Is the 'Linkage Principle' Valid?: Evidence from the Field*

Sung-Jin Cho
School of Economics, Seoul National University
Harry J. Paarsch
Department of Economics, University of Melbourne
John Rust
Department of Economics, Georgetown University

November 1, 2012

First Draft: 19 December 2009

Abstract

We present field evidence concerning experienced bidders that supports the linkage principle—specifically, the prediction that in affiliated-values auction environments the expected revenues generated at open-outcry, ascending-bid auctions are higher than those under auction formats that reveal less information to participants. Using field data from a large seller of automobiles who experimented with different selling formats, we have found that average revenues were significantly higher under an English auction than under a dynamic Internet auction format that revealed less information to bidders.

Key words: revenue comparisons; auction choice; linkage principle; used-car auctions.

JEL classification: C14, D44, L1.

*The authors wish to acknowledge the generous financial support of the National Research Foundation of Korea under grant NRF-2010-200-20100055-B00005 through the Global Research Project and to thank Matthew Chesnes and David McArthur for exceptional research assistance. They are also grateful to John Asker, Srihari Govindan, Han Hong, Timothy P. Hubbard, Vijay Krishna, Dan Levin, David P. Porter, Benjamin S. Skrainka, and Robert B. Wilson as well as three anonymous referees for providing helpful comments and useful suggestions on earlier drafts of the paper. Paarsch gratefully acknowledges that some of the research for this paper was completed while he was a visiting research scholar at the Center for Economic Institutions in the Institute of Economic Research at Hitotsubashi University in Kunitachi, Japan as well as a visiting fellow at the Collegio Carlo Alberto in Torino, Italy.

1 Introduction

What practical insights can economists provide concerning how to structure auctions as well as how to bid at them? In the workhorse model of auction theory, which was first developed by Vickrey [1961], each of a known number of potential bidders draws an individual-specific random valuation independently from the same distribution. In Vickrey's model, the specific value of his draw is that bidder's private information; it represents the monetary value of the object to him. Economic theorists refer to this model as the symmetric *independent private-values paradigm* (IPVP) because the draws are independent and the valuations are bidder specific. Also, because each potential bidder has an identical chance of getting any specific draw before the valuations are drawn, the bidders are *ex ante* symmetric.

Different auction formats (open-outcry versus sealed-bid) and different pricing rules (pay-your-bid versus second-price) provide potential bidders with different incentives concerning how to bid. For example, under the pay-your-bid pricing rule, a bidder's action (his bid) determines what he pays should he win, while under the second-price rule, the action (bid) of his nearest rival determines what the winner pays.

In equilibrium, different functions map the private information of participants (their values) into their actions (their bids). For example, open-outcry (sometimes referred to as *oral*) auctions can be conducted in at least two different ways. In the first, the price is set very low, perhaps at zero, and then allowed to rise more or less continuously until only one participant remains active in the auction. That remaining active bidder is the winner, and he pays what the last other active bidder was willing to pay, often plus a small increment. Economic theorists have typically chosen to model these oral auctions as clocks, where the price rises continuously with the movement of a clock hand. In this case, the winner of the auction is the participant with the highest valuation and he pays what his nearest rival (that participant with the second highest value) was willing to pay. Thus, the oral, ascending-price auction guarantees the efficient allocation of the object: the participant with highest valuation wins the auction. Such an auction is sometimes referred to as a *second-price* auction because, in the absence of bid increments, the winning bid is the second-highest bid, which happens to be the second-highest valuation as well. In economics, this outcome has special meaning because the second-highest valuation represents the opportunity cost of the object for sale—its value in its next best alternative. As a technical aside, the equilibrium at an oral, ascending-price auction (sometimes referred to as an *English* auction) has a special structure: it is a dominant strategy equilibrium; each participant has the incentive to reveal his private information, to tell the truth concerning his value by continuing to bid up to his value, regardless of what his rivals do.

In the second form of oral auction, the price is set very high, and then allowed to fall continuously; the winner is the first participant to cry out a bid, and he pays his bid. In practice, these oral auctions are often implemented using a clock, where the hand (or a digital panel) lists the current price. Participants affirm their willingness to pay the current price by pushing a button which stops the clock at that price. These auctions are often referred to as *Dutch* auctions, perhaps because the format has been used frequently in the Netherlands to sell both fish and flowers. Because the price that the winner pays is related to his action (crying out or pushing the button), he has an incentive to shave his bid, to wait longer before pushing the button to stop the clock. The equilibrium at the Dutch auction is not a dominant strategy equilibrium, but rather a Bayes–Nash equilibrium, which is a much stronger form of equilibrium. While the Bayes–Nash equilibrium bid function is an increasing function of a bidder's value, it has a slope that is typically less than one: each bidder is deceptive when bidding; he does not tell the truth, but rather bids less than his value. Again, however, the winner is the participant with the highest valuation, so objects are allocated efficiently at Dutch auctions.

Under the assumption that participants are risk neutral with respect to winning the object for sale, a remarkable result obtains: expected revenue equivalence. That is, if the same good sold were auctioned under the two different institutions, then the average winning bid at the English auction would equal the average winning bid at the Dutch auction. To most people, this expected revenue equivalence result is at first somewhat surprising because considerable information is revealed during the course of bidding at English auctions, whereas at Dutch auctions no information is revealed until the winner has been determined. Within the IPVP, however, information plays no extra role in determining the average winning price because each bidder's private information (his value) is, by assumption, statistically independent of the private information of his rivals (their values): knowing something about the values of his rivals provides no extra information to a bidder concerning his own valuation. Thus, no bidder at an English auction can learn anything more about his valuation from the actions (bids) of his rivals. Once one realizes this fact, the equivalence of average winning bids is clear: at a Dutch auction, assuming he wins because he has the highest value, a representative participant forms his bid so that he will, on average, just beat his nearest rival, the bidder with the second-highest valuation.

Similar analyses have been performed for the sealed-bid format under different pricing rules. In fact, game theorists have shown that sealed-bid auctions at which the highest bidder wins the auction and pays what he bid are strategically equivalent to Dutch auctions. Consequently, the Bayes–Nash equilibrium bid function at a sealed, pay-your-bid, auction is identical to that at a Dutch auction. It has also been shown that sealed-bid auctions at which the highest bidder wins the auction, but pays the bid of his nearest rival, are strategically equivalent to English auctions. Under the assumption of risk-neutral participants, expected revenue equivalence follows. That is, if the same good sold were sold under the two different institutions, then the average winning bid at a sealed, pay-your-bid auction would equal the average winning bid at a sealed, second-price (also known as *Vickrey*) auction.

This result is the celebrated *Revenue Equivalence Theorem* (RET), which was first outlined by Vickrey [1961, 1962] and then proven by Riley and Samuelson [1981] as well as Myerson [1981]. In its full generality, the RET states that any auction format that has the same probability of assigning a winning bidder generates the same expected revenue to the seller. In particular, the RET predicts that the expected revenues earned by the seller at sealed-bid auctions will be the same as those earned at English auctions, at least for one-shot, single-object auctions when the distribution from which the values are drawn is the same for all potential bidders, who are also risk-neutral.

While economic theorists have thoroughly investigated relaxing each of the assumptions required for the RET, allowing the symmetric bidders to have valuations that are dependent is perhaps the most interesting case. When the valuations of bidders are dependent, the revelation of private information through bidding can be important to the equilibrium outcome. Specifically, the winning bids at English auctions are more informative than those at sealed-bid or Dutch auctions because considerably more information is revealed during the course of bidding at English auctions; see, for example, the work of Pesendorfer and Swinkels [2000]; Hong and Shum [2004] as well as Hong et al. [2009]. In order to construct equilibria to auction games under dependence, economic theorists have been forced to put a specific structure on the dependence. Following the path blazed by Karlin [1968], mathematicians refer to this structure as *multivariate total positivity of order two*, or MTP_2 for short, while in an influential and classic paper, Milgrom and Weber [1982] coined the term *affiliation* to describe this form of dependence.

Under symmetric affiliation, Milgrom and Weber derived a powerful result and coined the term *linkage principle* to describe it. In single-object auction models where the signals of the risk-neutral potential bidders are symmetrically affiliated, the linkage principle states that a seller can expect to increase revenues by providing more information to bidders, both before and during the auction. An implication of the linkage principle is that English auctions will, on average, earn more revenue for the seller than sealed-bid auctions, under which no information is released, or similar auction formats that reveal less information to potential buyers: the RET breaks down. According to Perry and Reny [1999], "the linkage principle has come to be considered one of *the* fundamental lessons provided by auction theory."

Thus, the presence of some degree of dependence, or a *common-value component*, in the signals of potential buyers is critical to the validity of the linkage principle. The affiliated-values model is a generalization of the common-value model developed by Wilson [1977], and nests the IPVP. Under affiliation, the conditional expectation of any monotonic function of the signals of all bidders is an increasing function of any individual bidder's own signal. When the signals of bidders are dependent in this manner, information released by the seller or information the seller provides concerning the bids made by other participants (by virtue of the seller's choice of auction format) helps bidders refine their beliefs concerning the true value of the object for sale, which in turn induces them to bid more aggressively than they would in the absence of such information.

As one might expect, the analysis of auctions at which several units of the same object are sold sequentially is complicated. Within the symmetric IPVP, when potential bidders have single-unit demand, Weber [1983] has demonstrated that the equilibrium price path under the four combinations of auction formats and pricing rules follows a martingale; when potential bidders have multi-unit demand (which follows a Poisson process), Donald et al. [2006] have demonstrated that the equilibrium price path follows a supramartingale—on average, the equilibrium price rises over consecutive auctions. To our knowledge, only Perry and Reny have investigated the effect of affiliation in multi-unit auctions; in fact, they have provided a counter-example that demonstrates the Milgrom–Weber ranking breaks down in multi-unit auctions with affiliation. Of course, auctions at which several objects sold sequentially are even more difficult to analyze than the multi-unit case.

Another reason why the linkage principle can fail is when bidders cooperate, that is, collude. In general, collusion is easier to sustain in environments that are rich in information: more information is released at English auctions than at sealed-bid ones, or other less open auction formats.

In any case, while one can imagine circumstances under which the release of information could adversely affect the outcome at an auction (for example, if the seller released information concerning problems with the object for sale, or when low bids by some bidders convince other bidders that the item is worth less than they originally thought), the remarkable feature of the linkage principle is that, *ex ante*, providing more information raises the expected revenues to the seller. Milgrom and Weber [1982] have summarized the implications of the linkage principle succinctly: "honesty is the best policy."

To our knowledge, the specific implication of the linkage principle under the Milgrom–Weber assumptions outlined above—namely, that English auctions should, on average, generate higher revenues than sealed-bid ones, or other less open auction formats—has never been subjected to a direct empirical test, at least not using data from "the field." All of the empirical tests that we know of have been conducted using controlled laboratory experiments. In an important series of papers, Kagel and Levin [1986] and Levin et al. [1996] analyzed the behavior of laboratory subjects at English and sealed-bid auctions in situations where a common-value component existed in their experimentally-generated values.

The results of these experiments, which have been summarized by Kagel and Levin [2002], are mixed. For relatively inexperienced subjects, they found a pronounced "winner's curse" caused by overbidding at sealed-bid auctions relative to English ones. On average, the overbidding caused the seller's revenues to be higher at sealed-bid auctions than at English auctions, contrary to the prediction of the linkage principle. In experiments involving experienced bidders, however, the winner's curse was ameliorated and the English auctions generated higher expected revenues than the sealed-bid ones, a finding consistent with the linkage principle.

Below, we present an empirical analysis of (uncontrolled) field experiments conducted by a large rental car company that sells hundreds of unwanted, used cars each month.[1] The seller is obviously quite interested in adopting a selling mechanism or an auction format that maximizes the revenues it can earn from the sale of its unwanted inventory of used cars.

While there are certainly individual-specific, private-value components in any automobile purchase ("I *really* want that pink *Cadillac* over there, you know, the one with the cream leather seats, because . . . "), common-value elements must surely exist, too. Specifically, a pre-owned vehicle's true quality is uncertain because the intensity with which it has been used and the care shown it by previous drivers are unknown. This unknown quality is basically the same to all potential buyers, but will remain undiscovered until the vehicle has exchanged hands and the new owner has experienced it on the road. In short, we do not think it unreasonable to assume affiliation among the signals or the valuations of potential buyers of used cars.

At any given point in time, the rental car company's fleet contains more than 30,000 vehicles; over the past decade, the company has sold approximately 400 vehicles each month. During this period, the company has sold used cars under several different selling mechanisms. We focus on two: first, the rental car company conducted computerized Internet auctions held in cyberspace; second, the rental company hired a large auction house to conduct oral, ascending-price auctions at a central location. We refer to these two different methods of selling

[1]The rental car company which provided us with the data has requested that it remain anonymous. In addition, we are restricted from providing information that could identify the company as well as any individual vehicles, customers, or bidders.

used cars as *sales regimes*, or regimes for short. In our empirical analysis, we attempt to determine which of the two sales regimes yielded the highest average revenue to the rental car company.

Prior to developing its own specialized Internet auction software, the rental car company had sold most used cars at oral, ascending-price auctions conducted at individual car rental outlets; in addition, a relatively few used cars were sold directly to individual customers after informal bilateral bargaining. In 2002, however, the rental car company began to suspect that collusion among some participants at some of its English auctions. The rental car company then invested in developing a unique auction format for selling used cars online. The participants under this Internet auction were strictly anonymous. Over the course of an Internet auction, which was two minutes in duration, an individual bidder would only see a single piece of information: whether his bid was the highest competing bid at the auction. Participants could not observe the bids of their opponents. In fact, an individual bidder did not even know what the highest bid was at any time during the auction, unless the bidder himself had the current highest bid.

By 2007, the volume of vehicles sold at its Internet auctions was so large that the enterprise began to occupy too much of its managers' time; management began to regard the Internet auctions as a distraction from their main business—renting cars. Thus, the company decided to contract with a large, prominent auction house to sell the used cars. The auction house employed an oral, ascending-bid auction that was virtually identical to an economist's notion of an English auction. In particular, unlike the company's Internet auctions, a bidder at an English auction conducted by the auction house could see the other participating bidders as well as their bids at each stage in the auction, including the highest bid at any point in the auction. The auction house charged a variable commission rate for its services, but the average commission rate was approximately ten percent of a vehicle's gross selling price.

We have analyzed empirically the traded prices received by the rental car company (including prices net of commission in the case of sales by the auction house) during the period 2003 to 2008 under the two different sales regimes. We have compared revenues for specific vehicle classes and individual makes/models of vehicles for which we have the largest number of observations. Although the company sells a large number of vehicles in total each month, the numbers of vehicles sold for specific makes and models are insufficient to employ a "regression discontinuity" approach where net revenues for specific makes/models are compared just before and just after the transition from one sales regime to another, such as the transition from the company's Internet auctions to sales through the auction house, which began on 1 January 2008.

Instead, we have averaged prices over the much larger numbers of vehicles sold during entire sales regimes, not just the much smaller numbers of vehicles sold around sales-regime transitions. We justify this approach by noting that, during the period of our analysis, there were no sigificant "macro shocks" or inflation in the used-car market in the country where the rental company operates, which we document in section 3. In addition, no significant changes occurred in the engine or other features and characteristics of the specific car models we analyzed. Thus, we feel we can rule-out these explanations for the significant shifts in prices across different sales regimes. In short, we believe that a simple comparison of average prices received for specific high-volume vehicle makes and models provides an appropriate basis for measuring the effect of the sales regime and selling mechanism on revenues earned by the seller.

In general, our empirical findings are consistent with the prediction of the linkage principle. Specifically, comparing traded prices for mid-sized vehicles under the two main sales regimes, where the vast majority of our observations exist (the company's own Internet auctions and the English auctions conducted by the auction house), *net* revenues earned by the rental car company were, on average, significantly higher at the English auction than at the Internet auctions that released less information.

We also found, however, that when we analyzed specific makes/models of cars (for example, we considered three for which we have the largest number of observations) the rankings of the two sales regimes differed across the three models. For car Model A (again the specific make/model has been suppressed due to confidentiality restrictions imposed by the rental car company), the average price earned was higher at the English auctions conducted by the auction house (again, net of commmission) than at the Internet auctions. For car Model B, the average revenues under Internet auctions and the English auctions conducted by the auction house were about the same and not significantly different from one another. For car Model C, the average net revenues earned at the English auctions conducted by the auction house were significantly greater than those at the Internet auctions; the difference was statistically significant at conventional p-values.

Overall, our findings support the conclusion that the oral, ascending-price auction earned the highest average net revenues for the rental car company, evidence consistent with the predictions of the linkage principle. Another possible explanation for the higher average revenues from the auction house could, however, be demand aggregation: the auction house may have succeeded in attracting more bidders than were present under the Internet auctions. We do not believe that the significantly greater average sales prices under this sales regime can be explained by a larger number of potential bidders at the auction-house auctions. In fact, based on other evidence presented by Kim and Lee [2008], we believe that the same set of potential buyers participated at both auctions. In the conclusions, we discuss this evidence in detail.

Perhaps the single most important message to take from our analysis is the following: as counselled by Milgrom and Weber [1982] as well as Ausubel [2004], information release is an important feature in auction design. Consistent with the prediction of the linkage principle, the average traded price of vehicles was significantly higher at the oral, ascending-price auctions conducted by the auction house than the closed Internet auction implemented by the rental car company. The Internet auctions may have been successful in thwarting the collusion potentially present at the English auctions conducted by the rental car company at each of its car rental outlets. If there were any collusion by participants at the English auctions conducted by the auction house, then it does not appear to have been successful because the average prices are the highest under this sales regime.

We believe our findings are significant because they represent the first empirical test of the linkage principle that we know of using field data concerning experienced bidders. Our findings are consistent with the evidence found by Kagel et al. [1987] concerning experienced bidders in laboratory experiments. After we completed this paper, we became aware of a paper by Tadelis and Zettelmeyer [2010], who reported results from a controlled experiment conducted at a different rental car company and designed to test a different implication of the linkage principle—namely, whether the *ex ante* release of information concerning the mechanical conditions and repair histories of vehicles being sold at wholesale automobile auctions increased the average traded price. Tadelis and Zettlemeyer found that this information release did increase average traded prices, which is also consistent with the linkage principle. In their research, however, they did not undertake experiments that show the effect of the selling mechanism on average traded prices, the main contribution of this paper.

Because our data are from dynamic, multi-object auctions, the reader might reasonably ask whether we can really learn anything about the linkage principle from this empirical exercise. Had we rejected the linkage principle because the expected revenues at pay-your-bid Internet auctions were greater than those at English auctions, then that evidence could have been a rejection of the hypothesis of affiliation, a rejection of the hypothesis of risk neutrality, a rejection of the hypothesis of the single-object auction, or rejections of any combination of the three hypotheses. In short, we would not have learned that much. On the other hand, because the expected revenues at English auctions were significantly greater than those at the pay-your-bid Internet auctions, such evidence provides fairly strong support for the hypothesis that information is important at auctions because both risk aversion and multi-object demand could have overturned the affiliation hypothesis.

The remainder of our paper has three sections: in section 2, we describe in some detail the four sales regimes as well as the data, while in section 3, we presents a summary of our empirical analysis and, in section 4, we conclude.

2 Data

During the period for which we have data, from the last quarter of 2002 onward, we examined two different sales regimes. For parsimony, we refer to them in order as Regimes 1 and 2, respectively. We provide summary descriptions of the sales regimes in table 1.

At the beginning of 2003, the company implemented Regime 1, which involved conducting electronic auc-

Table 1: Description of Sales Regimes

Sales Regime	Description
Regime 1	Internet auctions conducted in cyberspace by rental car company
Regime 2	English auctions conducted by auction house at large central site

tions over the Internet. These electronic auctions were held at pre-announced times each month; vehicles were sold one at a time in a particular order over the Internet at auctions lasting exactly two minutes each.[2] At these auctions, a potential buyer could submit as many bids as he liked. The only information available to any participant, however, was whether he was the highest bidder. Specifically, none of the participants knew how many bidders were active at the auction. Because of these institutional features, unlike at the electronic auctions used by eBay, it is virtually impossible to snipe effectively: except for the current highest bidder, none of the other participants knew what the current price was, so only a lucky sniper could sneak in at the last second to "steal" a vehicle away from the existing highest bidder. In fact, we found no evidence of the "last-minute" behaviour noted by Roth and Ockenfels [2002], which in our case would be the last ten seconds, or so. At the end of the auction, the highest bidder won, and paid what he bid. Thus, the pricing rule at these auctions was pay-your-bid rather than second-price.

By eliminating a public forum in which signals could be discreetly exchanged and in which cooperative behavior could be monitored (and, thus, potentially enforced) by the colluders, the rental car company believed it could thwart uncompetitive behavior among the potential buyers. What made these pay-your-bid Internet auctions different from other sealed, pay-your-bid auctions typically used is that a bidder could, by trial-and-error, discover what the highest current tender was. By allowing participants to increase their bids sequentially, some information release was permitted, unlike in the models of standard pay-your-bid auctions studied by Milgrom and Weber [1982].

The company also restricted who could participate at the Internet auctions. In particular, in our dataset, 124 unique bidder identification numbers exist that map to specific used-car dealers. These dealers were obvious resellers of pre-owned vehicles: historically, they had purchased many vehicles from the rental car company, solely for the purpose of resale.

Because the Internet auctions were electronic, data collection was relatively easy. In principle, we have access to virtually every piece of relevant information concerning the auctions; in practice, missing odometer readings and other factors make some of the data incomplete. Also, the company was unable to provide us access to any transaction data for a three-month period in 2004. We do not believe there is any hidden agenda here: the most likely explanation is that the data were simply lost in a computer crash.

In figure 1, we present a graph of the sequence of bids observed at a representative Internet auction. There were nine bidders participating at this auction, which lasted two minutes. The solid line plots the highest bid received at each instant, and the various symbols plot the actual bids submitted by the nine bidders. Three of the bidders—6, 8, and 9—tendered only a single bid at the auction. Bidder 6 submitted the highest bid $20,000$ at the 77.953-second mark of the auction. This remained the highest bid for the remainder of the two minute auction and, consequently, bidder 6 won the auction and paid $20,000$.

The reader will note the large number of "dominated" bids being placed at this auction; this presumably occurred because of the limited information that the auctioneer provides to the bidders. As we noted, no bidder can observe either the number of other bidders or the bids they have placed in the auction: the only information a bidder observes is whether his bid is the highest. Consequently, we see obvious "testing strategies" being used by the other six bidders, who gradually increased their bids in an attempt to become the high bidder and possibly also to learn what the high bid was at that moment of the auction. It is, however, evident that several bidders never succeeded in learning what the high bid was since their bids were always below the high bid at the auction. Examples include bidders 2 and 3, whose bids are plotted as circles as well as five-pointed stars, respectively, in

[2] In practice, the time stamps in the electronic files document that some of the auctions were, in fact, as long as 121 seconds, but we believe that this heterogeneity is unimportant.

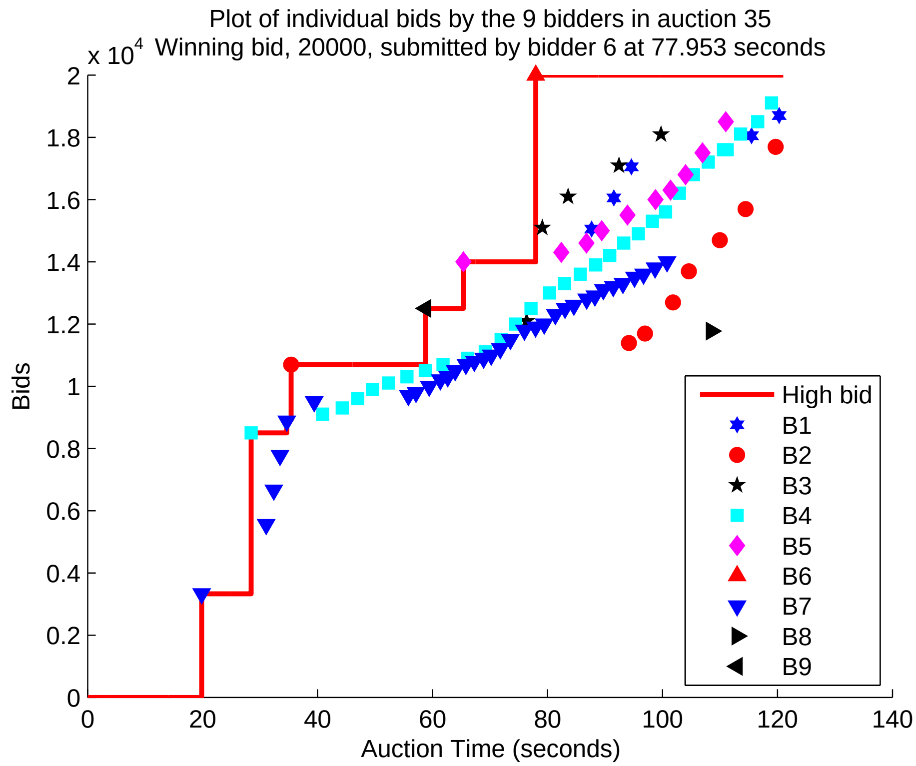

Figure 1: Sequence of Bids Observed at a Representative Internet Auction

figure 1.

Clearly, the Internet auction provides more information to bidders than what they would receive at one-shot sealed-bid auction. In particular, a bidder can start out with low bids and increase them gradually in attempt to learn what the high bid is. But, as we see in figure 1, this strategy is not always successful. In fact, most of the bidders who won most frequently at the Internet auctions placed only a small number of bids, often just a single bid.

It is also clear that the information provided to bidders at the Internet auction is less than what they would observe at an English auction, such as the auctions conducted by the auction house where all bidders see all bids placed by other bidders, including the winning bid. Furthermore, bidders can also potentially know the identities of the competing bidders because they are physically present and are calling out their bids on the auction floor. Thus, the information provided to bidders at the Internet auction is greater than the information provided at a sealed-bid auction, but less than the information provided at an English auction. If bidders do not collude and their values are affiliated, then the linkage principle predicts that the English auction should generate higher expected revenues to the seller than the Internet auction, and the Internet auction should generate higher expected revenues than a sealed-bid auction, at least under the Milgrom–Weber assumptions.

Unfortunately, the rental car company did not sell any of its vehicles at sealed-bid auctions. Consequently, we are unable to test the latter implication of the linkage principle. Our intuition, however, is that the value of using "testing strategies" and attempting to learn the value of the high bid is limited in these fast-moving auctions. We conjecture that the Internet auctions are "strategically close" to sealed-bid auctions in the sense that expected revenues are not much higher than those that would obtain at sealed-bid auctions. In separate work, we plan to test this conjecture by solving a model of equilibrium bidding strategies at the Internet auction and comparing expected revenues to those that arise at a sealed-bid auction. To our knowledge, the Internet auction used by the rental car company is a unique auction format which has never been analyzed either theoretically or empirically in the previous literature concerning auctions.

By 2007, conducting the electronic auctions had become a distraction to managing the company and management sought to exit this business by hiring an auction house to conduct the sales on its behalf.[3] In return, the firm selling the vehicles would receive a commission that varied according to the make of vehicle; the average commission rate was about ten percent of the gross sales price. We refer to this period, which began on 1 January 2008, as Regime 2.

The auction house chose to sell the vehicles using a selling mechanism that used-car dealers know best—the English auction, an oral, second-price auction. Because the auction house's commission is proportional to total sales revenue, it presumably had an incentive to design the auctions well. Although the rental car company has been quite generous with providing us information and in answering our questions, we have no close relationship with the auction house. One of the authors has attended several Regime 2 auctions. From this field research, we saw no obvious differences from other English auctions used to sell pre-owned vehicles.

The information gathered under Regime 2 is quite different. Under its contract with the rental car company, the auction house is only required to report the date and time of an auction as well as the winning bid received for each vehicle sold as listed on a manifest. We know from the auction house that the potential buyers under Regime 2 are essentially the dealers who participated under Regime 1. We have learned from the research of Kim and Lee [2008] that, while private citizens can *sell* vehicles at the auction-house English auctions, only registered dealers can *purchase* vehicles at these auctions. Evidence from the paper of Kim and Lee corroborates this claim. Specifically, in our database, we have 124 distinct buyer identification numbers, while Kim and Lee report 134 distinct buyer identification numbers. Presumably, both the rental car company and the auction house excluded private buyers because it would have been an administrative nightmare to deal with a large number of potentially inexperienced bidders.[4]

We organized all of the data concerning the $30,621$ sales that we acquired from the rental car company into a dataset. Because different amounts and kinds of information were generated under the two sales regimes, in making empirical comparisons between the two, we are constrained by the least-complete data-collection scheme. Specifically, the only information we have that is comparable across all of the selling regimes is the following:

1) date of sale;
2) vehicle identification number;
3) vehicle model;
4) vehicle age;
5) purchase price of vehicle;[5]
6) sale price of vehicle;
7) type of sale;
8) identification number of the winner for Regime 1.[6]

For some vehicles, we have an odometer reading for the vehicle and know whether that vehicle has been in an accident, but these data are unavailable for *many* vehicles; nearly fifteen percent of the vehicles sold have missing odometer readings. The information concerning accidents is reliable for around fifteen percent of the

[3]This is that same auction house from which Kim and Lee [2008] procured data for their analysis of used-car auctions, although their data concern different vehicles from ours.

[4]As we shall describe below, even some of the experienced participants made costly errors, albeit infrequently.

[5]For around 0.15 percent of the vehicles in the dataset, the initial purchase price is unknown.

[6]For around five percent of these sales, the identity of the winner is unknown. Under Regime 2, the winner is simply listed as 16, the firm who conducted the auctions, rather than the actual winner.

observations. Put another way, if we constrain ourselves to observations that have complete mileage and accident histories, then the remaining samples are extremely small.

In principle, under Regime 1, we should know the complete bidding histories of all participants, but no such information exists for Regime 2. At none of the auctions did a reserve price exist. None of the vehicles went unsold. Under Regime 1, however, some bidders made errors: infrequently, a bidder made a keystroke error, which resulted in his winning the auction at a ridiculously high price—for example, several hundreds of thousands of dollars for a vehicle worth less than ten thousand dollars. At the close of the auction, this mistake was realized. At this point, the company, voided the sale, and resold the vehicle at a later auction. The practical importance of such cases is likely very small.

3 Empirical Results

While we have data concerning the sales of nearly $31,000$ vehicles, most of these data are not strictly comparable with one another. In addition, as was alluded to above, trying to control for differences in observed covariates collected across each of the regimes is difficult because different types of information were gathered under the two regimes. For example, in principle, every action of every bidder was recorded under Regime 1, but only the winning bid is reported under Regime 2. Moreover, while under Regime 1 we know the entire set of potential bidders and actual participants, under Regime 2 we do not even know the identity of the winner, let alone the other participants: under Regime 2, the winner is always listed as the auction house.

Also, in order to avoid the potential biases that can arise when, for example, comparing the sale of a *Mercury Sable* with the sale of a *Jeep Cherokee* (namely, comparing apples and oranges), we have chosen to focus on relatively homogeneous products. Of course, there are limits to how fine we can go; these limits are largely determined by the information provided us by the rental car company concerning models. Note, too, that by restricting ourselves in this way, we have also reduced the potential samples sizes in our analysis: we must trade-off decreased bias with increased sampling variation.

Over thirty-four percent ($10,422$ of $30,621$) of the sales in our dataset involved mid-sized vehicles of various models sold under either Regime 1 or Regime 2. Thus, we focused on those first.

In the top four rows of table 2, we report the sample descriptive statistics for mid-sized vehicles under the two regimes. Switching from Internet auctions of Regime 1 to the English auctions of Regime 2 made profits for the rental car company; see the column labelled “Mean” for the sample averages. In real terms, the average traded price rose around 4.72 percent.[7] Of course, we do not know what it costs to run either of these auctions, but a nearly five percent improvement is substantial, and impressive. As was noted above, however, the rental car company pays the auction house a commission for conducting the auction, which averages out to be around ten percent of the gross revenues. It is important to note that the price data we received from the auction house are *net* of that commission, the rate of which varies from vehicle to vehicle. Thus, under Regime 2, the rental car company does not have to incur selling costs, such as those incurred when running the Internet auctions under Regime 1: all auction-related costs under Regime 2 are borne by the auction house. In short, while this increase in prices is relatively small, it is a lower bound on the profit that the rental car company made by switching auction formats and pricing rules.

While these differences are obviously economically important, the question of whether any one is statistically significant remains. Conventional standard errors for the sample means can be calculated using the information provided in the table; that is, simply divide the reported standard deviation (labelled “St.Dev.”) under each regime by the square root of the sample size (labelled “No.Obs.”) reported for that regime to get the standard error. We also calculated the asymptotic test statistic for the pair-wise difference: the p-value is 0.001, a difference which is unlikely to be the result of sampling error.

The main point to take from this part of the analysis is the following revenue ranking: English Auction > Pay-Your-Bid Internet Auction. Thus, at least at this granularity, the field evidence is consistent with the

[7] We made the CPI 1.00 for July 2005, around the midpoint of our sample.

Table 2: Sample Descriptive Statistics—Mid-Sized Vehicles, Broken Down by Regime and Model

Cut of Data	Regime	Variable	Mean	St.Dev.	L.Q.	Med.	U.Q.	No.Obs.
Mid-Sized, Main	1	Traded Price	7,255	2,189	6,012	7,183	8,280	6,214
Mid-Sized, Main	1	Age (in days)	1,080	179	972	1,141	1,178	6,214
Mid-Sized, Main	2	Traded Price	7,606	2,020	6,574	7,770	8,787	4,208
Mid-Sized, Main	2	Age (in days)	1,121	135	1,041	1,102	1,147	4,208
Mid-Sized, Sub.	1	Traded Price	7,352	2,107	6,049	7,190	8,317	4,557
Mid-Sized, Sub.	1	Age (in days)	1,069	176	962	1,064	1,138	4,557
Mid-Sized, Sub.	1	Odometer (miles)	49,172	23,260	36,108	46,894	60,242	4,557
Mid-Sized, Sub.	1	Purchase Price	15,555	3,666	13,673	15,004	16,752	4,557
Mid-Sized, Sub.	2	Traded Price	7,631	2,035	6,609	7,799	8,840	3,759
Mid-Sized, Sub.	2	Age (in days)	1,120	135	1,041	1,102	1,147	3,759
Mid-Sized, Sub.	2	Odometer (miles)	53,245	25,220	37,399	51,466	64,697	3,759
Mid-Sized, Sub.	2	Purchase Price	17,149	3,173	15,413	17,430	19,061	3,759
Mid-Sized, A	1	Traded Price	6,980	2,064	5,559	6,804	8,108	2,023
Mid-Sized, A	1	Age (in days)	1,073	167	972	1,074	1,138	2,023
Mid-Sized, A	1	Odometer (miles)	50,800	20,689	37,655	48,578	61,239	2,023
Mid-Sized, A	1	Purchase Price	15,538	4,691	13,298	14,437	16,672	2,023
Mid-Sized, A	2	Traded Price	7,615	1,797	6,895	7,749	8,570	1,563
Mid-Sized, A	2	Age (in days)	1,127	118	1,064	1,111	1,148	1,563
Mid-Sized, A	2	Odometer (miles)	55,486	28,488	40,336	53,780	65,803	1,563
Mid-Sized, A	2	Purchase Price	17,309	2,632	15,801	17,247	18,660	1,563
Mid-Sized, B	1	Traded Price	7,937	1,816	6,830	7,569	8,724	1,937
Mid-Sized, B	1	Age (in days)	1,094	176	983	1,081	1,154	1,937
Mid-Sized, B	1	Odometer (miles)	48,789	25,937	35,232	46,601	59,652	1,937
Mid-Sized, B	1	Purchase Price	15,633	2,196	14,121	15,179	16,814	1,937
Mid-Sized, B	2	Traded Price	7,609	2,246	6,667	7,990	8,996	1,219
Mid-Sized, B	2	Age (in days)	1,144	158	1,052	1,105	1,168	1,219
Mid-Sized, B	2	Odometer (miles)	53,696	19,542	40,517	52,816	65,280	1,219
Mid-Sized, B	2	Purchase Price	17,351	3,566	15,491	17,719	19,342	1,219
Mid-Sized, C	1	Traded Price	5,409	1,638	4,703	5,521	6,352	242
Mid-Sized, C	1	Age (in days)	1,012	165	915	989	1,105	242
Mid-Sized, C	1	Odometer (miles)	53,441	20,769	40,072	50,768	64,958	242
Mid-Sized, C	1	Purchase Price	14,503	3,169	13,020	13,905	15,312	242
Mid-Sized, C	2	Traded Price	5,612	1,196	5,127	5,698	6,340	117
Mid-Sized, C	2	Age (in days)	1,115	143	1,009	1,071	1,148	117
Mid-Sized, C	2	Odometer (miles)	70,088	34,624	46,297	62,137	88,316	117
Mid-Sized, C	2	Purchase Price	12,545	2,085	10,798	11,690	14,658	117

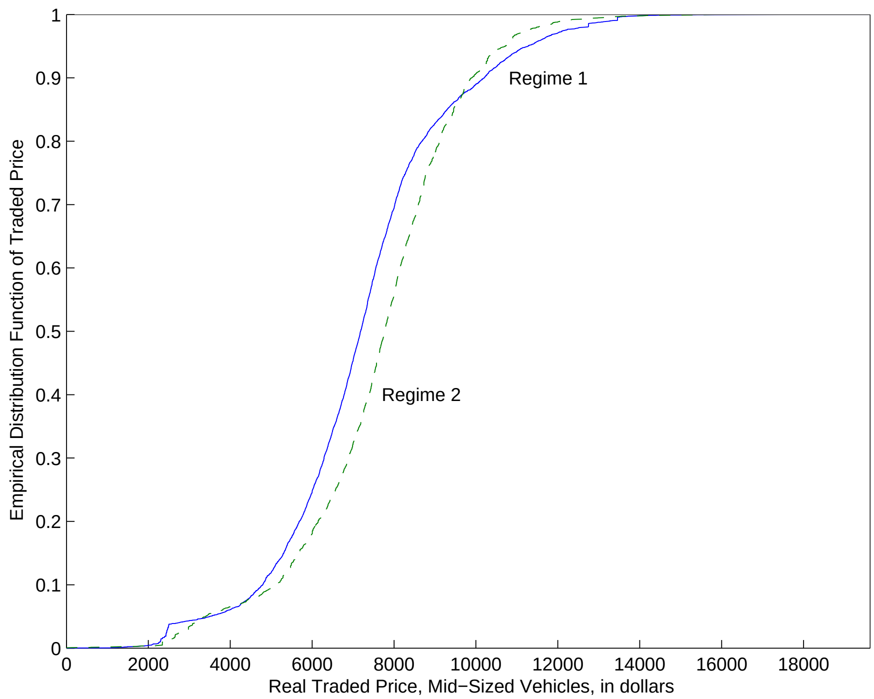

Figure 2: Empirical Distribution Functions of Traded Prices—Mid-Sized Vehicles

linkage principle. As Milgrom and Weber [1982] as well as Ausubel [2004] have counselled, information release matters.

In figure 2, we depict the empirical distribution functions (EDFs) of traded prices under the two regimes. Except at the very top end, above about the 85$^{\text{th}}$ percentile, the EDF of Regime 2 is to the right of that of Regime 1. When, however, Milgrom and Weber [1982] used the linkage principle to prove the revenue ranking of the auction formats and pricing rules, they did not characterize the effect that the formats and rules have on the distributions of traded prices, just the averages of traded prices.

We note, however, that in single-object models, within the symmetric IPVP, with risk-neutral bidders, the RET holds. In addition, the distribution of winning bids at pay-your-bid auctions and that at second-price auctions can be ranked in terms of second-order stochastic dominance. The latter involves a mean-preserving spread of the former. Within the Milgrom–Weber model, with affiliated signals, we know of no formal result along these lines. Nevertheless, under affiliation, the right tail of the winning bid distribution at a second-price auction is likely longer than that at a pay-your-bid auction, suggesting an inconsistency between the data and the theory.

Of course, the reader may worry that contamination, in the form of mis-reported traded prices or mis-classified vehicles, could affect our empirical results because, as an estimator, the sample mean has a very low breakdown point; see, for example, Belsley et al. [1980] as well as Huber [1981]. Contamination also has implications for what can learned from the data, as was noted out by Horowitz and Manski [1997].

In an effort to demonstrate the robustness of our results, we have reported the samples medians (labelled "Med.") as well as lower and upper quartiles (labelled "L.Q." and "U.Q.," respectively) in table 2.[8] For example, the estimated sample median of Regime 2 is greater than that of Regime 1 at size 0.01. But this is not an implication of the linkage principle, simply corroborating evidence supporting the notion that the English auction

[8]To calculate standard errors of the sample percentiles, we used the following first-order approximation for the q^{th} population per-

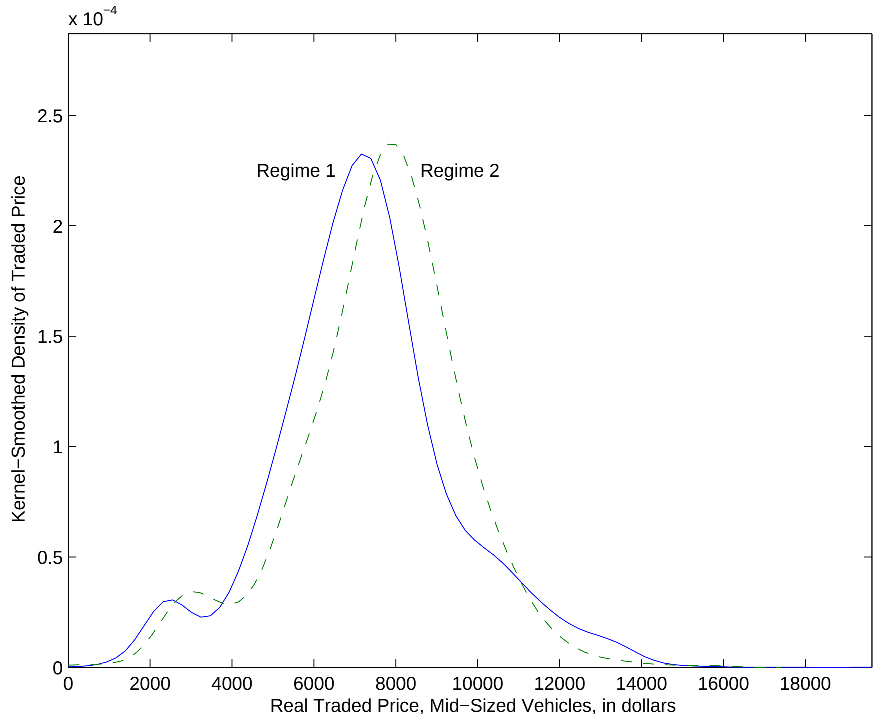

Figure 3: Estimated Kernel-Smoothed Densities of Traded Prices—Mid-Sized Vehicles

generates more information than the sealed-bid auction, and this release of information increases the average revenues garnered under the English auction.

In figure 3, we plot the estimated kernel-smoothed densities of traded prices using a Gaussian kernel with the bandwith parameter recommended by Silverman [1986]—namely., $4T_i^{-1/5}\hat{\sigma}_i/3$. Here, $\hat{\sigma}_i$ denotes the estimated standard deviation of trade prices, while T_i denotes the sample size of Regime $i = 1, 2$. Nothing new concerning the traded-price processes under Regimes 1 and 2 is really learned from this exercise, but we include this graph for completeness.

One obvious limitation of this analysis derives from the aggregation of all mid-sized vehicles into one sample. Within the mid-sized category, however, the top three models account for around two-thirds (7,101 of 10,422) of mid-sized sales, almost one quarter of total sales. Thus, we next disaggregated and focused on the top three models, individually. In the bottom three quarters of table 2, we present descriptive statistics for the top three models of mid-sized vehicles for which we have complete purchase price as well as age and odometer data.

In general, the descriptive statistics for Models A and C, under Regimes 1 and 2, are similar to those for the data concerning all mid-sized vehicles; that is, for these models, the ranking of Regime 2 over Regime 1 remains. The results for Model B are different: for this model, the average revenues are higher under Regime 1

centile $\xi^0(q)$ estimated by the sample percentile $\hat{\xi}(q)$:

$$\sqrt{T}[\hat{\xi}(q) - \xi^0(q)] \xrightarrow{\mathrm{d}} \mathcal{N}\left(0, \frac{q(1-q)}{f^0[\xi^0(q)]^2}\right)$$

where we used $\hat{f}(w)$, the kernel-smoothed estimate of the population probability density function of traded prices $f^0(w)$, evaluated at the sample percentile $\hat{\xi}(q)$, to approximate $f^0[\xi^0(q)]$. We should note, however, that under contamination this standard error statistic is not robust, even though the sample percentiles are, because (like the sample mean) the kernel-smoothed density estimator has a very low breakdown point as well.

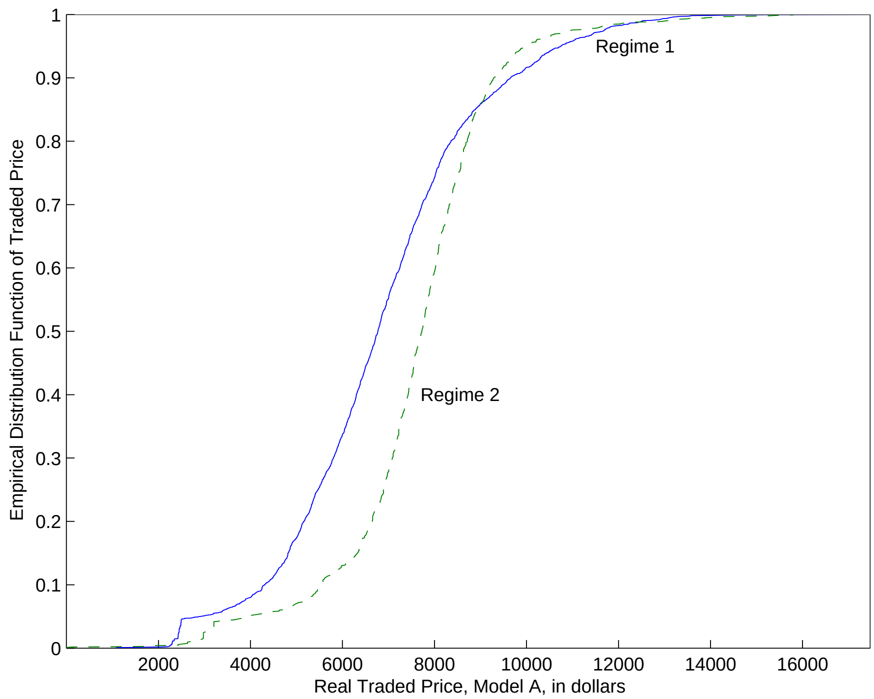

Figure 4: Empirical Distribution Functions of Traded Prices—Mid-Sized, Model A

than under Regime 2.

One obvious, but compelling point emerges from the previous analysis: the vehicles sold could be different in ways that the potential buyers can observe, but which we (as econometricians) cannot. We sought to use observed covariates to control for such factors. One important source of heterogeneity is in the new vehicle itself. While new model vehicles are remarkably homogeneous by some standards, considerable variation can exist in the features those vehicles possess. For example, we may not know whether a vehicle has the optional *Powder White Pearl Paint* or a sunroof or the *Bluetooth Hands-Free Phone System*, but the purchase price will probably reflect a good portion of this heterogeneity. Thus, in order to deal with this heterogeneity, we used p_t, the real purchase price of the t^{th} vehicle, as a control for unobserved features of the vehicle. While we believe that the real new-car price is a reasonable sufficient statistic for all of the unknown features of a vehicle, we should note that this data series is all we have to control for this type of heterogeneity. Also, we know that a vehicle's age is important. For all observations in our dataset, we know when the vehicle was bought and when it was sold—vehicle age, in days, which we then converted to years; we denote this variable by Age. We know, too, that past usage is important. For around eighty-five percent of the vehicles sold, we know the odometer reading when the vehicle left the fleet, which we converted to tens of thousands of miles; we denote this variable by Mileage.

When introducing observed covariate heterogeneity (denoted by the vector $\boldsymbol{x}$, below) into econometric models of auctions, only certain functional forms will lead to tractable empirical specifications. In particular, two different structures have typically been used to introduce observed covariates into the valuations (denoted by Vs, below) of potential buyers. The first is an additive form,

$$V_{nt} = g(\boldsymbol{x}_t) + \varepsilon_{nt}$$

for the n^{th} potential buyer at the t^{th} sale where $g(\cdot)$ is some (typically unknown) function, while the second is a

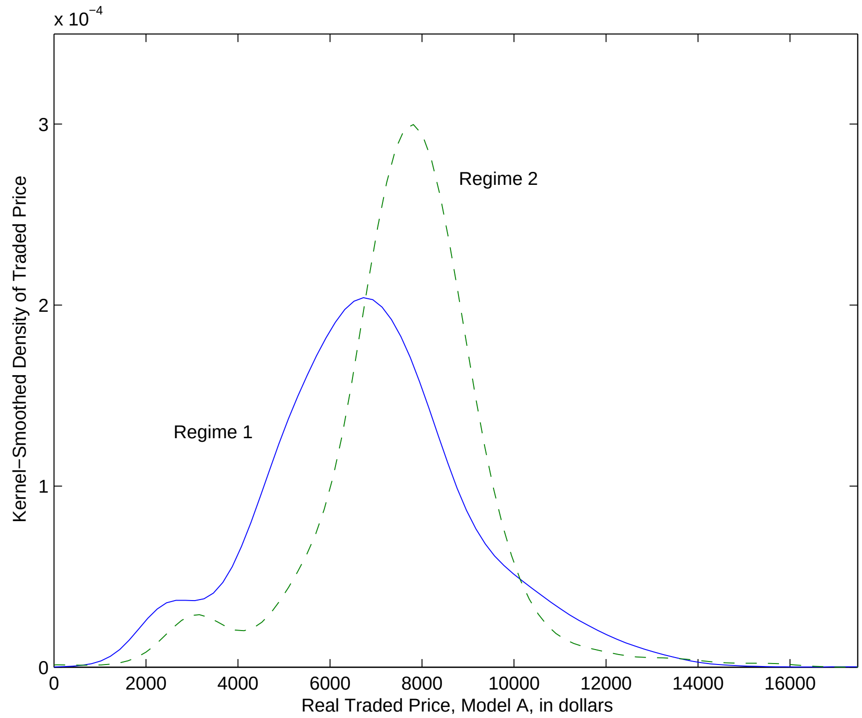

Figure 5: Estimated Kernel-Smoothed Densities of Traded Prices—Mid-Sized, Model A

multiplicative form,

$$V_{nt} = h(\boldsymbol{x}_t)\varepsilon_{nt}$$

where $h(\cdot)$ is some (typically unknown) function. Here, ε_{nt} denotes the unobserved bidder-specific heterogeneity in valuations.

Under these functional-form assumptions, the Bayes–Nash equilibrium bid function is of the form

$$\beta(V_{nt}) = g(\boldsymbol{x}_t) + \beta(\varepsilon_{nt})$$

in the first case, and

$$\beta(V_{nt}) = h(\boldsymbol{x}_t)\beta(\varepsilon_{nt})$$

in the second. When it comes to implementing these specifications, researchers often assume a single-index structure, like

$$g(\boldsymbol{x}) = \boldsymbol{x\theta},$$

$$\log[h(\boldsymbol{x})] = \boldsymbol{x\eta}$$

where $\boldsymbol{\theta}$ and $\boldsymbol{\eta}$ are vectors of unknown parameters conformable to $\boldsymbol{x}$.

We imagine the following separable empirical specification relating $(p_t, \mathsf{Age}_t, \mathsf{Mileage}_t)$ to W_{it}, the traded price of the t^{th} vehicle under Regime i:

$$W_{it} = \rho_1(p_t)\rho_2(\mathsf{Age}_t)\rho_3(\mathsf{Mileage}_t)\lambda_i(S_t) \tag{1}$$

Here, $\rho_1(p_t)$ represents an unknown transformation of the real purchase price, $\rho_2(\mathsf{Age}_t)$ an unknown transformation of Age_t, $\rho_3(\mathsf{Mileage}_t)$ an unknown transformation of $\mathsf{Mileage}_t$, and $\lambda_i(S_t)$ an unknown transformation

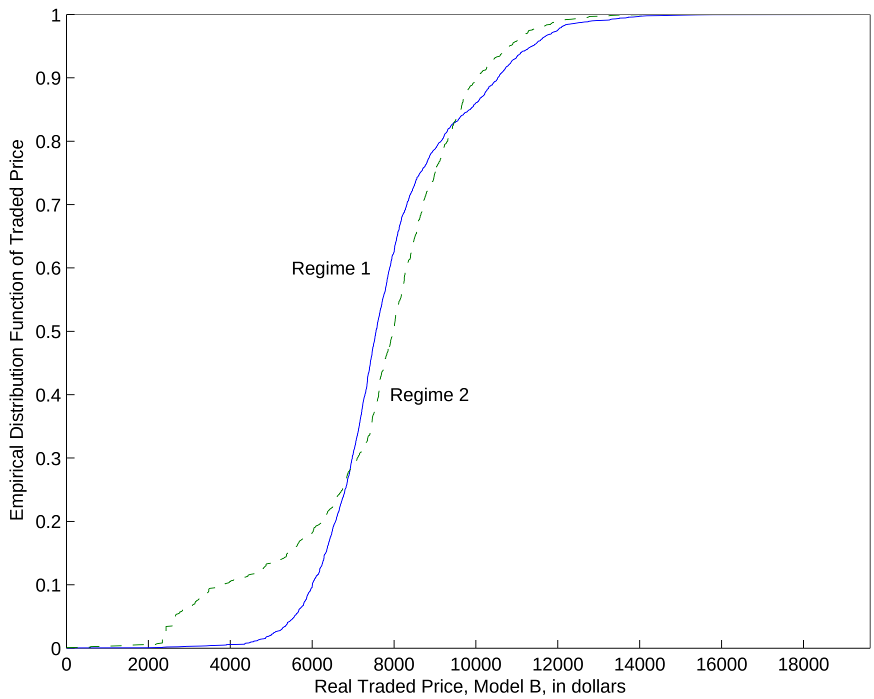

Figure 6: Empirical Distribution Functions of Traded Prices—Mid-Sized, Model B

of the sale-t specific unobserved error term S_t. This latter transformation can vary across the selling regimes $i = 1, 2$. Taking logarithms of both sides of equation (1) yields

$$\log W_{it} = \mu_1(p_t) + \mu_2(\mathsf{Age}_t) + \mu_3(\mathsf{Mileage}_t) + \lambda_0 + (\log[\lambda_i(S_t)] - \lambda_0) \tag{2}$$

where $\mu_j(\cdot)$ denotes $\log[\rho_j(\cdot)]$, $j = 1, 2, 3$. Here, the unknown parameter λ_0 is introduced as a centering parameter: under the null hypothesis that the selling regime does not matter, the random variable $(\log[\lambda_i(S_t)] - \lambda_0)$, which we shall denote below as U_{it}, has mean zero and is uncorrelated with p as well as Age and Mileage.

Suppose

$$\rho_1(p) = A_1 p^{\gamma_1},$$

then

$$\mu_1(p) = \alpha_1 + \gamma_1 \log p.$$

Also, when

$$\rho_2(\mathsf{Age}) = A_2 \delta_2^{\mathsf{Age}}$$

and

$$\rho_3(\mathsf{Mileage}) = A_3 \delta_3^{\mathsf{Mileage}},$$

so constant but different "depreciation" rates with age and mileage, then

$$\mu_2(\mathsf{Age}) = \alpha_2 + \gamma_2 \mathsf{Age}$$

and

$$\mu_3(\mathsf{Mileage}) = \alpha_3 + \gamma_3 \mathsf{Mileage}.$$

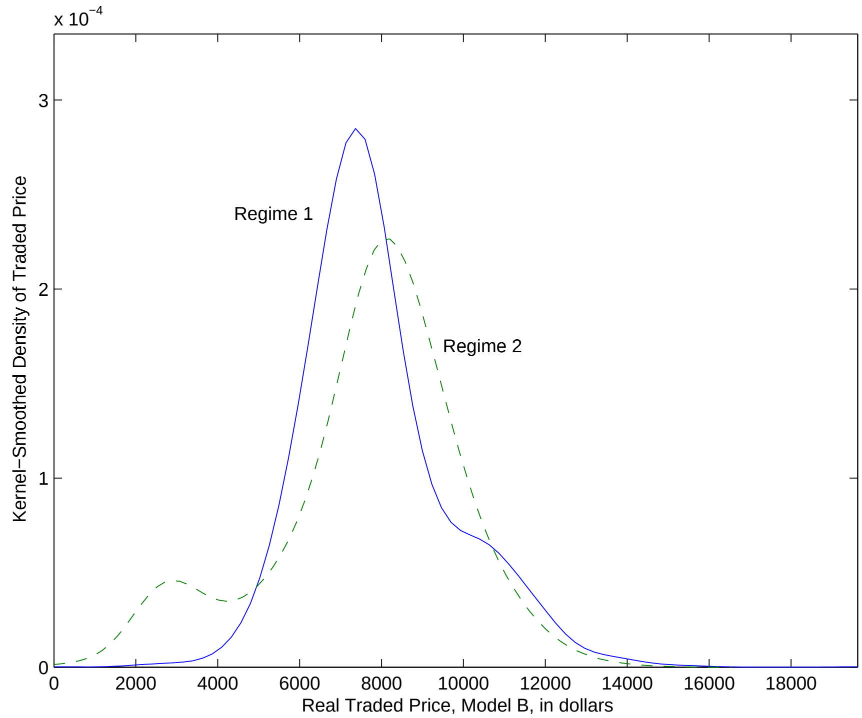

Figure 7: Estimated Kernel-Smoothed Densities of Traded Prices—Mid-Sized, Model B

We estimated the following empirical specification:

$$\log W_{it} = \gamma_0 + \gamma_1 \log p_t + \gamma_2 \mathsf{Age}_t + \gamma_3 \mathsf{Mileage}_t + U_{it} \tag{3}$$

by the method of least squares using all data concerning mid-sized vehicles for which complete observations concerning age and mileage as well as the purchase price were available. These data are listed in table 2 for the rows labelled "Mid-Size, Sub.": the subsample for which complete data concerning covariates exist. We report our parameter estimates as well as robust standard errors in table 3. The estimated "depreciation" parameters for Age (measured in years) and Mileage (measured in tens of thousands of miles) make sense: in the first year, a vehicle is predicted to lose 22.34 percent of its value; controlling for vehicle age, an extra ten thousand miles is predicted to reduce the vehicle's value by around 1.72 percent. In figure 10, we present the EDFs of the fitted residuals (by Regime), while in figure 11, we present the estimated kernel-smoothed densities of the fitted residuals (by Regime). In table 4, we present the descriptive statistics. The most important statistic to notice in this table is the mean: under Regime 2, the average residual is positive, while under Regime 1 it is negative. The average difference in the reported prices is 0.0033 (so 0.33 percent or 33 basis points), having a standard error of 0.0062, which implies a p-value of 0.6 for the hypothesis that the mean under Regime 2 is greater than that for Regime 1. Of course, the Regime 2 traded prices are *net* of commissions, which average around ten percent, while the Regime 1 prices are *gross* prices. When we adjusted the Regime 2 prices by the average commission, the difference rose to 0.0986 (or 9.86 percent), again having a standard error of 0.0062, so a p-value of less than 0.001. These results suggests that the additional information released under the English auction relative to the Internet auction is important, evidence supporting the linkage principle. Another striking feature of figures 10 and 11 is that the estimated residuals from the Engish auctions have much less variation than those from the Internet auctions, a result consistent with the effect of affiliation on English versus other auction formats that release less information.

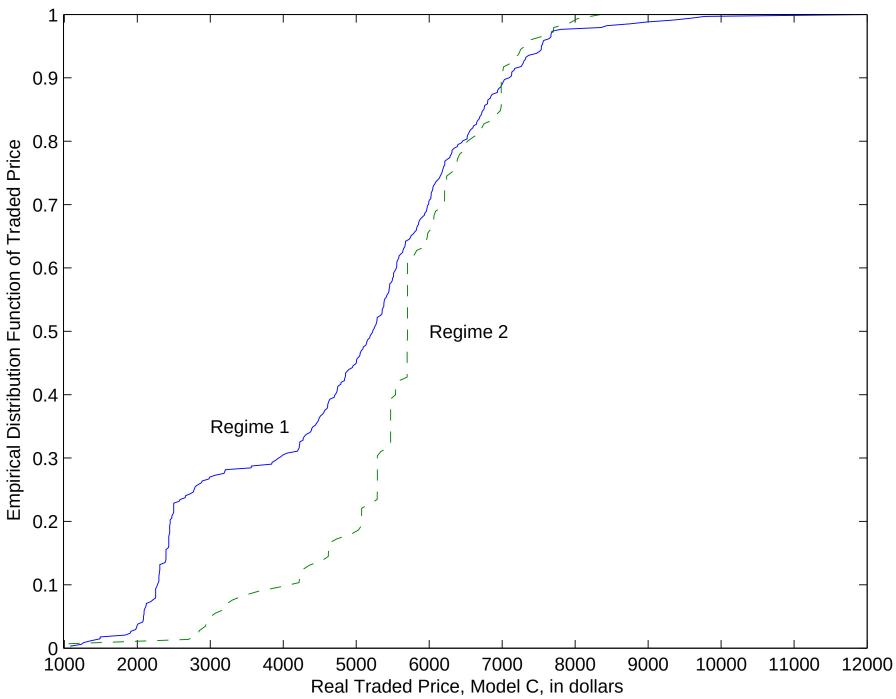

Figure 8: Empirical Distribution Functions of Traded Prices—Mid-Sized, Model C

Table 3: Least-Squares Estimates—Mid-Sized Vehicles

Parameter	Estimate	Std.Error
γ_0	2.0884	0.2557
γ_1	0.7886	0.0255
γ_2	−0.2528	0.0103
γ_3	−0.0174	0.0024
$T = 8,316$	$R^2 = 0.43$	$\hat{\sigma} = 0.2429$

As an epilogue, we note that some odometer readings were quite low, given the vehicle's age—for example, less than ten thousand miles. We do not know whether an odometer reading of $\mathtt{20}$ is really $\mathtt{1,000,020}$ miles, or a mis-reported observations; for example, someone recorded $\mathtt{20}$ instead of $\mathtt{20,000}$. Thus, we constrained ourselves to vehicles having mileages of greater than $\mathtt{10,000}$. At the suggestion of a referee, we also experimented with lower thresholds; our regression results were not robust to including observations whose odometer readings were less than $\mathtt{6,000}$ miles. Basically, observations with unusually low odometer readings became leverage points in the data, so those observations were given usually high weights in the regression and were excessively influential; for more on this, see Belsley et al. [1980].[9] While the estimated regression coefficients did, however, change, the main result that the average price was higher under Regime 2 than under Regime 1 remained unchanged.

[9]Specifically, look in the index for the so-called *hat matrix*.

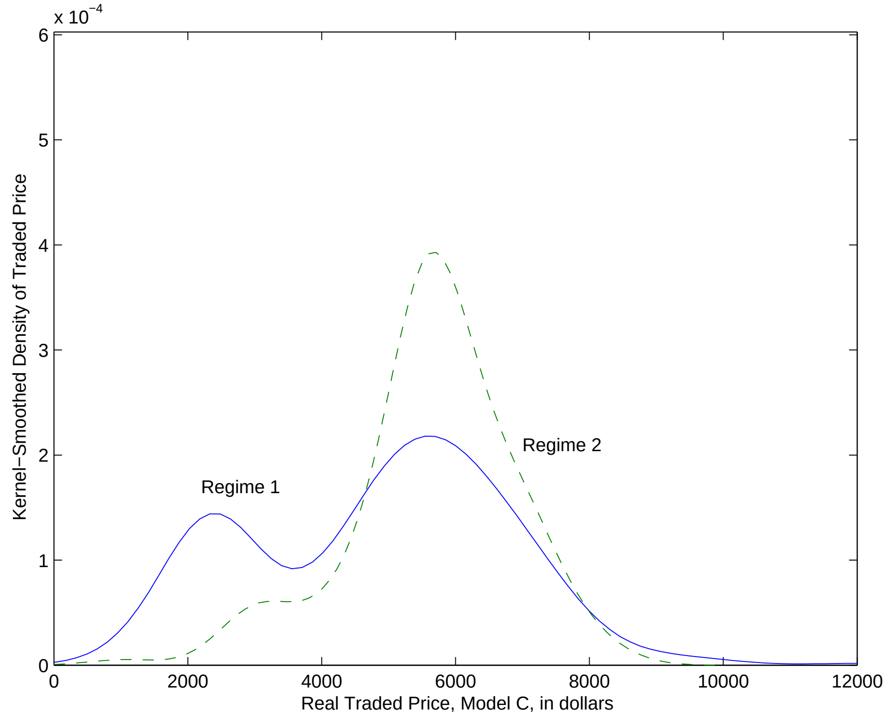

Figure 9: Estimated Kernel-Smoothed Densities of Traded Prices—Mid-Sized, Model C

Table 4: Sample Descriptive Statistics—Least-Squares Fitted Residuals, By Regime

Sales Regime	Mean	St.Dev.	L.Q.	Med.	U.Q.	No.Obs.
Regime 1	−0.001347	0.287709	−0.130008	0.034165	0.190377	4,557
Regime 2	0.001633	0.173741	−0.041629	0.036839	0.099573	3,759

4 Conclusions

In this paper, we have presented empirical results derived from a unique new dataset concerning the revenues earned by a large rental car company that sold used cars at two different kinds of auctions. This company experimented with several different selling mechanisms to dispose of unwanted, used vehicles, including designing a unique new Internet auction. To our knowledge, this Internet auction has never been analyzed theoretically, or empirically. Using simple empirical methods, we have analyzed these data to examine the effect that different auctions rules had on the average revenues earned by the rental car company, especially the role of information release.

Our empirical results are potentially subject to alternative interpretations. On the one hand, in general, we found that average traded prices were highest when vehicles were sold under the oral, ascending-price format (the standard English auction) as opposed to the Internet auction under which very little information was released. A potential, alternative explanation exists for the increase in average traded prices: the increase in average traded prices occurred because the number of potential bidders increased. This increased number of potential bidders alone is sufficient to explain why the average traded prices increased. Although we know the exact number of participants at each Internet auction as well as the set of potential bidders, unfortunately, the data do not allow us to determine the number of potential bidders at the English auctions, let alone which bidders participated at the English auction. From other sources (specifically, Kim and Lee [2008]) we have learned that

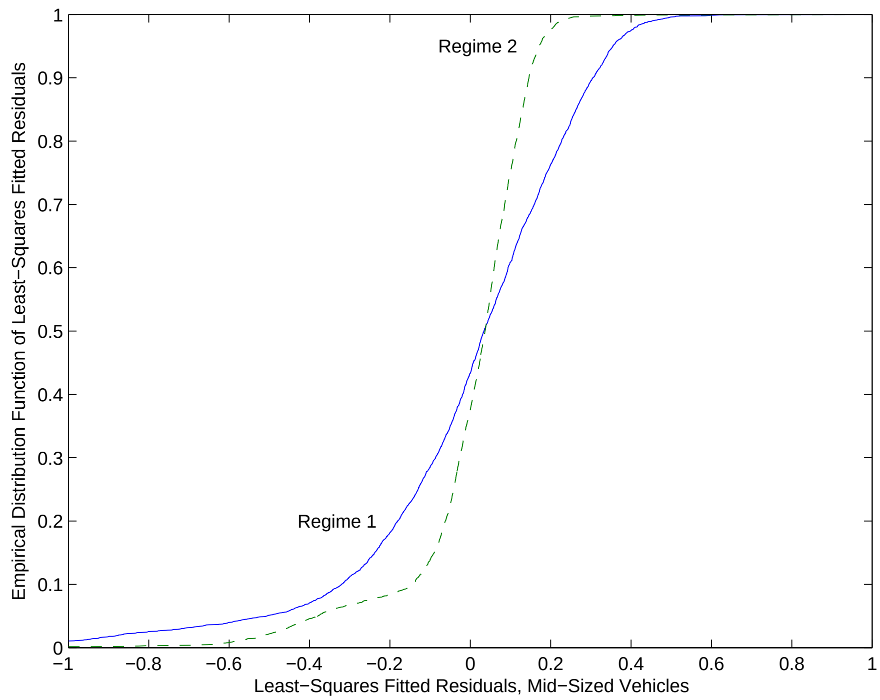

Figure 10: Empirical Distribution Functions of Least-Squares Fitted Residuals—Mid-Sized Vehicles

134 distinct bidder identification numbers were recorded by the auction house, which is very close to our 124.

Thus, we believe the most likely explanation for what we have found is that, when the company switched from the Internet auctions to the English auctions conducted by the auction house, the increase in information provided to participants at these auctions (that is, the linkage principle) is a key reason why the average traded prices increased at these auctions. While we do not know the number of participants at any given auction, we do know that the pool of potential bidders was larger, around 134, than the pool of potential bidders at the Internet auctions conducted by the rental company where there were 124 potential bidders. When the number of potential bidders is large, however, the relative effects of competition on traded prices is much smaller than when the number of potential bidders is small. Unfortunately, we have no way of knowing how many bidders actually participated (that is, called-out bids) at the English auctions conducted by the auction house. It is possible that the English auction induced increased participation by members of the set of potential bidders, but it is that increased participation which is predicted by the linkage principle under affiliation: emboldened by the actions of their rivals, additional potential bidders tender bids (participate) and this increased participation increases the average traded price.

Thus, while the results of our empirical analysis are relatively unambiguous—the average traded prices earned by the rental car company at the English auctions conducted by the auction house were the highest, especially when we note consider the *gross* traded prices rather than the *net* traded prices received by the rental car company—we cannot be absolutely certain whether the increase in average traded prices reflects primarily the linkage principle or a demand-aggregation effect (that is, an increased number of potential bidders at the auctions conducted by the auction house).

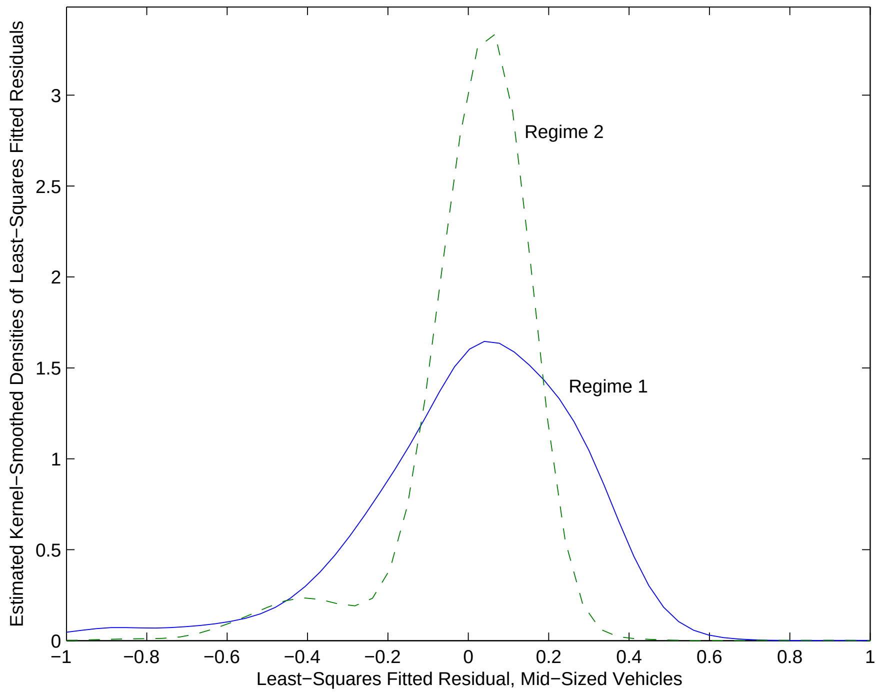

Figure 11: Estimated Kernel-Smoothed Densities of Least-Squares Fitted Residuals—Mid-Sized Vehicles

References

Lawrence M. Ausubel. An efficient ascending-bid auction for multiple objects. *American Economic Review*, 94:1452–1475, 2004.

David A. Belsley, Edwin Kuh, and Roy E. Welsch. *Regression Diagnostics: Identifying Influential Data and Source of Collinearity*. New York: John Wiley & Sons, 1980.

Stephen G. Donald, Harry J. Paarsch, and Jacques Robert. An empirical model of the multi-unit, sequential, clock auction. *Journal of Applied Econometrics*, 21:1221–1247, 2006.

Han Hong and Matthew Shum. Rates of information aggregation in common value auctions. *Journal of Economic Theory*, 116:1–40, 2004.

Han Hong, Harry J. Paarsch, and Pai Xu. On the asymptotic distribution of the transaction price in a clock model of a multi-unit, oral, ascending-price auction within the common-value paradigm, typescript, Department of Economics, University of Melbourne, 2009.

Joel L. Horowitz and Charles F. Manski. What can be learned about population parameters when the data are contaminated? In G.S. Maddala and C.R. Rao, editors, *Handbook of Statistics, Volume 15*. Amsterdam: Elsevier, 1997.

Peter Huber. *Robust Statistics*. New York: John Wiley & Sons, 1981.

John H. Kagel and Dan Levin. The winner's curse and public information in common value auctions. *American Economic Review*, 76:894–920, 1986.

John H. Kagel and Dan Levin. *Common Value Auctions and the Winner's Curse*. Princeton: Princeton University Press, 2002.

John H. Kagel, Ronald M. Harstad, and Dan Levin. Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica*, 55:1275–1304, 1987.

Samuel Karlin. *Total Positivity*, Volume 1. Stanford, California: Stanford University Press, 1968.

Kyoo il Kim and Joonsuk Lee. A structural analysis of wholesale used-car auctions: Nonparametric estimation and testing of dealers' valuations with unknown number of bidders, typescript, Department of Economics, University of Minnesota, 2008.

Dan Levin, John H. Kagel, and Jean-François Richard. Revenue effects and information processing in english common value auctions. *American Economic Review*, 86:442–460, 1996.

Paul R. Milgrom and Robert J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50:1089–1122, 1982.

Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:55–73, 1981.

Motty Perry and Philip J. Reny. Efficiency and information aggregation in auctions. *Econometrica*, 67:895–900, 1999.

Wolfgang Pesendorfer and Jeroen Swinkels. Efficiency and information aggregation in auctions. *American Economic Review*, 90:499–525, 2000.

John G. Riley and William F. Samuelson. Optimal auctions. *American Economic Review*, 71:381–392, 1981.

Alvin E. Roth and Axel Ockenfels. Last-minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the Internet. *American Economic Review*, 92:1093–1103, 2002.

Bernard W. Silverman. *Density Estimation for Statistics and Data Analysis*. London: Chapman and Hall, 1986.

Steven Tadelis and Florian Zettelmeyer. The effect of information on auction outcomes: A large scale field experiment, typescript, University of California, Haas School of Business, 2010.

William S. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

William S. Vickrey. Auctions and bidding games. In *Recent Advances in Game Theory,* Princeton Conference Series, 29. Princeton, New Jersey: Princeton University Press, 1962.

Robert J. Weber. Multiple-object auctions. In Richard Engelbrecht-Wiggans, Martin Shubik, and Richard M. Stark, editors, *Auctions, Bidding, Contracting: Uses and Theory*. New York: New York University Press, 1983.

Robert B. Wilson. A bidding model of perfect competition. *Review of Economic Studies*, 44:511–518, 1977.