I distinguish between no mention of the phrase

8

Table 2: Hedonic Regressions

	Log Price					
	(1)	(2)	(3)[a]	(4)[b]	(5)[c]	(6)[d]
Log Miles	-0.130	-0.127	-0.132	-0.123	-0.080	-0.183
	(0.005)[e]	(0.005)	(0.005)	(0.007)	(0.006)	(0.007)
Number of Photos	0.020	0.009	0.028	0.013	0.032	0.009
	(0.001)	(0.002)	(0.001)	(0.002)	(0.003)	(0.001)
Photos Squared / 100	-0.023	-0.016	-0.035	-0.013	-0.027	-0.009
	(0.002)	(0.003)	(0.003)	(0.003)	(0.007)	(0.002)
Number of Options	0.015	0.014	0.019	0.011	0.102	0.008
	(0.001)	(0.001)	(0.001)	(0.001)	(0.006)	(0.001)
Log Feedback	-0.009	-0.009	-0.011	-0.010	-0.001	-0.013
	(0.002)	(0.002)	(0.002)	(0.004)	(0.006)	(0.002)
% Negative Feedback	-0.004	-0.004	-0.003	-0.005	0.001	-0.003
	(0.001)	(0.001)	(0.001)	(0.001)	(0.003)	(0.001)
Age X Photo		0.001				
		(0.000)				
Warranty		0.059				
		(0.017)				
Warranty X Photo		-0.000				
		(0.001)				
Log Book Value						0.587
						(0.015)
Model/Year/Week FE	yes	yes	yes	yes	yes	yes
R^2	0.695	0.699	0.675	0.701	0.479	0.784
N	71292	71292	33232	38060	13688	47148

[a]Estimated on subsample of "private sellers" (list only a single vehicle in sample period).
[b]Estimated on subsample of "dealers" (list multiple vehicles in sample period) .
[c]Estimated on subsample of collectible vehicles (selected models with model-year ¡ 1980).
[d]Estimated on book value subsample.
[e]Standard errors are clustered by seller.

Table 3: Endogeneity

	Log Price		
	(1)	(2)^a	(3)^b
Log Miles	-0.131	-0.126	-0.088
	(0.005)^c	(0.004)	(0.006)
Number of Photos	0.019	0.022	0.019
	(0.001)	(0.001)	(0.002)
Photos Squared / 100	-0.023	-0.025	-0.014
	(0.002)	(0.002)	(0.004)
Number of Options	0.015	0.015	0.012
	(0.001)	(0.001)	(0.001)
Log Feedback	-0.011	-0.008	-0.015
	(0.002)	(0.002)	(0.014)
% Negative Feedback	-0.004	-0.004	0.001
	(0.001)	(0.001)	(0.003)
Model/Year/Week FE	yes	yes	yes
Seller Fixed Effects	no	no	yes
Number of Bidders FE	yes	no	no
N	71292	82538	38060

^{a}Tobit model: used to account for censoring of bids below the minimum bid. Full sample (including auctions with no bids) is used.

^{b}Estimated on dealer sub-sample, with seller fixed effects.

^{c}Standard errors are clustered by seller.

(the omitted group), an unqualified mention (e.g. "car has rust"), a negated mention (e.g. "car is rust free"), a positively qualified mention (e.g. "car has very little rust") and a negatively qualified mention (e.g. "car has a lot of rust"). The coefficients are consistent with buyers responding to the information presented: there are positive coefficients on the negated mention, and increasingly negative coefficients across positively qualified mentions ("small rust") through negatively qualified mentions ("big rust"). This is also true across a number of subsamples.

That said, the disclosure model doesn't directly match the data here. In theory, making no statement at all should be regarded as a very bad signal of quality; for if not, those with lemons would simply keep silent. Here, the group of cars with negative information sells on average for lower prices than those with no statement. Why then shouldn't owners of cars with rust simply choose not to reveal it? One explanation is that if the car is riddled with rust, then non-disclosure means both not disclosing in text *and* putting up few photos, and as already shown, cars with few photos obtain low prices. It is better to disclose. A different explanation is that the seller will struggle to enforce his purchase contract with the buyer if it appears he has deliberately omitted large and material details from the car description (e.g. large scratches, dents), and so, anticipating this, reveals it upfront to avoid costly ex-post renegotiation. Finally, sellers may have a reputation to preserve, or simply place value in behaving honestly. My sense is that all of these factors play a role.

3.3   Disclosure Costs and Disclosure

Previously, I argued that disclosure costs were theoretically important in determining market performance because they determine the extent of adverse selection. In this section, I look at how disclosure costs are related to the level of disclosure. To do this, I need a proxy for the latent disclosure costs. A natural candidate is the listing software used by the seller to create the webpage. In the data I have sellers who use the standard eBay software, and those who use the professional listing platforms provided by CARad, Auction123 and eBizAutos. These technologies promise users that they will simplify and streamline the process of creating a listing, through simple user interfaces, templates and free photo hosting and management services. It seems reasonable then that they should lower disclosure costs.

The downside with using this as a cost-shifter is that it is potentially correlated with seller unobservables. As shown by the summary statistics reported in Table 1, dealers are overwhelmingly more likely to use the professional platforms. There are a couple of reasons for this. First, there is a large initial fixed cost associated with setting up the templates properly (e.g. most dealerships include an "about us" part of the template, which private sellers would not bother with). Second, the platforms have a menu of prices, where one-off listings are relatively expensive ($10 for CARad, $15 for Auction123, not available for eBizAutos), but unlimited monthly listing plans

Table 4: Text Analysis

| | Log Price | | | |
	Full Sample	Private Seller	Dealer	Book Value
No scratch[a]	0.097	0.125	0.060	0.038
	$(0.020)^b$	(0.022)	(0.028)	(0.016)
Small scratch	0.028	0.055	0.005	0.011
	(0.015)	(0.017)	(0.021)	(0.012)
Scratch	0.007	0.031	-0.015	-0.016
	(0.016)	(0.014)	(0.023)	(0.012)
Big scratch	-0.018	0.013	-0.047	-0.054
	(0.018)	(0.024)	(0.026)	(0.016)
No dent	0.003	-0.010	0.017	0.028
	(0.017)	(0.021)	(0.023)	(0.015)
Small dent	-0.051	-0.103	-0.034	-0.062
	(0.032)	(0.052)	(0.036)	(0.030)
Dent	-0.110	-0.109	-0.110	-0.085
	(0.012)	(0.014)	(0.020)	(0.011)
Big dent	-0.150	-0.151	-0.152	-0.140
	(0.029)	(0.037)	(0.044)	(0.031)
No rust	0.080	0.064	0.091	0.009
	(0.014)	(0.016)	(0.020)	(0.018)
Small rust	-0.252	-0.211	-0.276	-0.158
	(0.022)	(0.024)	(0.044)	(0.028)
Rust	-0.275	-0.279	-0.257	-0.162
	(0.015)	(0.018)	(0.024)	(0.017)
Big rust	-0.457	-0.461	-0.431	-0.298
	(0.023)	(0.027)	(0.037)	(0.030)
Number of Photos	0.020	0.029	0.013	0.009
	(0.001)	(0.001)	(0.002)	(0.001)
Photos Squared	-0.024	-0.037	-0.014	-0.010
	(0.002)	(0.003)	(0.003)	(0.002)

[a]Dummies take the form "no x", meaning any negation; "small x", meaning any favorable qualifier; "x", meaning the phrase used without qualification; and "big x" implying an unfavorable qualifier.

[b]Standard errors are clustered by seller.

may be cost effective for high volume sellers (they range from \$200-\$300 a month). Fortunately, for dealers I have a panel of observations. So I can ask whether dealers who upgrade software tend to post more photos.

The results of regressing photos on characteristics and software are reported in the first column of table 5, under the "first-stage" column. They indicate that dealers who switch to professional listing software subsequently put up significantly more photos than those that don't, around 10 more. Those selling cars with more options, or lower mileage, also tend to put up more photos. Now, given that software is correlated with photos, it could potentially be used as an instrument in a regression of price on photos. It will be uncorrelated with the error term in equation (1) if there is no "marketing" effect whereby a better looking webpage induces higher bids for the same car, and if the software upgrade is not caused by selection on unobservable car quality. Neither of these are testable, but I find it a little reassuring that there is no significant difference in *observable* car characteristics pre and post upgrade (see supplementary appendix).

With these caveats in mind, the second column reports an IV regression of price on photos, with seller fixed effects and software as an instrument. The first-stage F test of the instruments is a respectable 34.2. In the main regression, the estimated coefficient on photos is significant and positive. Compared to the OLS results shown in column (3), the coefficient is smaller, as one would expect. This result suggests that disclosure costs have a causal effect on equilibrium prices, through affecting the level of disclosure.

4   Conclusion

Given the growth of online transactions in used goods markets, it is important to understand what makes these markets work. This paper shows that certain kinds of information asymmetries in these markets can be endogenously resolved, so that adverse selection need not occur. The required institutional features are a means for credible disclosure and institutions that allow for contractual enforcement. With these in place, sellers have both the opportunity and the incentives to remedy information asymmetries between themselves and potential buyers. Disclosure costs are important in determining how effective this remedy is. Where bandwidth and technology are available to tightly define the contract between buyer and seller through rich media such as photos and videos, adverse selection problems are mitigated.

13

Table 5: Cost and Equilibrium Outcomes

	First Stage[a]	IV[b]	OLS[c]
Log Miles	-0.197	-0.088	-0.087
	(0.046)[d]	(0.003)	(0.006)
Number of Options	0.076	0.012	0.012
	(0.012)	(0.001)	(0.001)
Log Feedback	0.658	-0.011	-0.013
	(0.206)	(0.012)	(0.014)
% Negative Feedback	-0.016	0.001	0.001
	(0.035)	(0.002)	(0.003)
Number of Photos		0.008	0.011
		(0.003)	(0.001)
CARad	9.130		
	(1.064)		
Auction123	12.886		
	(1.407)		
eBizAutos	10.392		
	(1.597)		
Model/Year/Week FE	yes	yes	yes
Seller FE	yes	yes	yes
First stage F	32.65	—	—
N	38060	38060	38060

[a]OLS Regression of photos on covariates, using dealer sub-sample.
[b]IV regression of log winning bid on covariates, with photos instrumented for with software.
[c]OLS regression of log winning bid on covariates.
[d]Robust standard errors are reported.

14

References

Akerlof, George A. 1970. "The Market for "Lemons": Quality Uncertainty and the Market Mechanism." *The Quarterly Journal of Economics*, 84(3): 488–500.

Baker, George P., and Thomas N. Hubbard. 2004. "Contractibility and Asset Ownership: On-Board Computers and Governance in US Trucking." *Quarterly Journal of Economics*, 119(4): 1443–1479.

Bond, Eric W. 1982. "A Direct Test of the "Lemons" Model: The Market for Used Pickup Trucks." *The American Economic Review*, 72(4): 836–840.

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

Grossman, Sanford J. 1981. "The Informational Role of Warranties and Private Disclosure about Product Quality." *Journal of Law and Economics*, 24(3): 461–483.

Grossman, Sanford J., and Oliver D. Hart. 1980. "Disclosure Laws and Takeover Bids." *Journal of Finance*, 35(2).

Houser, Daniel, and John Wooders. 2006. "Reputation in Auctions: Theory and Evidence from eBay." *Journal of Economics and Management Strategy*, 15(2): 353–370.

Jin, Ginger Zhe. 2005. "Competition and Disclosure Incentives: An Empirical Study of HMOs." *Rand Journal of Economics*, 36(1): 93–112.

Jin, Ginger Zhe, and Andrew Kato. 2006. "Price, Quality and Reputation: Evidence from an Online Field Experiment." *Rand Journal of Economics*, 37(4).

Jin, Ginger Zhe, and Phillip Leslie. 2003. "The Effect of Information on Product Quality: Evidence From Restaurant Hygiene Grade Cards." *Quarterly Journal of Economics*, 118(2): 409–451.

Jovanovic, Boyan. 1982. "Truthful Disclosure of Information." *The Bell Journal of Economics*, 13(1): 36–44.

Mathios, Alan D. 2000. "The Impact of mandatory disclosure laws on product choice: An analysis of the salad dressing market." *Journal of Law and Economics*, 43(2): 651–677.

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

Resnick, Paul, and Richard Zeckhauser. 2005. "Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System." Vol. 11 of *Advances in Applied Microeconomics*, Chapter 5. Elsevier Science.

Yin, Pai-Ling. 2008. "A Survey-Based Procedure for Measuring Uncertainty or Heterogeneous Preferences in Markets, with Application to a New Test for Common Values." working paper, MIT Sloan School of Management.