The Pursuit of Revenue Reduction: An Experimental Analysis of the Shanghai License Plate Auction

Evan Zuofu Liao and Charles A. Holt*

May 2013 (revision in progress, comments welcome)

Abstract

Monthly auctions of vehicle licenses in Shanghai were reconfigured in 2008 in an attempt to reduce prices by giving bidders the option of modifying initial bids. The modified bids are required to be within a narrow band around the lowest accepted bid at that point. The bidding constraints limit price competition, and therefore, allocation efficiency may be degraded if intended revenue reductions are achieved. The effectiveness of the new procedures, however, is unclear in light of sharp subsequent increases in license prices, even after the license quota was doubled. This paper reports a laboratory experiment designed to evaluate the revenue and efficiency consequences of the new Shanghai auction format.

I. Introduction

Auctions are commonly used in the private sector to create competition by grouping buyers together in a thick market with a clear time focus. In the public sector, auctions are attractive because they combine non-economic benefits (fairness, speed, and transparency) with economic benefits (revenue generation and allocative efficiency), bypassing the transactions costs of non-market procedures, e.g. waiting in line (Holt and Sherman, 1982) or "rent-seeking" in all-pay lobbying contests.[1] Even when other

---

* University of Chicago (Liao) and University of Virginia (Holt) Corresponding author: Charles Holt, Department of Economics, University of Virginia, Charlottesville, VA 22904-4182, cah2k@virginia.edu
This paper is an extension of a thesis that Liao wrote for the Distinguished Majors Program at the University of Virginia. We wish to thank Nancy Cheng, Virginia Hawkins, Nathan Larson, John Pepper, Mike Schreck and Sean Sullivan for helpful suggestions. This research was funded in part by the University of Virginia Bankard Fund.

[1] See Coase (1959) for a seminal discussion of the advantages of using auctions, with appropriately defined property rights, to allocate segments of the electromagnetic spectrum. Coase's auction proposal,

objectives are paramount, e.g. limiting greenhouse gas emissions, the hope is that auctions will provide price signals that minimize the efficiency consequences of the constraints being imposed. In public settings, however, revenue generation may take a back seat to political or other long-range objectives.[2] This paper is motivated by an auction that was designed to reduce revenue, while maintaining some degree of price competition.

About 20 years ago, the city of Shanghai began auctioning a limited number of automobile license plates in an effort to curb traffic congestion. These multi-unit, sealed-bid auctions were "discriminatory" in the sense that winning bidders paid their own bids. The cutoff (lowest accepted bid) was announced after each auction. Since cars were priced as luxury items and most consumers could not afford them, supply and demand conditions were relatively stable at first. Prices began rising steadily about a decade ago, which led to an abandonment of reserve prices in 2003. Prices rose from about 15,000 Renminbi (RMB) in 2002 to about 56,000 RMB in 2007, which corresponds to an increase from $1,800 to $7,500 in US dollars, as shown by the price spike in the center of Figure 1.

In the following year, the city adopted a new "Shanghai Auction" format that permits limited revisions of initial bids. This is a continuous-time auction in which prospective bidders can observe the current lowest accepted bid prior to submitting a bid. The auction begins with a period in which bidders may submit a single initial bid while observing the lowest accepted bid at that point. When the initial bid phase ends at a pre-

---

met with derision at time, was finally implemented more than three decades later. McMillan (1994) cites an estimate from the U.S. Federal Communications Commission (FCC) showing that the cost of assigning cellular licenses was six times higher with administrative proceedings than with auctions. Hazelett and Michaels (1993) used resale prices for cell phone licenses following an early FCC distribution by lottery to estimate that 35-40% of the rents for these licenses were dissipated by application costs. The paperwork produced by this lottery resulted in the collapse of the floor of an FCC warehouse. Anderson, Goeree, and Holt (1998) show that bounded rationality can cause rent dissipation in contests to be even greater than Nash equilibrium predictions.

[2] For example, the planned US Department of Interior auctions of offshore alternative energy generation tracts are not intended to generate unreasonably high prices, perhaps driven by winner's curse effects that could later impede the "build out" of wind towers on the outer continental shelf. Similarly, in early briefings of officials from the electric power industry, officials of the Regional Greenhouse Gas Initiative (RGGI) sometimes expressed a desire to sell emissions allowances at a "fair price" (Holt, et al., 2007). More recently, the US Treasury's objective in purchasing toxic mortgage backed securities from banks in a planned TARP auction was not to acquire these securities at "fire-house prices." Here the goal was to inject funds into the fragile banking system, and officials were looking for an auction that would provide a fair price for those assets (Armantier, Holt, and Plott, 2012).

announced time, there is a bid revision phase in which bidders are only permitted to make up to two bid changes within a 30 minute bid-revision period. The most severe constraint, however, is the requirement that any changed bid must fall within a very narrow price window (equivalent to about $50) around the lowest accepted bid in effect at that point in time. It is, of course, possible to submit a bid and not revise it, even though it falls above or below the window of competition. The lowest accepted bid may change over time, and hence, the window becomes a moving target. In the event of a tie, preference is given to the earliest submitted bid(s), which might tend to discourage excessive sniping at the lowest accepted bid in the final seconds. Even after the change was implemented, it was difficult to assess its effects, since the city doubled the license quota in the same month that they introduced the new auction format (see the spike in the middle part of Figure 1). Moreover, auction prices have risen steadily despite an approximate doubling of license quotas from pre-2008 levels, as can be seen from the trend in the dashed quota line in Figure 1. The auction raised about a billion dollars in 2012, and the price of a license plate exceeded $14,000 in March 2013, surpassing the price of many cars. In view of these price increases, the revenue consequences of the new auction format are unclear.

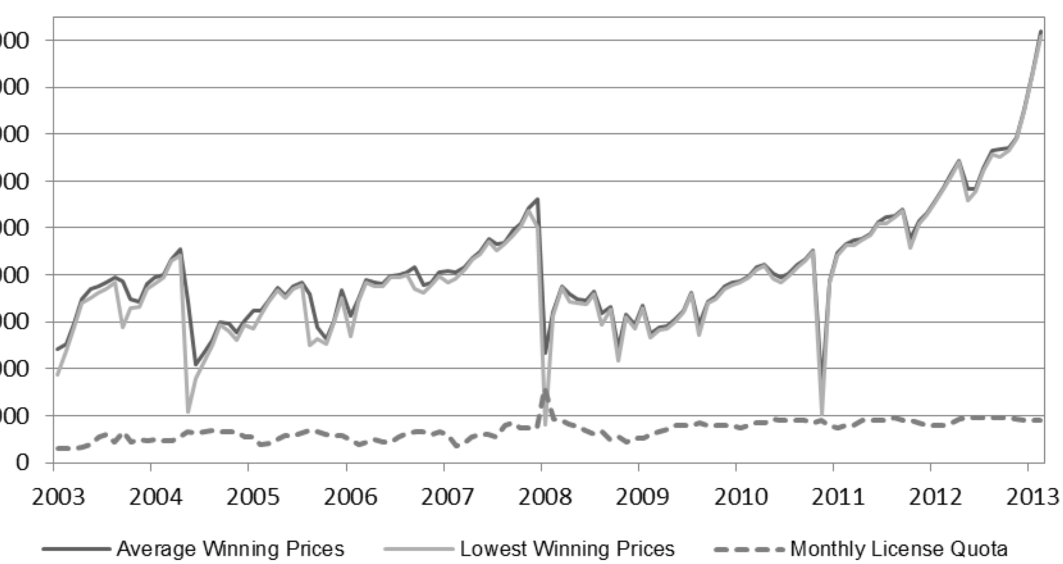

Figure 1. Shanghai Auction Prices (solid lines) and License Quotas (dashed line).

The Shanghai government did not explicitly tout the new auction as a way of reducing license prices, but the price reduction motive is widely cited in the Chinese media and in academic circles (Luo, 2008 and Chen, 2008). It was difficult to assess the effects of the new auction procedures on revenue and price levels in advance, given the complexity of the strategic setting. One naïve approach would be to reason that the bid revision period would create competition between provisionally excluded bidders, which may tend to raise revenues and efficiency measures. An alternative conjecture would be that bidders tend to bid low initially with less fear of being excluded given that bids can be revised subsequently. Then the narrowness of the range of permissible bid modifications might dampen subsequent price competition, resulting in lower final sale prices and reduce the efficiency of final allocations.

Any attempt to reduce revenue by curtailing bid competition, of course, could backfire by generating inefficiencies due to sniping or other factors that blur correlations between bids and values.[3] The most recent auction in April 2013 also implemented a biding price ceiling for the initial bids, set below the cutoff price in the March auction. This ceiling would tend to further reduce efficiencies. It is difficult to assess the efficiency consequences of the Shanghai auction, however, since private values cannot be observed, and secondary markets for licenses are subject to some constraints. Someone who wishes to sell a license legally must also sell the vehicle to which it is registered, and even this cannot be done for a specified time after the auction (the current wait period is 3 years). Moreover, the purchaser of a license/car must wait an additional year before transferring the license to another vehicle. The purpose of these restrictions was presumably to limit speculation. Some indirect evidence of inefficiencies is indicated by accounts of citizens who purchase license plates from neighboring provinces and use them on Shanghai streets, despite the need to avoid traffic checkpoints (Inch, 2012). On the other hand, even an auction with some price inertia would probably yield higher efficiencies than would result from a random lottery allocation, such as the one that was recently initiated in Beijing.

[3] See Holt et al. (2007) for an analysis of the detrimental effects of sniping in a continuous time laboratory auction patterned after auctions for emissions permits.

This paper reports a laboratory experiment designed to evaluate the revenue and efficiency effects of the new Shanghai auction format. Although a considerable amount of context is sacrificed in a small-scale laboratory setting, the advantages of replication and control over structural changes and individual private value differences can provide a deeper perspective on the workings of this new type of auction. In particular, efficiency can be directly observed in the lab, which is not possible in the field, and key features of the auction design can be changed independently to help deconstruct the cause of revenue changes.[4] The experiment design and web-based bidding procedures are described in the next section. Sections III and IV present aggregate revenue and efficiency measures and an evaluation of bidding patterns that provide an understanding of the differences that are reported. The concluding section contains a brief summary and comparison with the lottery system recently adopted in Beijing.

II. Experiment Design and Procedures

A useful classification of multi-unit auctions makes distinctions based on whether the auction is static ("sealed bid") or dynamic (with bid revisions), and whether the winning high bidders pay their own bids or a uniform market-clearing price. A static pay-as-bid auction is called "discriminatory" since the winning bidders pay different amounts, as compared with a "uniform price" auction in which all bidders pay a common, clearing price, e.g. the highest rejected bid. Dynamic auctions may permit continuous bid revisions, as with many on-line auctions, or may offer discrete opportunities to reduce desired quantities as a uniform price is raised by increments, as if by the ticks of a clock. The experiment that we report involves a comparison between a sealed-bid discriminatory auction (used in Shanghai prior to 2008) and a continuous discriminatory

[4] Auction experiments are surveyed in Kagel and Levine (2012). In particular, laboratory experiments have been used to assess and compare different spectrum auction designs in the US and the UK, e.g.: Plott (1997), Banks, Olsen et al. (2001), Goeree and Offerman (2002), and Goeree and Holt (2008). Abbink, Brandts, and McDaniel (2003), Rassenti, Smith, and Wilson (2003), and Fabra et al. (2003) present experimental comparisons of different auction mechanisms for the electricity market. Also, see Abbink, Brandts, and Pezanis-Christou (2006) for a laboratory study of three alternative methods of auctioning off government securities. Holt et al. (2007) and Burtraw et al. (2010) evaluate different auction formats for the Regional Greenhouse Gas Initiative (RGGI) auctions for selling CO_2 emission allowances. Many of these studies involve comparisons of discriminatory and uniform price auction formats, which has been a major focus of attention since Milton Friedman (1960) argued for the adoption of competitive, uniform-price procedures in Treasury Bill auctions.

auction with limited opportunities for bid revision, as implemented in the current Shanghai auction. In addition, we included a sealed-bid uniform price auction, using the highest rejected bid to provide an incentive for bidders to bid “at value” when they only demand 1 unit, so the predicted clearing price equals the intersection of demand and the vertical supply at the auction quantity. The uniform price auction thus provides a kind of Walrasian benchmark.[5]

We ran three sessions for each of these auction formats, which will be referred to as “Discriminatory”, “Uniform Price” and “Shanghai.” In addition, we ran three sessions with a “Shanghai Sealed Initial Bid” partial implementation of the new auction. In this case, initial sealed bids were submitted without any information provided about the cutoff, as with a discriminatory auction, prior to a constrained bid revision phase that matches the Shanghai auction. The purpose of this sealed initial bid treatment was to assess the relative importance of the two components of the new Shanghai auction, the continuous initial bid phase and the continuous bid revision phase.

Each session consisted of 12 participants, who were bidders in a series of 10 auctions, with 6 licenses up for sale in each.[6] Participants could bid for a single license, referred to as a “unit” in the instructions, with a value that was randomly determined from a discrete uniform distribution on the interval from $50 to $150, with each integer in this range being equally likely. Although bidders’ valuations in most auctions have some elements of both private and common values, the sharp restrictions on resale of the Shanghai plates suggest that a private value formulation is most appropriate.[7] We used

[5] Milton Friedman (1960) advocated the use of uniform price auctions for U.S. Treasury Bills, on the basis that such procedures may be more attractive to non-specialists who prefer to rely on the auction for price discovery. In addition, Friedman argued that the discriminatory auction was more prone to collusive manipulations of dealers. These factors are not likely to be important in the Shanghai auction, which attracts many thousands of bidders each month. In laboratory experiments with only 6 bidders and rich opportunities for collusive “chat,” both uniform price and discriminatory auctions performed comparably in terms of revenue and efficiency (Burtraw et al., 2009). In contrast, the dynamic clock auctions exhibited significant revenue reductions in these collusion-friendly environments, as compared with sealed-bid uniform price and discriminatory auctions. The collusive discussions in clock auctions were often focused on mutual quantity reductions. In the Shanghai setting, each bidder only bids for a single license, so collusion driven by demand-reduction would be highly unlikely. However, a clock auction would probably be difficult to manage in the Shanghai setting, with thousands of on-line bidders, each bidding for a single license and updating bids as the auction price is raised in a large number price increments.

[6] The ratio of the license quota to the number of bidders in each auction has been declining over time, but the average of these quota-to-bidder ratios over the period from 2003 to 2013 is about 0.5.

[7] The standard common value model might apply if bidders received independent signals about upcoming changes in the attractiveness of driving or taking the subway that would affect everyone’s values associated

the same random number seed to ensure that participants in each treatment received the same sequences of random value draws (random number sequences differed by ID number but were identical across treatments). These parameters are shown in Table 1.

Table 1. Experiment Design

Uniform Price	Discriminatory	Current Shanghai Auction
3 sessions 10 auctions 12 bidders 6 licenses private values: u[50, 150]	3 sessions 10 auctions 12 bidders 6 licenses private values: u[50, 150]	6 sessions 10 auctions 12 bidders 6 licenses Private values: u[50, 150]
sealed bids price paid = highest rejected bid	sealed bids price paid = own accepted bid	3 sessions with a 3 minute bid interval, and 3 sessions with sealed bids cutoff price = lowest accepted bid 2 minute bid revision period at most 2 bid changes within plus or minus $1 of current lowest accepted bid price paid = final accepted bid

The bottom part of Table 1 provides the auction details. In each auction, the final bids were ranked to determine the 6 winning bidders, with ties decide by random draw, except for the Shanghai auction with an initial bid interval, for which we used the bid submission time to break ties. Earnings for winning bidders were determined by the difference between the person's private value and the price paid. The 6 winning bidders paid their own bid amounts in the discriminatory and Shanghai auctions. In contrast, the winning bidders in the uniform price auctions paid the highest rejected bid, which serves as a Walrasian market clearing price. In all formats, initial bids were required to be in the interval from $10 to $200. The initial bidding phase for the Shanghai auction lasted for 3 minutes, with information provided about the lowest accepted bid and number of initial bids received at that point. Such information was not provided in the sealed initial bid variation of the Shanghai auction treatment. Regardless of whether initial bids were submitted as sealed bids or in a time interval with information feedback, bid revisions in the Shanghai auction were required to be within plus or minus $1 of the cutoff price (in 10 cent increments), with the cutoff determined by the lowest accepted bid at that point in

with commuting by car. It seems reasonable to expect such informational differences to be dwarfed by private value differences due to the distance of one's commute, the convenience of subway stations to one's residence, possibilities for car pooling, parking costs at one's work site, etc.

the auction.[8] After submitting initial bids in the Shanghai auction, subjects could refresh their submit screens to see the status of their bids (provisionally winning or provisionally losing), the remaining time in the two-minute revision period, as well as the cutoff price determined by the lowest accepted bid at that point in time (as is the case for the post-2008 Shanghai auction).

The 144 participants were recruited from a database of students at the University of Virginia. Participants were paid a $6 show-up fee, plus 10% of their earnings at the end (in cash).[9] Total earnings (including the show-up fee) ranged from about $10 to about $40, with average earnings of about $20-$30, depending on the treatment. Earnings variation was higher in the Shanghai auction, where variations in individuals' bidding strategies seemed to be important. Session lasted from 40 to 90 minutes, depending on the treatment.

III. Performance Comparisons

It is useful evaluate auction outcomes by comparing them with Walrasian predictions. In any symmetric equilibrium, the 6 licenses will be allocated to the bidders with the highest values, and the sum of those values constitutes the maximum surplus (area under demand to the left of the vertical supply at the auction quantity). The actual revenue sequences, as a percentage of this maximum surplus, are shown in Figure 2, for the three sessions with Sealed Bid Discriminatory auctions (thin lines) and for the current Shanghai auction format that permits limited bid revisions after the initial bid interval (dashed lines).

[8] The bid revision range of $2 was selected to be a small fraction of the actual winning bids, to replicate the narrow range used in the Shanghai auction, with a bid revision range of 600 RMB as compared with an average winning price of about 50,000 RMB.

[9] The auctions were run at the University of Virginia with the web-based Veconlab software (http://veconlab.econ.virginia.edu/admin.php), using the Multi-unit Auction program listed on the Auctions menu. The three auction formats were implemented by selecting the "Sealed Bid Discriminatory," "Sealed Bid Uniform Price," or "Continuous Discriminatory" options, with other parameter choices as indicated in table 1. A copy of the instructions is available in the Appendix.

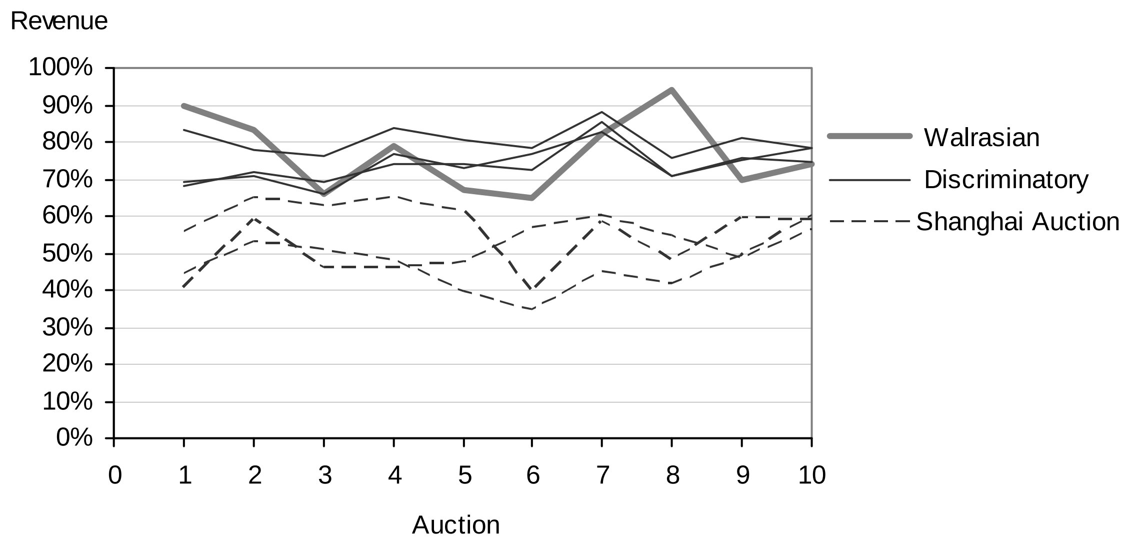

Figure 2. Average Revenues: Walrasian Prediction (thick line), Sealed Bid Discriminatory Auctions (thin lines), Shanghai Auctions with Revisions (dashed lines).

As a benchmark, the Walrasian revenue is shown as a thick line in Figure 2. This is the revenue, as a percentage of the maximum surplus, which could result if all six units were sold at a common uniform price that equals the 7^{th} highest private value (in other words, at a price that equates supply and demand for the random value sequences used). It is apparent from Figure 2 that the revenue percentages for the sealed-bid discriminatory auction sessions tend to track Walrasian predictions somewhat loosely, with the notable exception of the eighth auction, where the predicted revenue spike is not observed.[10] In contrast, the new Shanghai format yields lower revenues for all auctions in all sessions. Hence our first conclusion:

Revenues: T*he Shanghai auction with a narrow band of limited bid revisions results in lower sales prices and revenues, in comparison with the sealed-bid discriminatory format used previously.*

[10] The Walrasian revenue prediction would be generated by bidding at value in a uniform price auction. The divergence between Walrasian revenue and observed revenue in auction 8 for the *discriminatory* format can be explained by considering the Nash equilibrium bids for that format. In the next section, we will derive the Nash equilibrium bid function for the multi-unit discriminatory auction, which can be used to construct a series of round-by-round revenue predictions. Taking the highest 6 value draws for auction 8 and inserting them into the Nash equilibrium bid function derived below, one obtains a revenue prediction (77%) that is bracketed by the observed revenues for the discriminatory auction sessions in that round.

The experiment provides evidence that the change in format had its intended effect. This difference is economically and statistically significant, using session level revenue averages from Table 2 for comparisons between the discriminatory auction and either variation of the Shanghai auction.[11] It is also apparent from Table 2 that there are no significant differences between revenues for sealed-bid discriminatory and uniform price auctions, and that the uniform price auctions also yield higher revenues than either the Shanghai Auction with the initial bid interval or the deconstructed version with initial sealed bids. Finally, it is apparent from the bottom two rows of Table 2 that most of the revenue reduction with the new Shanghai auction format is due to the constrained bid revision interval (present in both Shanghai auction treatments) and not to the switch from sealed initial bids (next to last row) to the initial bid interval (last row).

Table 2. Revenues and Efficiencies by Session and by Treatment (in **Bold**)

	Session Revenues: **Treatment Average**	Session Efficiencies: **Treatment Average**
Uniform Price Sealed Bid	74.6, 76.6, 82.2: **78.5**	96.3, 99.2, 96.6: **97.4**
Discriminatory Sealed Bid	74.9, 73.6, 80.5: **76.3**	97.1, 97.6, 98.3: **97.6**
Shanghai Auction (sealed initial bids)	57.1, 54.9, 53.6: **56.2**	87.1, 91.9, 85.7: **86.2**
Shanghai Auction (initial bid interval)	46.6, 51.5, 57.3: **51.8**	85.7, 84.6, 88.1: **86.1**

Auctions are generally fast and transparent, and they may provide significant revenue-generating benefits, as compared with other non-market allocation mechanisms like lotteries. But the primary attraction of auctions is that price-based mechanisms can distribute goods efficiently to those who value them the most. Efficiency is typically measured by taking the ratio of the actual realized surplus (sum of private values of winning bidders) to the maximum surplus. It is well known that such measures can be artificially inflated by adding a constant to all values, and therefore a comparison with the surplus generated by a purely random allocation should also be considered. The thick

[11] A conservative test can be constructed from the revenue averages by session. The three lowest revenue averages were observed in the continuous-revision Shanghai format. There are "6 take 3" = 20 ways that the six session averages for these two treatments could have been allocated across treatments, and of these, the most extreme outcome was observed. Since there is no overlap, a test based on session averages would yield the same result as a test based on ranked averages. Hence, the p value for a Wilcoxon test of the null hypothesis of no effect is 1/20 = 0.05 (one-tailed test) or 0.10 (two-tailed test).

solid line in Figure 3 shows the efficiency generated by a random allocation, which serves as a lower bound on auction performance comparisons.[12]

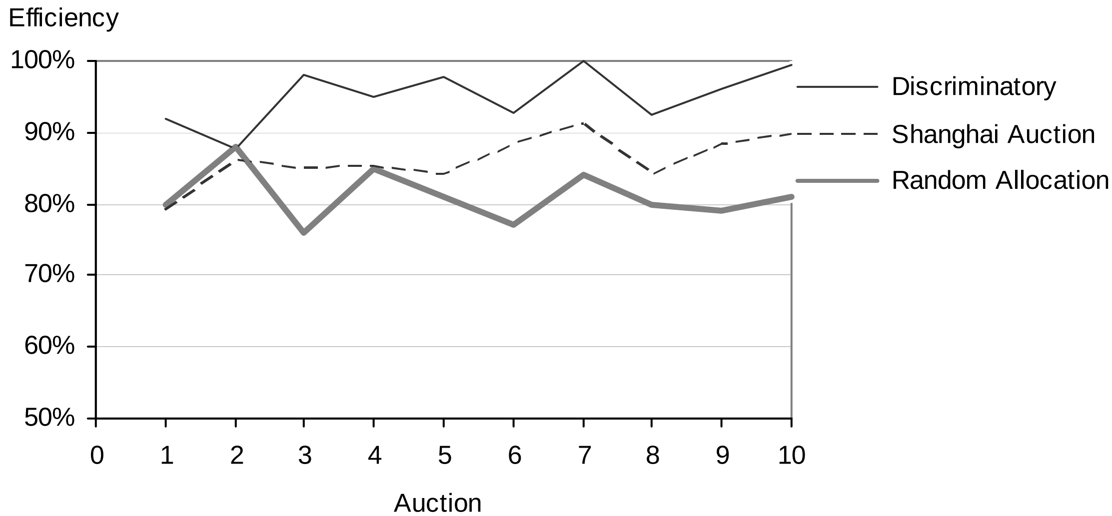

Figure 3. Average Efficiencies for All Sessions in Each Treatment
Random Allocation Benchmark (thick), Sealed Bid Discriminatory Auctions (thin),
Shanghai Auctions Initial Bid Interval and Continuous Revisions (dashed).

Even though the Shanghai auction permits some opportunities for efficiency-enhancing revisions of initial bids, these opportunities are strictly limited, which may reduce efficiency. It is apparent from Figure 3 that the efficiency reduction forces dominate in the sense that average efficiency (across all sessions) is lower for the Shanghai auction (dashed line) than for the sealed bid discriminatory auction used previously (thin line). In fact, the Shanghai auction does not do much better than the random allocation in the first 5 auctions, but there is improvement in the last 5 auctions, where the efficiency is about halfway between the random benchmark and full efficiency. Despite the fact that the sealed-bid discriminatory auction does not permit *ex post* competition, efficiency is higher with this format, near 100% in some auctions. This difference is economically significant, and statistically significant using session-level

[12] The random allocation efficiency of about 80% may seem high, but this is due to the high lower bound on the range of private values (50 to 150). Roughly speaking, a random allocation will generate an average value of about 100, whereas the average of the top half of the draws (the efficient allocation) would be at about 125, so the random allocation generates an efficiency of about 80%.

measures shown on the right side of Table 2.[13] An analogous argument shows that efficiencies for uniform price auctions are also significantly higher than for the Shanghai auction. To summarize:

Efficiencies: *The Shanghai Auction with limited bid revision opportunities yields lower efficiencies than the sealed bid Uniform Price Auction and the Discriminatory Auction, with no bid revision opportunities.*

IV. Individual Bidding Behavior

In this section, we examine individual bidding patterns to explain the differences in efficiency. The simplest case is the uniform price auction. Since the price paid is the highest rejected bid and bidders can only bid for a single unit, standard arguments for a private value setting indicate that there is a dominant strategy to bid at one's value. Figure 4 shows that average and median bids are, on average, at value for the final 5 auctions. The points on the average bid line that lie above the dashed value line reveal the presence of some overbidding, which is actually quite common in second price auctions with a single prize (Kagel and Levine, 2012). Note however, that the median bids in Figure 4 are quite close to the value line. Since bids tend to track values, it is not surprising that the efficiency measures for each session are relatively high.

[13] As before, there are 20 ("six take three") ways that the session efficiency averages could have been assigned to these two treatments, and of these, the most extreme was observed, with the three highest efficiencies for the sessions with sealed bids. Therefore, the null hypothesis of no effect can be rejected with a p value of 1/20 = 0.05 (one-tailed test) or 0.10 (two tailed test).

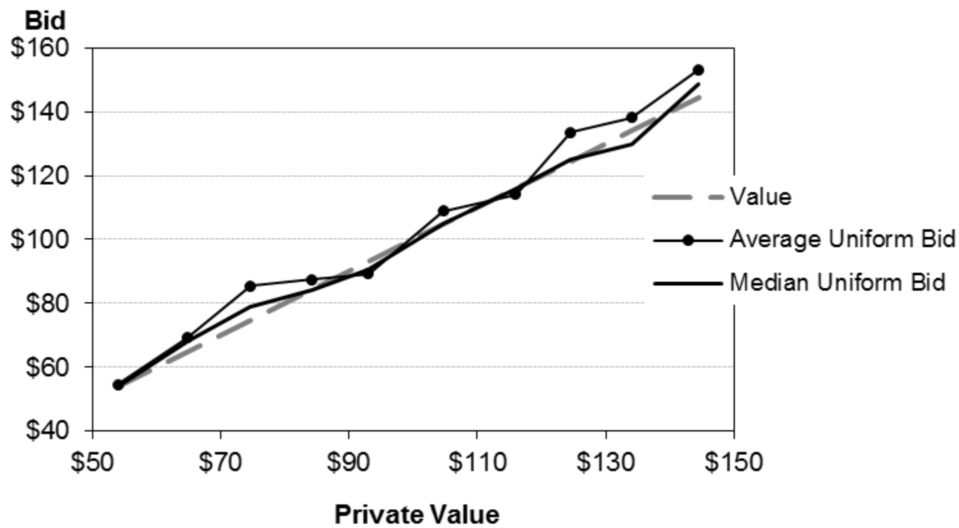

Figure 4. A Comparison of Average and Median Bids for the Uniform Price Auction with the Dominant Strategy Equilibrium Bid Function (Last 5 Auctions)

Of course, the key prerequisite for efficient auction outcomes is that the bids as a function of values be monotonic with enough of a slope to offset "noise" variations in bidding behavior. Therefore, the next step is to derive the Nash equilibrium bid function for the discriminatory auction. For a bidder with a value of v, the expected payoff with a bid of b is given by $(v-b)G(b)$, where $G(b)$ is the probability of winning. With N bidders and M prizes, this is the probability that at $M-1$ or fewer of the $N-1$ other bidders are above b. If the cumulative distribution of bids is denoted by $F(x)$, then the probability of winning with a bid of b is calculated by summing up the binomial probabilities that fewer than M other bidders are above b:

$$(1) \qquad G(b) = \sum_{k=0}^{k=M-1} \binom{N-1}{k} F(b)^{N-1-k}[1-F(b)]^k .$$

In a symmetric Nash equilibrium with an increasing bid function $B(v)$, it must be the case that $b > B(v)$ if and only if $B^{-1}(b) > v$. In particular, a bid of b will win one of the M units if there are no more than $M-1$ other bidders with values above the inverse of one's own bid: $B^{-1}(b)$. If the cumulative distribution function for the private value draws is represented by $H(v)$, the probability of wining in a symmetric equilibrium can be written as the probability that no more than $M-1$ other value draws are higher than a bidder's own draw:

(2) $Q(v) = \sum_{k=0}^{k=M-1} \binom{N-1}{k} H(v)^{N-1-k} [1 - H(v)]^k$,

with a density function denoted by $q(v)$. In equilibrium, $v = B^{-1}(b)$, so the probability of winning with a bid of b can be expressed in terms of the distribution of value draws:

(3) $G(b) \; = \; Q(B^{-1}(b)) \; = \; \sum_{k=0}^{k=M-1} \binom{N-1}{k} H(B^{-1}(b))^{N-1-k} [1 - H(B^{-1}(b))]^k$,

with a derivative denoted by $q(B^{-1}(b))\ B^{-1}{}'(b)$. With this notation, the first-order condition for maximizing $(v-b)G(b)$ is the derivative of $(v-b)Q(B^{-1}(b))$, which can be expressed: $-Q(B^{-1}(b)) + (v - b)\, q(B^{-1}(b))\, B^{-1}{}'(b) = 0$. Since the derivative of the inverse bidding function in the second term can be written as $1/B'(b)$, the derivative can be written: $B'(b) = (v-b)q(B^{-1}(b)) / Q(B^{-1}(b))$. In a symmetric equilibrium, $B^{-1}(b) = v$, so the first-order condition yields a nonlinear first-order differential equation in the equilibrium bidding function:

(4) $B'(v) \; = \; (v - B(v))q(v) / Q(v)$.

This differential equation can be written as $B'(v)Q(v) + B(v)q(v) = vq(v)$. Note that the left side of this equation is the derivative of $B(v)Q(v)$, so the equation can be integrated and rearranged to obtain the Nash equilibrium bid function:

(5) $B(v) = \dfrac{\int_{\underline{v}}^{v} sq(s)ds}{Q(v)} + B(\underline{v})$,

where $\underline{v}$ is the lower bound on the value distribution.

In order to calculate the values of Q(v), we used $M = 6$ prizes and $N = 12$ bidders, and a uniform distribution of private values on [50, 150) with $H(v) = (v - 50)/100$. The last step is to recognize that the lowest possible bid (at a value of 50) will lose for sure, and hence the lowest bid must be no more than a penny below the lowest private value. It turns out to be more convenient to use the differential equation in (4) than the integral expression in (5) to calculate the numerical bid function. We constructed bid equilibrium bid function recursively by starting with a bid of $49.99 for the lowest possible value, $\underline{v}$ = 50, and using the slope resulting from (4) to determine the approximate equilibrium bid for a value of 51, and so on. The resulting equilibrium bid function is shown as the smooth curve in Figure 5, which tracks the actual bid averages for each $10 range in private values, using the last 5 auctions in each session to allow for learning. The connected dots show these averages for the Discriminatory auctions.[14]

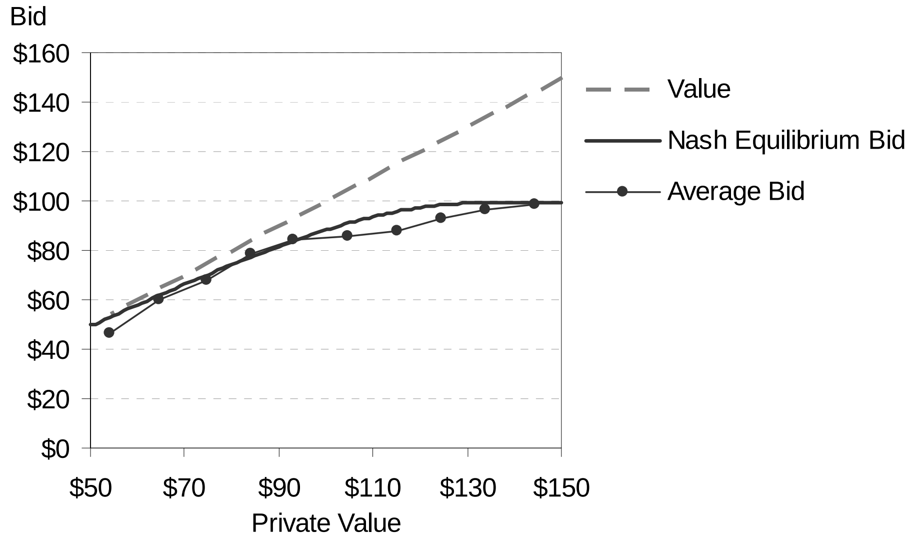

Figure 5. A Comparison of Average Bids for the Discriminatory Auction with the Nash Equilibrium Bid Function

Notice that there is no evidence of overbidding due to risk aversion or other factors such as a "joy of winning" in this auction. Indeed, it is remarkable that the curve in Figure 5 is a *theoretical prediction*, not a line estimated from data. To summarize:

Equilibrium Bidding: *Bids for Discriminative Auctions are close to the Nash equilibrium bidding function, assuming risk neutrality. In the Uniform Price auction, it is a dominant strategy to bid a value, regardless of risk preferences, and average bids are approximately equal to private values, as predicted.*

One interesting divergence between the theoretical predictions and average bids is the tendency to bid a little lower than predicted near the "edge of the cliff" in the $85 bid range in Figure 5. This raises the question of whether observed bids are close to being expected payoff maximizing best responses to the empirical distribution of bids. To answer this question, we ranked the 360 observed bids in all auctions to obtain $F(x)$.

[14] Median bids for each value category (not shown) are quite close to average bids for the Discriminatory auctions, with no systematic bias upward or downward. .

Then we used (1) to obtain the probability of winning for each bid, and using that, we listed the expected payoffs in a spreadsheet for all permitted bids from 10 to 200, to determine the best bid for each possible value *v*. The expected-profit-maximizing responses to the observed bid distribution are shown in Figure 6 as the "optimal response" line, which has steps due to flat segments of $F(x)$. For comparison, the bid averages for each $10 range of value draws are also shown, as connected dots, using all bids for the final five auctions in each session to allow for learning effects. Obviously, observed average bids are approximately equal to the risk-neutral best responses to the observed bid distribution.

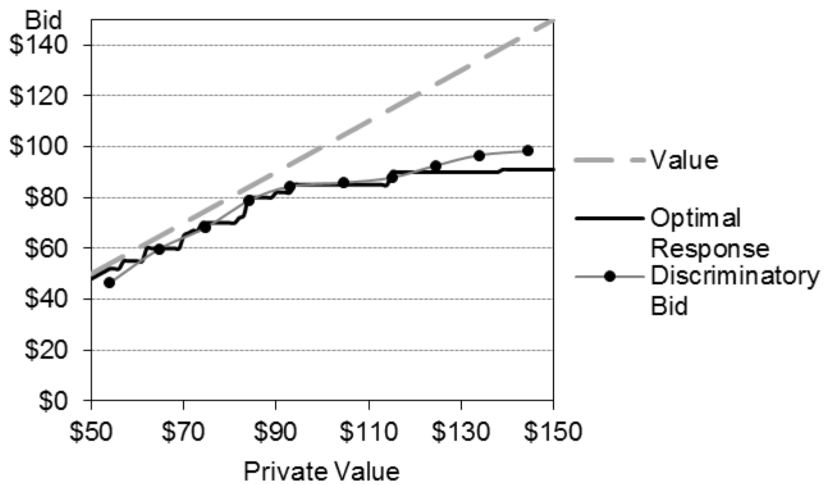

Figure 6. Expected Profit Maximizing Responses to the Cumulative Distribution of Observed Bids in Discriminatory Auctions (Solid Line with Steps) Versus Average Bids for the Final Five Rounds (Line with Dots)

Notice that the empirical bid function is slightly below and parallel with the value line for bids below $80, with a flatter but increasing segment for higher bids. Not surprisingly, the "edge of the cliff" is at about $80. In fact, only about 5 percent (19 of 360) of the bids below $82 are accepted, and 5 percent (18 of 360) of the bids at $82 or above are rejected. This relatively "bright line" cutoff has the effect of limiting aggressive bid reductions, and consequently, there is small chance that a high value bidder will lose the auction to a bidder with a lower value in the Discriminatory Auction.

In contrast, the final bid function for the Shanghai Auction (with an initial bid interval) is roughly a flat line for values above $80, as shown by the connected "x" marks in Figure 7. Both high value bidders and low value bidders make similar bids as a direct result of the narrow band of bid revision, which tends to reduce auction efficiency. Note that the average bids for the Shanghai Auction are below the empirical bid function for the Discriminatory Auction. This is consistent with the observation of lower revenues for the Shanghai Auction in the previous section.

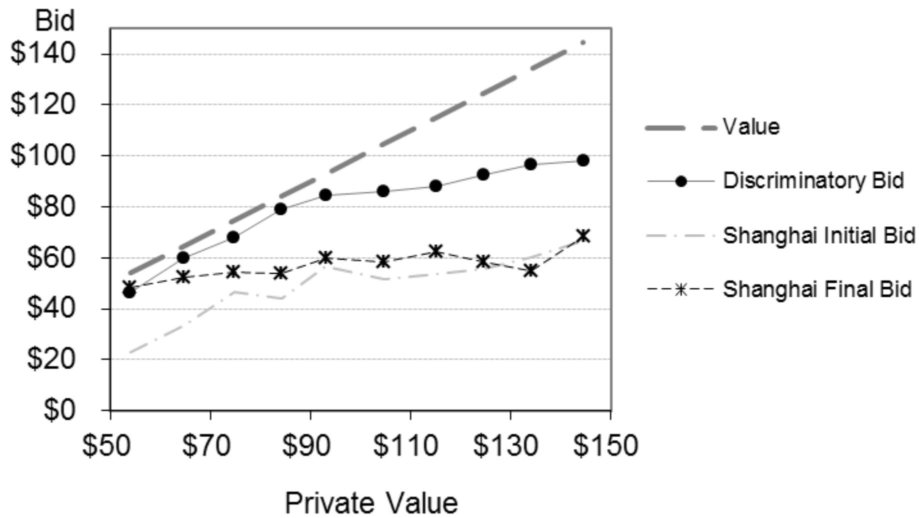

Figure 7. A Comparison of Average Bids for the Discriminatory Auction with Initial Sealed Bids and Final Bids for the Shanghai Auction with an Initial Bid Interval

We conduct further analysis to identify the causes of revenue and efficiency differences by regressing individual bids on private values for each of these auction formats. These regressions are not intended to be structural models of bidding behavior, but rather, to provide "big picture" views of underlying bid-value correlations across auction formats. For any given positive correlation between bids and values, a higher coefficient of determination is indicative of better correspondence between the order of bid prices and the order of values and therefore is associated with higher efficiency.

For the Shanghai Auction, we run separate regressions for the initial bids and the final bids. In order to assess the impact of the narrow band of bid revision, we only include final bids that are within plus or minus \$3 of the cutoff price.[15] These final bids are computed as deviations from the cutoff, allowing for a direct comparison of bids

[15] According to our data, the average gap between the final cutoff price and the initial cutoff price is \$2.94. So a range of plus and minus \$3 of the final cutoff price includes most bid revisions, given no large outliers in calculating the average.

across different bid windows. We expect to see a weak linear correlation between private values and the bids in the bid window, because the narrow band of bid revision induces off-margin bidders to bid around the market clearing price and therefore may result in allocative inefficiency.

Bids are regressed on the private value and controls for bidder-specific and round-specific effects. The four regressions are estimated separately with the linear regression model: $b_{ij} = \beta v_{ij} + \mu_i + \gamma_j + \varepsilon_{ij}$. In this equation, b_{ij} denotes the bid price that bidder i placed in auction j (or period j), v_{ij} represents bidder i's private value in auction j and ε_{ij} is an error term with mean zero. The regression also includes a time-invariant individual effect term μ_i and an auction-specific term γ_j.

Table 3. Fixed Effects Regression Results (Standard Errors)

	Number Observations	Intercept	Private Value Coefficient	R^2
Uniform Price Sealed Bid	360	-23.133 (5.436)	0.978 (0.040)	0.692
Discriminatory Sealed Bid	360	11.483 (3.267)	0.559 (0.024)	0.689
Current Shanghai Interval Initial Bids	351	9.474 (6.218)	0.436 (0.041)	0.556
Current Shanghai Final Bids Near Cutoff	253	0.491 (0.279)	0.003 (0.002)	0.269

Regression results in Table 3 provide evidence that for a given private value, bidders in the discriminatory auction bid higher than bidders in the first stage of the Shanghai auction. In particular, the intercepts for the Discriminatory and Shanghai initial bid rows are similar, but the slopes are somewhat different, and the difference is statistically significant at all conventional levels. As a result, the initial cutoff price for the Shanghai Auction tends to be significantly lower than the cutoff price for the discriminatory auction. This further leads to lower cutoff prices for the final bids of the Shanghai auction, because on average the final cutoff price does not increase by much during the bid revision phase. The fourth row of Table 3 reports the estimation results for the final bids that are within plus or minus \$3 of the cutoff price for the Shanghai auction. Based on the estimated slope, there is only a very slight upward trend in prices (about a three cents for each dollar increase in private value), suggesting that the final bids in the bid window are relatively stable across the range of private values. Not surprisingly, the R^2 is low for this regression, which implies that the allocation of prizes for the bids in the bid window rarely follows the positive linear relationship suggested by the regression results. Note that the number of observations for bids in the bid window is more than two-thirds of the total number of bids. Therefore, regression results presented in the fourth row are consistent with the clear revenue and efficiency differences observed in our experiments.

Given the inefficient allocations that result from the low bid-value correlations in the bid window, efficiency is further reduced by a tendency for high value bidders to "jump-bid" into the bid window. Since the decision to revise a bid can be treated as a binary dependent variable, we run a probit regression on the differences between the initial bids and the initial cutoffs determined by the lowest accepted initial bid. Figure 8 shows the predicted probabilities that high initial bidders will revise downward into the bid window.[16] The horizontal axis is the difference between the initial bid and the lowest accepted bid. The prediction dots are plotted for each of the bid differences actually observed.

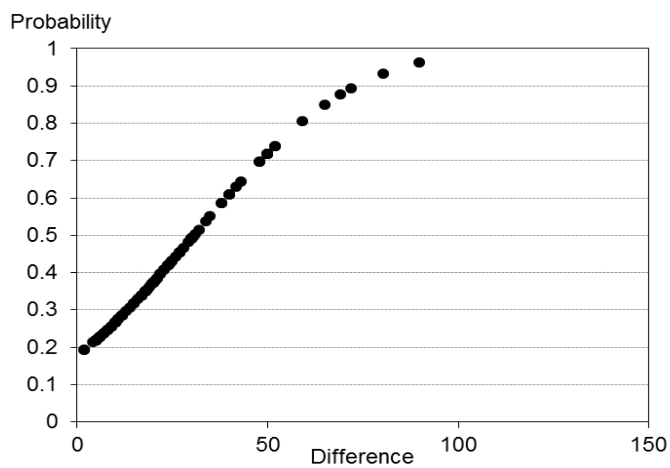

Figure 8. Probability Predictions for Jump Bidding into the Bid Window for the Shangahi Auction with an Interval of Initial Bids before Bid Revisions, as a Function of the Difference Between Initial Bids and the Initial Cutoff Bid

The predictions in Figure 8 indicate that the higher an initial bid is above the bid window, the more likely it will be revised downward. Five of the initial bids were intentionally placed above value, so these bids would have to be reduced to avoid losses.

[16] We exclude those bids that are higher than the initial cutoffs but are surpassed by the bid window as the cutoff price rises. We also exclude people with values that were below the initial cutoff.

After eliminating these bids, the average of the initial bids that are subsequently reduced is $29.50 above the initial cutoff. In contrast, the average of the initial bids that are not revised is only $17.50 above the cutoff. Admittedly, a lower initial bid does not necessarily imply a lower private value. Some high-value bidders in the last 5 rounds do tend to place a bid at a price slightly above the initial cutoff but also high enough to avoid competition in the bid window. Thus their bids are kept low but safe from being displaced in the bid revision period. Despite these exceptions, high initial bids are generally associated with high values, and the tendency for these high bidders to jump down into the bid window tends to reduce their chances of winning and further reduce auction efficiency. It is interesting to note that the bidder with the lowest earnings in the Shanghai auction treatment (who only earned about $4) is someone who revised the initial bid in all 10 auctions. These bid revisions put that bidder into the "mud pit" of bid sniping and scrambling, where the constraints on bid revisions prevent a bidder with a high value from expressing a strong desire to win by bidding high.

V. Discussion

This paper uses a laboratory experiment to evaluate the performance of the Shanghai Auction for automotive license plates, which was modified in 2008 in an effort to contain price increases. The modifications, which constrain the numbers and range of bid changes during the auction process, tend to produce a noisy, somewhat flat correlation between bids and bidders' private values in a very tight "bid window," which causes major reductions in the efficiency of the initial allocations. In fact, efficiencies for the initial auctions in each laboratory session were not much higher than would be achieved with a purely random (e.g. lottery) allocation. These efficiencies were considerably lower than those produced by standard multi-unit auction formats (discriminatory and uniform price auctions). One advantage of experiments is that an understanding of the efficiency-reducing processes (narrow bid window with low bid/value correlations) can suggest modifications that may improve auction performance.[17]

[17] One of the treatments considered in Liao (2012) involved a wider bid window, which generated some efficiency improvements without a large increase in auction prices.

Coase (1959) argued persuasively for clear definitions of property rights associated with auction-based allocations. In an effort to curb the influence of speculators and brokers, however, Shanghai officials have imposed serious constraints on second-hand license plate transactions, which now require a three-year wait and the sale of the car. These restrictions impede such market corrections. The license plate quotas do limit traffic congestion, of course, but the policy is best described as one of "cap and no trade."

The laboratory version of the Shanghai auction, with limited opportunities for bid revisions, does have its intended effect of reducing prices and revenues, as compared with the sealed bid format that permits no bid revisions. Although these price reductions are presumably motivated by political considerations, there can be economic benefits of auction-based allocations that generate significant revenues, over a billion dollars annually for the Shanghai auction. In particular, these revenues may displace distortionary taxes and facilitate transportation investments that provide better options to those who cannot afford the price of a license plate that often equals the price of the vehicle.

The primary policy alternative is a lottery-based allocation, which has been recently adopted in Beijing and the Southwest city of Guiyang, China. Popular press accounts of the hundreds of thousands of eager applicants (and their friends and relatives) suggest that rent-seeking costs and misallocations associated with lotteries may be considerable. Inch (2012, p. 215), for example, describes several scams involving bogus lawsuits and contractual indemnifications associated with quasi-legal title transfers, in combination with under-the-table payments on a magnitude that is comparable with Shanghai auction prices.

References

Abbink, K., Brandts, J., and McDaniel, T. (2003) "Asymmetric Demand Information in Uniform and Discriminatory Call Auctions: An Experimental Analysis Motivated by Electricity Markets," *Journal of Regulatory Economics*, 23(2), March, 125-44.

Abbink, K., Brandts, J., Pezanis-Christou, P. (2006) "Auctions for Government Securities: A Laboratory Comparison of Uniform, Discriminatory and Spanish Designs," *Journal of Economic Behavior & Organization*, 61(2) 284-303.

Anderson, S., J. Goeree, and C. A. Holt (1998) "The All-Pay Auction: Equilibrium with Bounded Rationality," *Journal of Political Economy*, 106(4), August, 828-853.

Armantier, Olivier, Charles Holt, and Charles Plott (2012) "A Reverse Auction for Toxic Assets," Discussion Paper, University of Virginia.

Banks, J., Olson, M., Porter, D., Rassenti, S., Smith, V.L. (2001) "Theory, Experiment and the Federal Communications Commission Spectrum Auctions," *Journal of Economics Behavior and Organization*, 51, 303-350.

Burtraw, D., J. Goeree, C. A. Holt, E. Myers, K. Palmer, and W. Shobe (2010) "Price Discovery in Emission Permit Auctions," in R. M. Isaac and D. A. Norton, eds., *Experiments on Energy, the Environment, and Sustainability, Research in Experimental Economics, Vol. 14*, Emerald Group Publishing, Bingley:UK, 11-36.

Burtraw, D., J. Goeree, C.A. Holt, E. Myers, K. Palmer, and W. Shobe (2009) "Collusion in Auctions for Emissions Permits: An Experimental Analysis," *Journal of Policy Analysis and Management*, 28(4), 672-691.

Chen, Tongkui (2008) "The Transformation of the License Plate Auction in Shanghai," *Decision Making*, April, (in Chinese).

Coase, Ronald H. (1959) "The Federal Communications Commission," *Journal of Law and Economics*, 2: 1-40.

Cox, J. C., Smith, V.L., and Walker, J.M. (1988) "Theory and Individual Behavior of First-Price Auctions," *Journal of Risk and Uncertainty*, 1(1) 61-99.

Friedman, Milton (1960) *A Program for Monetary Stability*, New York: Fordham University Press, 1960.

Goeree, J. and C. A. Holt (2010) "Hierarchical Package Bidding: A Paper & Pencil Combinatorial Auction," *Games and Economic Behavior*, 70, September, 146-169.

Goeree, J.K. and Offerman T. (2002) "Using First-Price Auctions to Sell Heterogeneous Licenses", *mimeo*, University of Amsterdam.

Hazlett, T. W. and Robert J. Michaels (1993) "The Cost of Rent Seeking: Evidence from Cellular Telephone License Lotteries," *Southern Economic Journal*, 59(3), January, 425-435.

Holt, C., W. Shobe, D. Burtraw, K. Palmer, and J. Goeree (2007) *Auction Design for Selling CO2 Emissions Allowances Under the Regional Greenhouse Gas Initiative*, RGGI website**:** http://www.rggi.org/docs/rggi_auction_final.pdf.

Holt, C. A. and R. Sherman (1982) "Waiting-Line Auctions," *Journal of Political Economy*, 90(2), April, 280-292.

Inch, Jason (2012) *China's Economic Supertrends*, Victoria, British Columbia: InChina Publishing.

Kagel, John and D. Levine (2012) "Auctions: A Survey of Experimental Research, 1995-2010," forthcoming in Kagel, J. and A. Roth, eds., *Handbook of Experimental Economics, Vol. II*, Princeton, N.J.: Princeton University Press.

Klemperer, Paul (2002) "What Really Matters in Auction Design," *Journal of Economic Perspective*, 16(1), 169-189.

Liao, Evan Zuofu (2012) "The Pursuit of Less Revenue: An Experimental Test of Shanghai Designs for the Automobile License Plate Auction," Distinguished Majors Thesis, Department of Economics, University of Virginia.

Lopomo, G., Marx, L., McAdams, D., Murray, B. (2011) "Carbon Allowance Auction Design: An Assessment of Options for the United States," *Review of Environmental Economics and Policy*, doi:10.1093/reep/req024.

Luo, Wei (2008) "Multi-Unit Auctions Theory and Its Application in Automobile License Plate," Shanghai Jiaotong University, (in Chinese).

McMillan, John (1994) "Why Auction the Spectrum," *Telecommunication Policy*, 19, 191-199.

Plott, C.R. (1997) "Laboratory Experimental Testbeds: Application to the PCS Auction," Jo*urnal of Economics and Management Strategy*, 6(3) 605-638.

Rassenti, S., Smith, V. and Wilson B. (2003) "Discriminatory Price Auctions in Electricity Markets: Low Volatility at the Expense of High Price Levels," *Journal of Regulatory Economics*, 23(2) 109-23.

Xu, B. (2005) "I Forecast Shanghai License's Price," http://www.ceibs.edu/facultyCV/xubin/ shanghai.pdf (in Chinese).