American Economic Journal: Microeconomics 2015, 7(3): 85–103
http://dx.doi.org/10.1257/mic.20120085

Cyber-Shilling in Automobile Auctions: Evidence from a Field Experiment†

By David Grether, David Porter, and Matthew Shum*

We run a large field experiment with an online company specializing in selling used automobiles via ascending auctions. We manipulate experimentally the "price grid," or the possible amounts that bidders can bid above the current standing price. Using two diverse auction sites, one in New York and one in Texas, we find that buyer and seller behavior differs strikingly across the two sites. Specifically, in Texas we find peculiar patterns of bidding among a small but prominent group of buyers suggesting that they are "cyber-shills" working on behalf of sellers. These patterns do not appear in the New York auctions. (*JEL* C93, D12, D44, L62, L81)

For the past two decades, the possibilities of online retailing have transformed the retail landscape. When Internet-based retailers and resellers first appeared, they were hailed as initiating a new era of retailing which would see an end to the search frictions and transactions costs that characterize the brick-and-mortar shopping experience. While it is true that online retailing has made shopping much more convenient for many consumers, there is mounting evidence that anonymity—a defining characteristic of online commerce—has introduced new possibilities for buyer and seller manipulation and mischief in retail markets that were not present in brick-and-mortar settings. As examples, in an online setting, it is easy for retailers to hide, or "shroud," certain attributes of a transaction (such as shipping costs), which can lead consumers to end up paying much more than they intended to for a given product. Also, in most online marketplaces, it is relatively straightforward for buyers and sellers to change their identities, or maintain multiple identities—"cyber-shilling," in the parlance of the popular press—thus increasing the possibilities for fraudulent transactions.[1] Indeed, the largest online marketplaces (including eBay and Amazon) maintain a large workforce to police and verify claims against the hordes of buyers and sellers using their platforms.

* Grether: Caltech, HSS (Mailcode 228-77), 1200 E. California Blvd., Pasadena, CA 91125 (e-mail: dmg@hss.caltech.edu); Porter: Economic Science Institute, Chapman University, One University Drive, Orange, CA 92886 (e-mail: dporter@chapman.edu); Shum: Caltech, HSS (Mailcode 228-77), 1200 E. California Blvd., Pasadena, CA 91125 (e-mail: mshum@caltech.edu). We thank the Gordon and Betty Moore Foundation for grant GBMF#1158, which facilitated this research. We also thank Amit Gandhi, Yaron Raviv, Jimmy Roberts, Tom Ruchti, Artie Zillante, and seminar participants at Florida State University and Chapman University.

† Go to http://dx.doi.org/10.1257/mic.20120085 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

[1] See Gabaix and Laibson (2006); Ellison and Ellison (2009); Jin and Kato (2006). There are also many stories in the popular press of sellers using fake accounts or contractors to write positive reviews of their own products.

In this paper, we present a case study from a large online used car auctioneer, and document several striking types of buyer and seller behavior within these online car auctions. The data used in our analysis derive from a set of field experiments that were performed on the online automobile auctions run by Copart, Inc., one of the largest sellers of salvage vehicles worldwide. We analyzed data from 24 online auctions at two auction locations (New York and Texas) at which approximately 15,000 vehicles were auctioned.

The auctions held in Texas revealed some puzzling behavior by the buyers. There, we found a frequent tendency for bidders to raise their own standing bids repetitively. While this behavior has been linked with strategic behavior (see Cramton 1997), the pattern we found does not fit that type of conduct. After a closer investigation, we found that most of this "repeat bidding" behavior can be explained by a small and peculiar group of buyers present in the auctions. Features of these bidders' behavior strongly hint at their being "shill" bidders confederate with sellers, who attempt to drive up the price of the cars. Each of these bidders bid almost exclusively for cars sold from a particular seller. Additionally, these bidders bid frequently, but almost never consummate a sale. However, we also find that this suspected shill bidding strategy is not clearly successful for the sellers who employ it; there were very few auctions in which the shill bidders submitted the second-highest bid and, thus, played the role of marginal price setters.

However, these behavioral patterns are not present in the New York-area auctions. One explanation for the differences between these two locales is the makeup of the car sellers: car dealers are much more prominent in Texas, selling almost half the cars, but dealers make up only 10 percent of the sales in New York. We conjecture that perhaps used car dealers—with lower storage costs and higher outside options—are more willing to engage in surreptitious misbehavior (such as shill bidding), and indeed, the "shill bidders" we observed in Texas were always bidding for cars sold by dealers. As far as we are aware, these results constitute some of the first direct evidence of shill bidding in online auction markets. Anecdotally, this possibility has been raised, and it has been studied in lab experiments,[2] but this paper appears among the first to present credible evidence of shill bidding in the field.

Our findings emerge within the context of a set of field experiments which were designed to examine the role of a price grid—or "click box bidding," as it is known—on bidding behavior in ascending auctions. Against the backdrop of the spectrum auctions run by the US Federal Communications Commission (FCC), in which bidders allegedly colluded using a "trailing digits" jump-bidding strategy[3] whereby the final three digits in a bid signaled to rival bidders the license where they would be punished if they persisted in bidding on a specific license,[4] we intended to examine how a price grid—which rules out this type of behavior by restricting the bids that could be submitted—would affect bidding behavior and revenue. In these

[2] See Kosmopoulou and de Silva (2007).

[3] See Cramton and Schwartz (2000) for details of this strategy.

[4] Market numbers are two or three digits and bids are typically six figures or more, so a bid could contain, at negligible cost, the market number as its last few digits, prefaced by leading zeroes to highlight the trailing digits.

auctions, we varied experimentally the price grid which restricts a bidder to select a bid on a grid that consists of integer multiples of a given bid increment.[5]

Manipulation of the price grid interlaces with several phenomena that have been commonly observed in real-world auctions. First, jump bidding is an endemic feature of real-world ascending auctions; this includes not only the FCC wireless spectrum auctions[6] mentioned above, but also online (eBay) auctions, and also conventional art and antiquities auctions run by Sotheby's and Christies for hundreds of years.[7] Indeed, the prevalence of jump bidding presents a puzzle for standard auction theory;[8] as a result, there is a small but growing theoretical literature explaining jump bidding from a strategic perspective.[9]

A second set of phenomena that will be affected by our manipulations of the grid size is the degree that bidders learn and update their valuations for the goods being sold by observing what rivals bid.[10] In auctions where the objects for sale have a common value component, bidders will typically increase their expectation of the object's value upon observing more aggressive bidding behavior by the rivals. Additionally, rapidly escalating prices may make the auction more emotionally exciting, which may encourage more aggressive bidding by other bidders.[11]

This, in turn, raises the possibility that wily sellers may employ confederate shill bidders who submit aggressive bids in order to drive up prices in an auction. Moreover, because a coarser price grid enables more aggressive bidding (i.e., enables bidders to place bids much above the current standing price), we may expect shills to bid more prominently under a coarser price grid. Shills may take advantage of the larger jump opportunities available under a coarser price grid in order to mislead or excite the nonshill bidders. Indeed, we find that our suspected shill bidders employ large jump bids more frequently than other bidders; patterns such as these demonstrate how the experimental changes in the price grid connects with the calculus of shill bidding.

[5] Thus, if the increment amount were 10 percent for a particular object, any bid submitted for that object was restricted to be equal to the Standing Bid times $(1 + 0.10\pi)$, where π is a positive integer greater than or equal to 1 but less than 10.

[6] See Isaac, Salmon, and Zillante (2007); Plott and Salmon (2004); and Cramton (1997) for details of this behavior.

[7] Jump bidding has also been observed in many experimental implementations of ascending auctions (McCabe, Rassenti, and Smith 1990; Banks et al. 2003; Coppinger, Smith, and Titus 1980; and Lucking-Reiley 1999).

[8] For book-length treatments of auction theory, see Milgrom (2004) and Krishna (2002). In one standard model of English auctions—the so-called "clock" or "button" auction (Milgrom and Weber 1982)—the price is set by a clock which rises automatically, and bidders indicate their willingness to pay the current price by holding down a button. Once a bidder releases his button, however, he "drops out" of the auction, and can no longer reenter. In such a setting, it is not obvious what benefits bidders derive from "speeding up the clock" by jumping.

[9] This includes Avery's (1998) model of jump bidding as a means for bidder intimidation in an affiliated values setting. Other models of jump bidding, primarily in the independent private values setting, include Daniel and Hirschleifer (1998); Easley and Tenorio (2004); Hoerner and Sahuguet (2007); and Isaac, Salmon, and Zillante (2007).

[10] See Milgrom and Weber (1982) or Hong and Shum (2003).

[11] cf. Ding et al. (2005).

I. Field Experiments: Used-Car Auctions at copart.com

We designed a set of unique field experiments[12] using an online ascending auction for automobiles. Specifically, we created an experiment with Copart Inc., a publicly-traded (NASDAQ: CPRT) company that is the largest auction house for salvage vehicles in the world. In these auctions, we manipulated the way bidders could engage in jump bidding, by restricting the maximum amount that bids could be submitted above the current standing price. Before describing our experimental design in detail, we begin with a description of Copart and its online auction mechanism.

Copart sells well over a million cars annually through its online virtual auction. On average, each business day, Copart auctions around 5,000 vehicles on its site. Copart is an intermediary that obtains vehicles from governments, charities, finance companies, banks, dealers, fleets, rental car companies, and the insurance industry; there are over 150 Copart facilities throughout the United States, Canada, and the United Kingdom. Buyers are located all over the world and auctions are conducted each business day at various Copart facilities. Our experiments utilize Copart's largest auction yard (in Houston, Texas) and another geographically different yard in upstate New York to examine the effect of price grid changes on observed auction outcomes. We use data from 24 auctions—13 run under the company's baseline parameters, and 11 run under altered parameters introduced by us. The volume varies across the sales, but each auction has approximately 500 vehicles offered for sale.[13]

A. *Copart Auctions: Main Features*

Copart auctions attract a large number of heterogeneous buyers located around the world (auto parts dismantlers, rebuilders, used car dealers, wrecking yards, and the public). Upon registering an account with Copart, a buyer can access the "current sales" button to view all of the auctions occurring that day, the locations of the auctions, and the start times. Buyers can join an auction at any time, and also view vehicles in upcoming auctions. Each auction shows pictures of the vehicle up for auction, its make, model, and year, along with the list of details shown in Table 1. Figure 1 shows a typical auction screen from the Copart auction site.

The Copart bidding process begins with a preliminary bidding (proxy bids) stage. The preliminary bidding process, which ends 60 minutes before the start of the virtual auction, allows participants from around the world to preview vehicles for sale in each of Copart's facilities in person or over the Internet. Using preliminary bidding, participants enter the maximum price (called the "Bid4U Max") they are willing to pay for a specific vehicle and the software incrementally bids for the vehicle on their behalf. In the preliminary bidding stage, all of the preliminary bids are

[12] cf. Harrison and List (2004).

[13] The scale of the experiment is comparable to that of the sequencing experiments with used car auctions reported in Grether and Plott (2009). Relatedly, Tadelis and Zettelmeyer (2009), use field experiments with a used automobile auction company to explore how providing more information (in the form of "Standardized Condition Reports" describing a used car's condition) to bidders affects auction outcomes, particularly revenues.

Table 1—Lot Details and Variable Definitions

Actual cash value	*Estimated* retail value of the lot as submitted to Copart by the seller. If the lot has been damaged, this is the value prior to the occurrence of the damage. The number is only informational.
Repair cost	*Estimated* cost to repair the vehicle as submitted to Copart by the seller of the vehicle.
Title state/type	Title type denotes the ownership documents that will be transferred to the buyer.
Odometer	Odometer codes are shown to reflect the known reliability of the odometer reading.
Primary damage	Location of the major damage on the car.
Secondary damage	Location of the minor damage on the car.
VIN	Vehicle Identification Number assigned by the manufacturer.
Body style	Body style is the manufacturer's designation of the vehicle's configuration.
Color	Color listed on this site is the common color name that reasonably represents the exterior color of the vehicle.
Engine	Engine is the motor.
Drive and cylinders	Manufacturer's designation of the vehicle's power train.
Fuel	Designates the fuel type used by the engine as designated by the VIN.
Keys	Indicates whether Copart is in possession of the keys to the lot.

Figure 1. Screenshot of Copart Bidder Screen

Table 2—Minimum Bid Increments

Bid range	Minimum increment
$0–$5	$1
$5–$40	$5
$40–$100	$10
$100–$1,000	$25
$1,000–$5,000	$50
$5,000–$25,000	$100
$25,000–$50,000	$250
$50,000–$100,000	$500
$100,000–$10,000,000	$1,000

incremented until only the highest preliminary bid is left. When the online bidding begins, the opening price is set equal to the second highest preliminary bid plus one increment. As the online bidding process starts, the remaining (highest) preliminary bidder has their bid controlled by Copart software, which automatically bids one bid increment above the current high bid (standing bid) for the vehicle, until their Bid4U Max is reached. Bidders in the online auction are not informed if a bid is coming from the preliminary bidder or not.

The car to be auctioned is called a *lot* and is sold sequentially in *lanes* at each facility called a *yard*. Once the starting price is determined, the *bid increment* is set based on the current bid. Table 2 shows how the bid increments change during the course of an auction, depending on the level of the current (or "standing") bid. Once the auction is underway, bidders can submit bids in real time that are equal to one of the following bid grid size:

- the current bid plus the minimum increment; or
- the current bid plus five times the minimum increment; or
- the current bid plus ten times the minimum increment.

As shown in Figure 1, the buttons for the different bid choices available to the bidders are located prominently on the lower right-hand side of the bidder screen. Bidders are identified by location: bidders in North America are identified by the state or province where they are located, and bidders outside of North America are identified by country. So from the perspective of market participants, a bidding competition between two or more buyers in the same state cannot be distinguished from a single buyer submitting a series of bids.

Once a bidder submits one of these three bids it becomes the new standing bid and if no new standing bid is made in two seconds, then there is a five second countdown displayed on the bidder screen. If no new standing bid is provided in those five seconds, the auction ends. Thus, if no bid is received in seven seconds the auction is over. In our data the actual median time between bids is about one second with the average time approximately 2.5 seconds. The distribution of interbid times is bimodal with a large mode at zero (presumably the automatic increments for the winning preliminary bidder) and a second smaller mode at seven seconds. Histograms of the distribution of interbid times for the New York and Texas yards are presented in, respectively, Figures 2 and 3. In some cases the time between

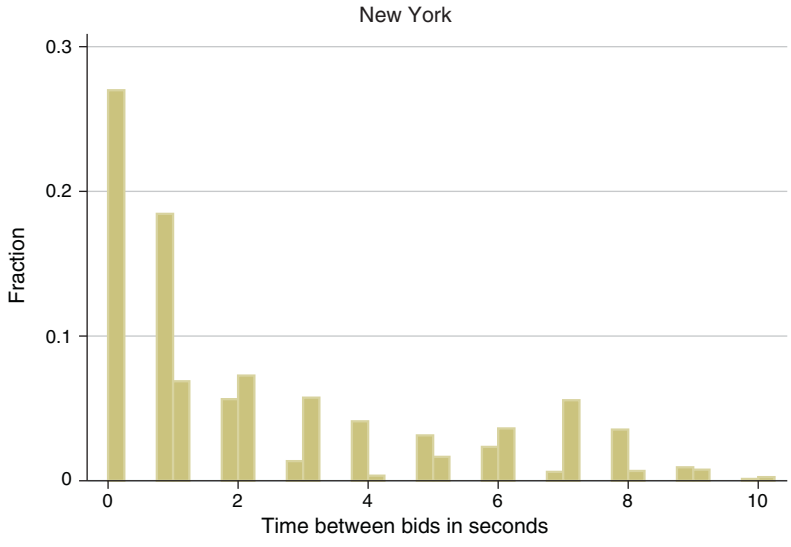

FIGURE 2. INTERBID TIMES

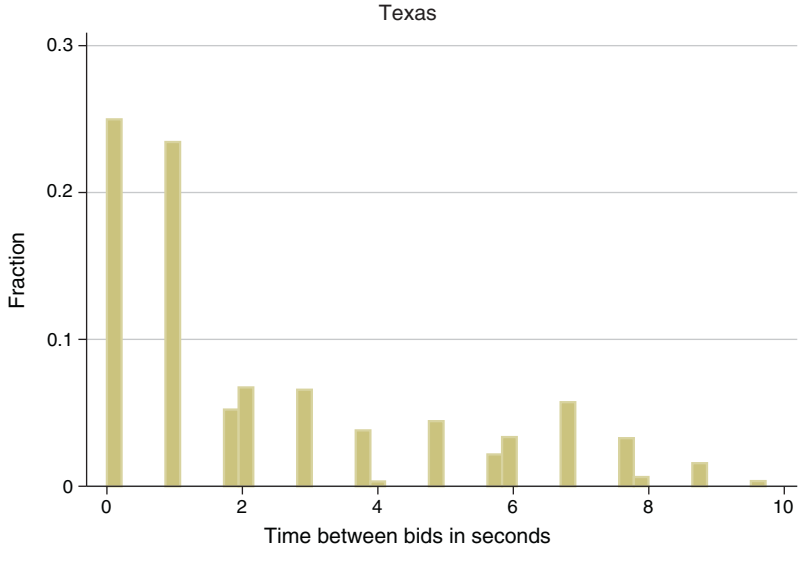

FIGURE 3. INTERBID TIMES

bids exceeds seven seconds, but is never greater than eleven seconds. These longer intervals are caused by delays due to the latencies with the Internet. These auctions move quickly with most not exceeding one minute.

Sellers in these auctions include insurance companies, dealers, charities, rental car companies, governmental units, and single car sellers. Sellers in Copart auctions can and typically do set a "minimum bid"; this is essentially a secret reserve price,[14]

[14] If the minimum bid is $0, this is listed as a "pure sale" in the auction. If there is a minimum bid required, it is always secret to the bidders (but they know that there is a reserve price on the lot).

which is unobserved to bidders at the time they choose their bids, such that if the highest bid in the auction falls below it, the seller has the option to not sell to the highest bidder.[15] Sellers who do not want to state a minimum bid may place the lot "on approval," thereby reserving the right to approve the sale at the high bid or to negotiate with the high bidder. In these cases at the start of the virtual auctions it is announced that the auction is "on approval." Importantly, if the minimum bid is met during the course of the virtual auction, an announcement is made that the lot is "*sellin' all the way*."

If the bidding does not reach the reserve price the seller may negotiate with the high bidder or in some cases with the second highest bidder. Copart's new revised auction site specifically highlights this feature noting that bidders may engage in negotiations with sellers who "reveal or eliminate their minimum bid requirement to speed up the final sale to you." This opportunity for post-auction negotiation is the source for some striking results below.

B. *Experimental Design: Price Grid Interventions*

In our field experiment we manipulated the *price grid* that bidders could choose when submitting their bid. As we noted above, the standard Copart auction rules allow bidders to submit jump bids, which are either five or ten times the bid increment above the current bid. We call this the *baseline treatment*, and denote it by (1,5,10). We introduced two contrasting treatments. First, we have a *fine grid treatment*, which restricts bids to only two or three times the bid increment above the current bid. We denote this treatment by (1,2,3). Second, we have a *coarse grid treatment*, which allows bidders to bid 10 or 20 times the bid increment above the current bid.[16] This treatment is referred to as (1,10,20). For convenience, in the remainder of this paper, we refer to bids submitted at a level more than one bid increment above the current bid as "jumps."

Two Copart yards—in Houston, Texas and upstate New York—were used in our study. The Texas yard has greater volume with two sales per week while the New York yard and all the other company yards host only one sale per week. The volume varies, but averages around 500 vehicles per sale. At both yards insurance companies are the owners of around 40 percent of the vehicles offered for sale. At the New York site, the other main sources of vehicles are governments and municipalities (20 percent) and charities (18 percent). Notably, car dealers account for only about 10 percent of sales in New York. The seller mix at the Texas site, however, is quite different; dealers have the most prominent presence there, and account for 45 percent of the lots offered for sale.

Table 3 lists the sequence of our treatments by date and yard. For each lot we have information about the item and summary bid data. The information on the lot

[15] Secret reserve prices are actually commonly observed in real-world auctions, but not completely understood from an auction-theoretic point of view. See Bajari and Hortaçsu (2003) and Katkar and Reiley (2007) for empirical and experimental work exploring secret reserve prices, and Elyakime et al. (1994), Vincent (1995) for theoretical analyses.

[16] We had initially proposed a treatment that eliminated jumps completely (i.e., "(1,1,1)"), but this was not feasible due to software limitations in Copart's online bidding system.

Table 3—Treatment Application

Yard	Date	Treatment	Number of lots in sample
Texas	2/19/10	1,5,10 (baseline)	407
Texas	2/23/10	1,2,3 (fine grid)	497
Texas	2/26/10	1,5,10 (baseline)	557
Texas	3/2/10	1,5,10 (baseline)	489
Texas	3/5/10	1,2,3 (fine grid)	548
Texas	4/20/10	1,5,10 (baseline)	513
Texas	4/23/10	1,2,3 (fine grid)	723
Texas	4/27/10	1,2,3 (fine grid)	483
Texas	4/30/10	1,5,10 (baseline)	637
New York	5/19/10	1,5,10 (baseline)	712
Texas	5/25/10	1,5,10 (baseline)	687
New York	5/26/10	1,2,3 (fine grid)	655
Texas	5/28/10	1,10,20 (coarse grid)	687
Texas	6/1/10	1,10,20 (coarse grid)	535
New York	6/2/10	1,5,10 (baseline)	583
Texas	6/4/10	1,5,10 (baseline)	524
New York	8/11/10	1,5,10 (baseline)	547
Texas	8/17/10	1,5,10 (baseline)	611
New York	8/18/10	1,10,20 (coarse grid)	546
Texas	8/20/10	1,10,20 (coarse grid)	698
Texas	8/24/10	1,10,20 (coarse grid)	447
New York	8/25/10	1,10,20 (coarse grid)	548
Texas	8/27/10	1,5,10 (baseline)	741
New York	9/1/10	1,5,10 (baseline)	574

includes the description (make, model year), damage including repair cost (seller's estimate), mileage, title type and state of registration, and the number of times the lot has been previously auctioned. In our empirical work below, these are the main variables used to control for heterogeneity across lots. Since these are virtual auctions and, for the most part, bidders are offsite and do not have the opportunity to inspect vehicles physically before bidding, we believe that these variables capture practically all heterogeneity observed by bidders before they bid. Moreover, our treatments were not announced to sellers before the auction, so that they do not have an opportunity to adjust their reserve prices in response to our treatments.

The observed variables for each car includes the minimum bid, the starting bid, number of bids and jump bids, the size of the jumps, the high bid, the selling price (listed as zero if the seller did not accept the price), the high bidder (coded) including the state and nationality of the high bidder and the seller's identity (coded), and the type of the seller. The final sale price may differ from the high bid as a result of negotiations between the seller and the first or second highest bidder. In addition, for each lot we observe the complete sequence of bids and bidder identities, allowing us to determine accurately whether a bidder submitted a bid more than one increment above the standing bid, and the amount of the jump.

II. Empirical Results

Summary statistics for the two yards and the three treatment conditions are given in Table 4. First, we confirm that the treatments are *effective* in that the proportion of jumps in both the coarse and fine grid treatments are significantly different from the

Table 4—Average Bidding Behavior for Each Yard/Lot/Treatment

Yard	Number of bidders	Number of bids	High bid	Jump bid frequency	Repeat-bid frequency	Size of jumps	Proportion sold	Sales price if sold
Texas (1,5,10)	4.15	23.8	5,330	0.024	0.184	440	0.576	3,778
Texas (1,2,3)	4.52	27.1	5,503	0.035	0.181	186	0.579	3,915
	(0.00)	(0.00)	(0.19)	(0.00)	(0.08)	(0.00)	(0.21)	(0.40)
Texas (1,10,20)	4.23	26.0	5,201	0.023	0.192	922	0.543	3,592
	(0.20)	(0.00)	(0.75)	(0.00)	(00)	(0.00)	(0.01)	(0.22)
New York (1,5,10)	3.02	15.2	2,217	0.015	0.013	322	0.812	2,122
New York (1,2,3)	2.87	11.3	2,478	0.019	0.014	154	0.811	2,466
	(0.08)	(0.49)	(0.12)	(0.01)	(0.47)	(0.00)	(0.96)	(0.06)
New York (1,10,20)	3.21	11.3	2,222	0.006	0.013	508	0.765	2,379
	(0.01)	(0.40)	(0.97)	(00)	(0.99)	(0.00)	(0.00)	(0.09)

Note: Figures in parentheses are significance levels for testing equality with baseline (1,5,10).

proportions in the baseline (1,5,10) treatment. The observed treatment effects are sensible with the number of jumps increasing when the jump size is restricted, and falling when the jump size is increased. The actual number of jumps varies somewhat across the auction sites. At the New York location approximately 1.5 percent of the bids are jumps with roughly 6 percent of the buyers jumping at least once. Jump bids are more frequent at the Texas site with about 9 percent of the bidders jumping at least once and jump bids accounting for approximately 2.4 percent of the bids. Moreover, the auctions in Texas have about twice as many bids as those at the New York yard, and the high bids are roughly twice those at New York.

Second, looking at average prices with the three treatments does not reveal any substantial revenue effects of changing the grid size. The average high bid does not vary significantly nor does the average sale price (conditional on the vehicle being sold). However, the overall revenue effects are ambiguous and consistent across the two locations. At both places the high bids and sale prices were slightly higher under the fine grid treatment and slightly lower with the coarse treatment, but the differences were not significant in Texas and only marginally so in New York. However, the proportion of vehicles sold decreases at both yards under the coarse grid treatment.

The features of strategic jump bidding emphasized in the existing theoretical literature are missing from our findings. These models (cf. footnote 9 above) often have equilibria in which auctions start or end with jump bids. In our data this does not happen. The fraction of first bids that are jumps is only slightly higher than the overall jump rates at both locations (0.022 in New York and 0.036 at the Texas site). At the New York site, the proportion of final bids that are jumps is about the same (1.5 percent) as the overall proportion of jump bids. In Texas, about 3.2 percent of the sales conclude with jumps. Overall, these modest figures do not support features of equilibria from strategic jump bidding models.

Table 5 summarizes the behavior of sellers at the two locations. There is a substantial difference in the fractions of cars unsold (20 percent in New York and 42.8 percent in Texas). Variation in the composition of the seller groups seems to account for the differences between the two sites. Insurance companies account for about 40 percent of the vehicles at both locations, while dealers sell only about

Table 5—Seller Behavior

	Texas			New York	
	Number of lots	Percent		Number of lots	Percent
All lots			*All lots*		
Sell at high bid	4,948	50.5	Sell at high bid	3,183	76.3
Negotiate price	660	6.7	Negotiate price	154	3.7
Withdraw	4,197	42.8	Withdraw	833	20.0
Total	9,805		Total	4,170	
Lots with high bid < minimum bid			*Lots with high bid < minimum bid*		
Sell at high bid	2,446	34.7	Sell at high bid	446	38.6
Negotiate price	660	9.4	Negotiate price	154	13.3
Withdraw	3,943	55.9	Withdraw	556	48.1
Total	7,049		Total	1,156	
Lots with high bid ≥ minimum bid			*Lots with high bid ≥ minimum bid*		
Sell at high bid	2,502	90.8	Sell at high bid	2,737	90.8
Negotiate price	0	0	Negotiate price	0	0
Withdraw	254	9.2	Withdraw	277	9.2
Total	2,756		Total	3,014	

Table 6—Seller Behavior: Car Dealers versus Insurance Companies

	Car dealers			Insurance companies	
	Number of lots	Percent		Number of lots	Percent
All lots			*All lots*		
Sell at high bid	1,291	27	Sell at high bid	5,065	83
Negotiate price	327	7	Negotiate price	301	5
Withdraw	3,121	66	Withdraw	742	12
Total	4,739		Total	6,108	
Lots with high bid < minimum bid			*Lots with high bid < minimum bid*		
Sell at high bid	1,093	24	Sell at high bid	1,341	62
Negotiate price	327	7	Negotiate price	301	14
Withdraw	3,069	68	Withdraw	535	24
Total	4,489		Total	2,177	
Lots with high bid > minimum bid			*Lots with high bid > minimum bid*		
Sell at high bid	198	79	Sell at high bid	3,724	95
Negotiate price	0	0	Negotiate price	0	0
Withdraw	52	21	Withdraw	207	5
Total	250		Total	3,931	

10 percent of the vehicles in New York but nearly half the vehicles in Texas. While insurance companies in our sample sell over 80 percent of their cars at auction, dealer sales rates are mainly in the 35 percent range. Table 6 gives results similar to those in Table 5, for vehicles auctioned by dealer and insurance companies. The re-auctioning[17] of cars at the Texas site is about twice the rate observed at the New York site. Looking at the number of times a vehicle has been auctioned, the median is one in New York and two in Texas and the numbers are about double at

[17] A re-auctioned lot is a lot that was previously offered in an earlier auction but not sold.

Table 7—Regression Results for New York and Texas Sales

	New York		Texas	
Dependent variable	OLS high bid	OLS sale price if sold	OLS high bid	OLS sale price if sold
FINEGRID	197.8	284.6	93.3	33.4
	(1.7)*	(2.3)**	(0.8)	(0.4)
COARSEGRID	71.8	51.6	77.0	−13.0
	(0.7)	(0.5)	(0.7)	(0.1)
Odometer	5.2	−4.3	9.7	1.5
	(0.8)	(0.6)	(1.8)*	(0.4)
Actual odometer	44.7	−51.9	909.9	461.0
	(0.5)	(0.5)	(9.9)***	(5.6)***
Seller book value	0.26	0.27	0.33	0.26
	(49.7)***	(46.3)***	(118.8)***	(65.4)***
Age	−168.7	−142.0	−73.7	−211.7
	(7.8)***	(62)***	(4.3)***	(11.3)***
Age squared	2.1	2.2	0.85	4.2
	(3.8)***	(3.7)***	(2.0)**	(7.3)***
Constant	6,518.3	−3,168.8	−132.5	2,156.8
	(0.6)	(0.3)	(0.1)	(2.6)**
Observations	4,161	3,329	9,781	5,593
Cubic trend	Yes	Yes	Yes	Yes
Make dummies	Yes	Yes	Yes	Yes
Lane dummies	Yes	Yes	Yes	Yes
Primary damage dummies	Yes	Yes	Yes	Yes
R^2	0.64	0.65	0.73	0.70

Notes: T-stats are in parentheses. FINEGRID denotes the (1,2,3) grid treatment, and COARSEGRID is the (1,10,20) treatment. The omitted treatment is (1,5,10).

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

the quartiles and, at the ninety-ninth percentile: 7 for New York and 14 for Texas. In the New York auctions the high bids usually exceed or meet the reserve prices (72.3 percent of the time), while in Texas the situation is almost the reverse with high bids being less than the reserve in 70.2 percent of the auctions. Cars that do not sell at an auction often return later, which makes the revenue effects of changing proportions sold more ambiguous.

Results from linear regressions, reported in Table 7, confirm the robustness of the results discussed above after controlling for a large number of car characteristics. In New York, high bids and sale prices are about $200 to $300 higher in the fine grid treatment with borderline significance levels. In this table, FINEGRID is an indicator for the (1,2,3)-price grid treatment, while COARSEGRID denotes the (1,10,20) treatment ((1,5,10) serves as the baseline treatment). None of the treatment variables are significant in the Texas equations.

Post-Auction Price Negotiation.—As we discussed before, due to the prevalent use of secret reserve prices, sellers are not required to sell the car at the high bid. From Table 5, we see that seller behavior varies substantially depending on whether the minimum bid (i.e., secret reserve price) is exceeded in the auction. For the Texas

sales, we see that when the final bid is below the minimum bid, sellers sell the car at the final bid only 34.7 percent of the time, and withdraw the car 55.9 percent. When the final bid exceeds the minimum bid, however, sellers sell the car 90.8 percent of the time. The remaining 9.2 percent are cars that are not sold even when the final bid exceeds the minimum bid, which arises from buyers reneging their winning bid.[18] When a buyer reneges, the seller can negotiate with the second highest bidder to sell the lot. This results in sales 70 percent of the time. Similar figures hold for the New York sales.

As we noted above, car dealers make up around 45 percent of the sellers in Texas, while they constitute only about 10 percent of the sellers in New York. In Table 6, we compare the behavior for car dealers versus insurance companies, which are the largest group of sellers in both Texas and New York. As we see there, the differential behavior of these two types of sellers can explain the differences in seller behavior between Texas and New York—car dealers are much less likely to sell their cars at the high bid than insurance companies (27 percent versus 83 percent).[19]

These differences in selling rates across seller types are associated with differences in the use of minimum bids. Table 8 gives regression results with the reserve price (ie., the minimum bid) as the dependent variable using the same independent variables as in Table 7 plus indicator variables for the seller being a dealer or an insurance company. As many vehicles are auctioned with a reserve price of $0, the so-called "pure sales" mentioned above, the equation is estimated using a Tobit specification. The results show that dealers' reserve prices are substantially higher (by $1,000 to $3,000) and insurance companies' reserve prices are much lower (by $3,000) than those of other sellers. Since the price grid treatments were not announced in advance of the auctions, the treatment effects on reserve prices should be zero, which is confirmed by the insignificance of their regression coefficients.

A. *A Puzzle: Repeat Bidding Behavior*

The results discussed above indicate that our manipulated changes in grid size appeared to have only modest (if any) effect on car revenues. However, here we present evidence on a particularly striking and unexpected bidding phenomenon that occurred with some frequency in Copart auctions—that of *repeat bidding*.[20] We say that a buyer engages in repeat bidding when he/she submits two consecutive bids. Thus, repeat bidders are *raising their own bids*; obviously it is difficult to rationalize such behavior because, by repeat bidding, a buyer essentially increases the price he pays if he were to win the auction. Repeat bids make up less than 2 percent of the bids in New York. In Texas, however, the proportion of repeat bids is much higher, ranging from 14 percent to over 21 percent.

[18] Buyers that renege on their bid must pay a penalty of $400 or 10 percent of the sales price, whichever is greater. This amount is paid to Copart. In addition, a seller can relist a car up to three times. After that they must pay an $85 fee to relist it again. Buyer fees for winning bid vary by state, the actual fees can be determined using the Copart Fee Calculator at http://gotparts.com/copart/Fees.html.

[19] Table 8 summarizes the characteristics of cars brought to auction by seller type.

[20] In the history of the Copart auctions, repeat bidding at one time was not allowed. However, an uproar by the bidders caused Copart to allow such bidding to be part of its current design.

Table 8—Reserve Price Regressions

Dependent variable	Texas reserve price	New York reserve price
FINEGRID	110.2	353.2
	(0.6)	(1.1)
COARSEGRID	−172.8	−11.8
	(1.0)	(0.0)
Odometer	26.6	12.21
	(3.2)***	(6.3)***
Actual	933.9	−350.0
	(6.7)***	(1.3)
Seller book value	0.48	0.30
	(113.9)***	(23.6)***
Dealer	1414.4	3250.0
	(5.9)***	(9.1)***
Insurer	−3621.8	−2882.6
	(15.0)***	(10.4)***
Age	129.2	−273.7
	(3.9)***	(4.7)***
Age squared	−3.1	3.2
	(3.9)***	(2.3)**
Cubic trend	Yes	Yes
Make, lane, and damage dummies	Yes	Yes
Make dummies	Yes	Yes
Lane dummies	Yes	Yes
Damage dummies		
Observations	9,117	2,998

Note: Excludes vehicles sold on approval.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

These auctions move very fast with the typical lot lasting under a minute, so that it is possible that bidders may accidentally and mistakenly bid against themselves: "trembling," in short. This explanation is plausible for much of the repeat bidding at the New York site. There, most bidders who repeat-bid do so only once or twice in a weekly sale, and the most frequent do so on the order of only 10 to 20 times.

Additionally, the repeat bids are also (partially) explained by examining their timing patterns. Repeat bids occur with a longer lag, i.e., with a longer time interval from the preceding bid, than nonrepeat bids. At a typical sale, the median time interval between repeat bids is six seconds while it is only two seconds for the nonrepeating bids. Recall that reserve prices at Copart auctions are secret (that is, unknown by bidders when they choose their bids). Moreover, if the bidding starts above the secret reserve price or goes above it during the bidding, an audible announcement that the lot is "selling all the way" is emitted on the website. However, virtually all of the repeat bids take place below the reserve price. At the Texas site the proportion of bids that are repeats drops by more than half when the minimum bid is passed. This suggests that some repeat bidding (and, to a lesser extent, jump bidding) may be symptomatic of a type of search behavior by which some bidders try to discover the reserve price, but do not want to risk going over it.

TABLE 9—OUTCOMES FOR THE SUSPECTED SHILL BIDDERS IN TEXAS AUCTIONS

Buyer	Repeats	Bids	High bid	Buys	High bid > reserve
1	9,451	20,479	426	14	4
2	8,746	12,986	373	24	15
3	7,661	16,580	378	40	24
4	5,296	8,723	301	6	5
5	3,542	5,529	151	18	8
6	615	1,543	277	8	3
7	18	7,244	315	5	5
Total	35,329	73,084	2,221	115	64

B. *Repeat Bidders in Texas. Are These Shills?*

While trembles and discovering the reserve price may explain some of the repeat bidding behavior, a deeper examination of the repeat bidding patterns in Texas reveals some peculiar and suspicious behavior among a small subset of buyers. In our Texas auctions, we observe approximately 260,000 bids, of which 18 percent (about 47,000) are repeat bids. There are 5,659 different bidders in the sample, but *six* buyers account for over 35,000 of the repeat bids. The number of bids, repeat bids, number of high bids, and cars purchased by the six buyers are shown in Table 9. These six buyers, all from Texas, had the high bid in nearly 20 percent of all the Texas auctions in our sample. If we look at the auctions in which these buyers had the highest bid, another sharp pattern emerges. With only one exception, each buyer was the high bidder in auctions from a single seller. The exception is buyer 4 (made the fourth largest number of repeat bids) who won multiple auctions from three sellers.[21] There is another buyer (tagged as number 7 in Table 9) who, while not engaging in repeat bidding, was the high bidder in 315 auctions, of which 285 pertained to the same seller (different from those shown in Table 9); only 5 out of these 315 auctions resulted in sales.

These behavioral patterns among these seven bidders are strongly suggestive as actions of bidders working in cahoots with specific sellers; in short, these seven bidders could be shill bidders. Kaufmann and Wood (2005, 2003), in their study of eBay coin auctions, list four types of behavior which they use to detect shill bidding:

(i) Shill bidders are more likely to buy items from few sellers;

(ii) They are more likely to drop out early from the auction;

(iii) They are more likely to bid at larger increments

(iv) Less likely to win an auction.

[21] Buyers and sellers are identified in the data by coded numbers. We know that some large sellers, for example insurance companies, have multiple seller numbers. It could be that buyer 4 is dealing with a single large seller who may have multiple lots or makes of cars, but we are not able to verify this.

For the seven bidders mentioned above, we have already shown (i). Our manipulations of the price grid are ideal for testing (iii), and the evidence here is immediate: this group of seven (out of 5,000+) bidders utilize the ten increment option (when it is available) much more often than other bidders; these bidders made about one-third of all ten increment jumps in the (1,5,10) and (1,10,20) treatments.

Apparently, these bidders prefer a coarser grid treatment because it provides an opportunity for more visibly aggressive bidding. Such bidding may not have much of an effect on auction outcomes in a private value environment in which bidders face little common uncertainty. But in the Copart environment, when bidders face common uncertainty about the level of competition, as well as potential common value components of the used cars, more aggressive bidding on the part of some bidders may be interpreted by other bidders as "good news" which might lead them to raise their estimate of the level of competition or the object's value, leading to a higher price.[22] Relatedly, a behavioral economics literature has also argued that more aggressive bidding can induce excitement or a frenzy among bidders and drive up prices, even in noncommon value settings.[23] This intuition highlights an interesting connection between the grid size manipulations in our experiments and how they impact the calculus of shill bidding.

Strictly speaking, Kaufmann and Wood's criteria numbers (ii) and (iv) do not apply to our context, because winning a Copart auction does not imply winning an object, due to the possibility of post-auction price negotiation with the seller. However, the logic of (ii) and (iv), adapted to our setting, implies that shill bidders should be less likely to consummate a sale, regardless of whether they win the auction. Looking at the data, we see sales rates are systematically lower for our suspected shills. Only about 5 percent of the auctions in which these seven buyers were the high bidders (115 out of 2,200) resulted in sales; in over 95 percent of these auctions, the high bid was lower than the reserve price, which enabled the seller to relist the car without paying Copart's relisting fee (up to three times; cf. footnote 19). Thus, these suspected shill bidders don't often take possession of cars which they "win" in these auctions, while the sales rates for these same sellers on auctions with high bids submitted by other (nonshill) buyers is nearly 80 percent. What happened to the 115 cars that the shill bidders actually bought? We have no evidence that these were relisted by the sellers (which would entail a reneging fee for the buyer); presumably, the sellers (who are all used car dealers) could sell these cars through other channels.

Furthermore, while these seven shill bidders win cars which are never relisted only around 5 percent of the time, they win (i.e., are the high bidder in) over 40 percent of the auctions of vehicles that appear more than once. (At one extreme, these bidders were the high bidders in 59 out of the 95 auctions for the 19 cars which were re-auctioned five times in our data!) Since the typical reason for relisting a car is that the high bid did not exceed the reserve price, this suggests that the shill bidders are not able to drive up the nonshill bids "high enough" for these cars.

[22] See Izmalkov (2004).

[23] See Ding et. al. (2005); Ku, Galinsky, and Murnighan (2008); Adam et. al. (2011); Malmendier and Lee (2011).

In short, these seven prominent bidders in Texas do behave in ways which have been associated with shill bidding in existing studies.

As a side note, we also find that these seven bidders in the Texas auctions are more likely to "win" an auction (resulting in a purported sale of a car) under a coarse rather than finer grid. Specifically, grouping these seven bidders together, we find that they win 34 percent of the auctions in the baseline (1,5,10) condition; this rises to 40 percent ($t = 2.4$) in the coarse grid (1,10,20) treatment but drops to 12 percent in the fine grid (1,2,3) treatment ($t = 9.0$).

III. Quo Vadis? Summaries and Conclusions

In this paper, we report the results from a field experiment in which we varied the size and scope of the price grid in ascending auctions run by a large online used car seller. We used these experimental manipulations to examine the effects of grid size on auction outcomes, including the probability of sale and revenue from a sale. We analyzed data from 24 online auctions at two auction locations (New York and Texas) at which approximately 15,000 vehicles were auctioned. In both New York and Texas, a coarsened price grid (which allows bidders to "jump" to amounts much higher than the current bid) lowers the probability of sale. In Texas, our findings can be explained by a small and peculiar group of buyers present in the auctions. These bidders exhibit behavior which, in the existing literature, has been associated with "shill" bidders confederate with sellers, who attempt to drive up the price of the cars.

As far as we are aware, this constitutes some of the first field evidence for shill bidding in the economics literature. Indeed, our knowledge of the Copart online auction marketplace and conversations with company representatives suggests that eliminating shill bidding in online auctions is virtually impossible since one can easily sign up for multiple different accounts. The ease in which one can disguise initial intentions in these auctions is almost costless.

Given the difficulties in getting rid of shill bidding in these auctions, we conclude with some evidence regarding the disruption caused by shill bidding in these auctions. How much of an advantage does shill bidding yield to the sellers who employ them? Our evidence suggests that these bidders had only a small and negligible effect on sellers' revenues. Most of the auctions in which they were the high bidders had bids below the sellers' reserve prices, and did not result in cars being sold. There were just 1,045 auctions in which one of the suspected shills was the second highest bidder, and, of these, 795 resulted in cars being sold (with median price $6,800). Thus, there were few auctions in which these shill bidders played the role of the marginal price-setting bidder.[24] Indeed, when we restrict attention to the sellers who employed shill bidders, we find that the distributions of purchase prices and winning bids for auctions won by the shill bidders are not significantly different from those

[24] For all auctions in which these bidders participated that resulted in sales, the median high bid was $7,700 (interquartile range $4,700 to $11,400), as compared to a median of $1900 (interquartile range of $875 to $3,750) for all other sellers. However, it is likely that these differences arise mainly because the sellers who would employ shills tend to specialize in high-end cars.

won by other bidders.[25] Thus, while our analysis shows the ease with which shill bidders can operate in the online auction environment, the evidence here suggests that they do not disrupt the functioning of this marketplace.

More broadly, "cyber-shilling" constitutes just one among many strategies that sellers may employ in an online auction setting. Two other important strategies used by sellers in the auctions we study are secret reserve prices, as well as post-auction price negotiation. In ongoing work, we are conducting additional field experiments to gauge the effect of seller strategies on bidder behavior and auction outcomes.

REFERENCES

Adam, Marc T. P., Jan Kräemer, Caroline Jäehnig, Stefan Seifert, and Christof Weinhardt. 2011. "Understanding auction fever: A framework for emotional bidding." *Electronic Markets* 21 (3): 197–207.

Avery, Christopher. 1998. "Strategic Jump Bidding in English Auctions." *Review of Economic Studies* 65 (2): 185–210.

Bajari, Patrick, and Ali Hortaçsu. 2003. "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from *eBay* Auctions." *RAND Journal of Economics* 34 (2): 329–55.

Banks, Jeffrey, Mark Olson, David Porter, Stephen Rassenti, and Vernon Smith. 2003. "Theory, experiment and the federal communications commission spectrum auctions." *Journal of Economic Behavior and Organization* 51 (3): 303–50.

Coppinger, Vicki, Vernon Smith, and Jon Titus. 1980. "Incentives and Behavior in English, Dutch and Sealed-Bid Auctions." *Economic Inquiry* 18 (1): 1–22.

Cramton, Peter C. 1997. "The FCC Spectrum Auctions: An Early Assessment." *Journal of Economics and Management Strategy* 6 (3): 431–95.

Cramton, Peter C., and Jesse A. Schwartz. 2000. "Collusive Bidding in the FCC Spectrum Auctions." *Journal of Regulatory Economics* 17: 229–52.

Daniel, Kent D., and David A. Hirshleifer. 1998. "A Theory of Costly Sequential Bidding." University of Michigan Business School Working Paper 98028.

Ding, Min, Jehoshua Eliashberg, Joel Huber, and Ritesh Saini. 2005. "Emotional Bidders: An Analytical and Experimental Examination of Consumers' Behavior in a Priceline-Like Reverse Auction." *Management Science* 51 (3): 352–64.

Easley, Robert F., and Rafael Tenorio. 2004. "Jump Bidding Strategies in Internet Auctions." *Management Science* 50 (10): 1407–19.

Ellison, Glenn, and Sara Ellison. 2009. "Search, Obfuscation, and Price Elasticities on the Internet." *Econometrica* 77 (2): 427–52.

Elyakime, Bernard, Jean-Jacques Laffont, Patrice Loisel, and Quang Vuong. 1994. "First-Price Sealed-Bid Auctions with Secret Reservation Prices." *Annales d'Économie et de Statistique* 34: 115–41.

Gabaix, Xavier, and David Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *Quarterly Journal of Economics* 121 (2): 505–40.

Grether, David M., and Charles R. Plott. 2009. "Sequencing strategies in large, competitive, ascending price automobile auctions: An experimental examination." *Journal of Economic Behavior and Organization* 71 (2): 75–88.

Grether, David, David Porter, and Matthew Shum. 2015. "Cyber-Shilling in Automobile Auctions: Evidence from a Field Experiment: Dataset." *American Economic Journal: Microeconomics*. http://dx.doi.org/10.1257/mic20120085.

Harrison, Glenn W., and John A. List. 2004. "Field Experiments." *Journal of Economic Literature* 42 (4): 1009–55.

Hörner, Johannes, and Nicolas Sahuguet. 2007. "Costly Signalling in Auctions." *Review of Economic Studies* 74 (1): 173–206.

[25] Of course, this evidence is only suggestive, as we do not directly observe the counterfactual of what prices cars won by the shill bidders would have fetched had they not participated. It could be that in the absence of the shills, these prices would have been much lower.

Hong, Hong, and Matthew Shum. 2003. "Econometric Models of Asymmetric Ascending Auctions." *Journal of Econometrics* 112 (2): 327–58.

Isaac, R. Mark, Timothy C. Salmon, and Art Zillante. 2007. "A Theory of Jump Bidding in Ascending Auctions." *Journal of Economic Behavior and Organization* 62 (1): 144–64.

Izmalkov, Sergei. 2004. "Shill Bidding and Optimal Auctions." http://economics.mit.edu/files/1085.

Jin, Ginger Zhe, and Andrew Kato. 2006. "Price, Quality, and Reputation: Evidence from an Online Field Experiment." *RAND Journal of Economics* 37 (4): 983–1005.

Katkar, Rama, and David H. Reiley. 2007. "Public Versus Secret Reserve Prices in eBay Auctions: Results from a Pokémon Field Experiment." *Advances in Economic Analysis and Policy* 5.

Kaufmann, Robert J., and Charles A. Wood. 2003. "Running up the Bid: Detecting, Predicting, and Preventing Reserve Price Shilling in Online Auctions." http://misrc.umn.edu/workingpapers/fullpapers/2003/0304_022003.pdf.

Kaufmann, Robert J., and Charles A. Wood. 2005. "The effects of shilling on final bid prices in online auctions." *Electronic Commerce Research and Applications* 4*:* 21–34.

Kosmopoulou, Georgia, and Dakshina G. de Silva. 2007. "The effect of shill bidding on prices: Experimental Evidence." *International Journal of Industrial Organization* 25 (2): 291–313.

Krishna, Vijay. 2002. *Auction Theory.* Burlington, MA: Academic Press.

Ku, Gillian, Adam D. Galinsky, and J. Keith Murnighan. 2008. "Arousal, Interest, and Auction Bidders." Paper presented at 21st Annual International Association for Conflict Management (IACM), Chicago.

Lucking-Reiley, David. 1999. "Using Field Experiments to Test Equivalence between Auction Formats: Magic on the Internet." *American Economic Review* 89 (5): 1063–80.

Malmendier, Ulrike, and Young Han Lee. 2011. "The Bidder's Curse." *American Economic Review* 101 (2): 749–87.

McCabe, Kevin A., Stephen J. Rassenti, and Vernon L. Smith. 1990. "Auction Institutional Design: Theory and Behavior of Simultaneous Multiple-Unit Generalizations of the Dutch and English Auctions." *American Economic Review* 80 (5): 1276–83.

Milgrom, Paul R. 2004. *Putting Auction Theory to Work.* Cambridge: Cambridge University Press.

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

Plott, Charles R., and Timothy C. Salmon. 2004. "The Simultaneous, Ascending Auction: Dynamics of Price Adjustment in Experiments and in the U.K. 3G Spectrum Auction." *Journal of Economic Behavior and Organization* 53 (3): 353–83.

Tadelis, Steve, and Florian Zettelmeyer. 2009. "The Effect of Information on Auction Outcomes: A Large-Scale Field Experiment." https://escholarship.org/uc/item/7t4834h2.

Vincent, Daniel R. 1995. "Bidding Off the Wall: Why Reserve Prices May Be Kept Secret." *Journal of Economic Theory* 65 (2): 575–84.