Explaining Early Bidding in Informationally-Restricted Ascending-Bid Auctions

Sungjin Cho, Seoul National University
Harry J. Paarsch, University of Central Florida
John Rust, Georgetown University

September 25, 2022

Overview of *work in progress!*

- We analyze a large number of continuous time ascending bid auctions of rental cars conducted online over 2 minute intervals.
- Due to concerns about bidder collusion, the rental company designed a unique dynamic auction format.
- Bidders only know the amount of their own bids and an indicator whether their bid is the highest so far.
- *Bidders are never able to observe the bids placed by competing bidders, and thus do not even know the number of other bidders bidding in any given auction.*
- We are aware of only one other auction that has informational restrictions similar to the Korea auctions: auctions of certificates of deposit (CDs) by the state of Texas.
- Groeger and Miller *Journal of Econometrics* (2021) provide an empirical analysis bidding under of this auction format under the assumption that bidding in the auction follows a Perfect Bayesian Equilibrium (PBE).

How does this relate to market design?

- Auctions are markets (often "small" ones, i.e. with a relatively small number of bidders). Some of the earliest work on market design can be traced back to work on auctions.
- There are many different formats for auctions (i.e. rules for running auctions) such as English (open acscending bid) auctions, first price sealed bid auctions, etc.
- One could also use non-auction mechanisms to sell an item: e.g. selling lottery tickets
- Which is the best format to use? Choosing a 'best" auction mechism is clearly a classic example of *market design*
- Landmark paper: *Optimal Auction Design* (1982) by Roger Myerson *Mathematics of Operations Research*
- Characterized the optimal auction mechanism for selling a single object to a finite (known) set of bidders who have valuations for the object with known distributions that are independently distributed (the *Independent Private Values* (IPV) assumption)
- Optimal market design (maximizes expected revenue to the seller): A second-price auction (Vickrey auction) with a reservation price.

How do second-price auctions work? Theory vs Practice

- According to *Wikipedia* "The auction was first described academically by Columbia University professor William Vickrey in 1961 though it had been used by stamp collectors since 1893. In 1797 Johann Wolfgang von Goethe sold a manuscript using a sealed-bid, second-price auction."
- The second price auction is *ex post efficient* i.e. the bidder with the highest valuation wins the auction.
- If the auction is *symmetric* (all bidder's values are *IID*) the first price auction is also *ex post* efficient too.
- Under symmetry, the *Revenue Equvalence Theorem* predicts that expected revenue in a first price auction and second price auction are the same.
- However in practice (e.g. laboratory experiments) show *overbidding in second price auctions*

Overbidding in second-price auctions

Cooper and Fang, 2008, "Understanding Overbidding in Second Price Auctions: An Experimental Study"

> *In laboratory experiments, however, subjects are found to exhibit a consistent pattern of overbidding. Kagel et al. (1987) found that the actual bids are on average 11% above the dominant strategy bids. Kagel and Levin (1993) found that about 62% of all bids in their five-bidder SPA sessions exceed the bidder's value, while only 8% of all bids were below it. Both Kagel and Levin (1993) and Harstad (2000) further reported that experience has only a small effect in reducing overbidding in SPA.*
>
> *We found that small and medium overbids are more likely to occur when bidders perceive their rivals to have similar values, supporting a modified "joy of winning" hypothesis but large overbids are more likely to occur when bidders believe their opponents to have much higher values, consistent with the "spite" hypothesis.*

Collusion in second-price auctions

Paul Klemperer (2002) "What Really Matters in Auction Design" *Journal of Economic Perspectives*

> *Economists are proud of their role in pushing for auctions; for example, Coase (1959) was among the first to advocate auctioning the radio spectrum. But many auctions — including some designed with the help of leading academic economists — have worked very badly.*

George Mailath and Peter Zemsky (1991) "Collusion in Second Price Auctions with Heterogeneous Bidders" *Journal of Economic Theory*

> *We show that efficient collusion by any subset of bidders in second price private value auctions is possible, even when the bidders are heterogeneous. An important property of efficient collusion is that a bidder's net payoff from participating in collusion is independent of her valuation.*

Robust market/mechanism design

- Mechanism design theory is brilliant theory, but requires strong assumptions to get crisp results and characterizations
- Economists who actually do auction design in practice (e.g. Paul Milgrom and Larry Ausubel, etc) use *judgement* because it is not clear that that strong assumptions necessary are valid in practical situations.
- Also, many real world auctions are far more complex and involve wider considerations (e.g. social welfare) that are difficult to handle in theory than, say, expected revenue maximization.

The Wilson Critique of Mechanism Design

It is noteworthy that our optimality result concerns a property of this simple, well-motivated mechanism across a range of possible environments and sizes of the market, rather than simply for a single, fixed environment and size of market. Our result thus responds to the Wilson Critique (Wilson (1987) of mechanism design. Wilson criticized this field for focusing upon the problem of designing a mechanism explicitly for each specific problem (e.g., as determined here by the specification of an environment and a market size). An economic consultant asked for advice on the selection of a mechanism may not know all the parameters that specify the problem, and the parameters may change over time; theoretical results that describe how the mechanism should be chosen assuming detailed knowledge of the problem may thus have little value to the consultant. A more meaningful task for mechanism design is to establish the sense in which a simple mechanism performs reasonably well across the variety of problems that might be encountered in practice, which is the nature of our results.

The k-double auction

Satterthwaite and Williams (2002) "The Optimality of a Simple Market Mechanism" *Econometrica*

> *Strategic behavior in a finite market can cause inefficiency in the allocation, and market mechanisms differ in how successfully they limit this inefficiency. A method for ranking algorithms in computer science is adapted here to rank market mechanisms according to how quickly inefficiency diminishes as the size of the market increases. It is shown that trade at a single market-clearing price in the k-double auction is worst-case asymptotic optimal among all plausible mechanisms: evaluating mechanisms in their least favorable trading environments for each possible size of the market, the k-double auction is shown to force the worst-case inefficiency to zero at the fastest possible rate.*

Illustration of the k-double auction

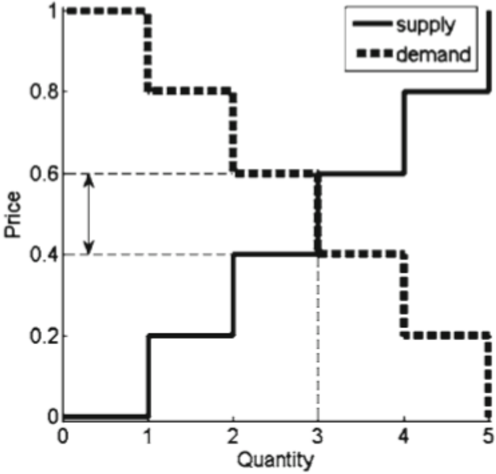

Static vs Dynamic Auction Mechanisms

- The k-double auction is an example of a *static mechanism* — it is a generalization of the idea of a sealed bid auction for an single sided auction (one object for sale)
- But there is also the *continuous double auction* (CDA, also known as an "open outcry" auction) — it is a generalization of an open ascending bid auction, i.e. an English auction
- The CDA is easy to play in the lab and shows high efficiency and convergence to close to the "Walrasian" outcome, a result demonstrated in hundreds of lab experiments by economists such as Vernon Smith and Charlie Plott
- Paradox: The CDA is a continuous time game of incomplete information and has never been fully solved, not even by the most brilliant auction theorists such as Robert Wilson.
- Which mechanism is better? The static k-DA or the CDA?
- Designs of many auctions are motivated by the *linkage principle* (release of information during the auction leads to higher bids) and indeed the FCC spectrum auctions are dynamic combinatorial auctions. No theorist can solve these auctions either.

Preliminary Results of Our Analysis

- We characterize bidding behavior in the Korean auction and show *frequent early bidding.*
- We conjecture that *early bidding will not occur in a PBE.* That is, we conjecture that the only equilibrium is an *uninformative equilibrium* where all bidders wait to the last instant to submit bids.
- The uninformative equilibrium always exists and and is strategically equivalent to the equilibrium of a static first price sealed bid auction.
- We illustrate a two bidder, two period example where the only PBE is the uninformative equilibrium.
- We introduce a behavioral boundedly rational model of bidding behavior that can explain early bidding.
- Bidders learn enough from early bidding to avoid overpaying later in the auction. As a result, expected revenue in the Korean auction is lower than expected revenue in a first price sealed bid auction.

Auction Format

- We study a data set of over 8600 auctions of cars the rental company sold between 2003 and 2007.
- There is a universe of 90 professional bidders who participate in these auctions: most are auto dealers.
- The auctions are conducted over the internet in back-to-back 2 minute auctions. Bidders have the ability to physically inspect the cars being auctioned prior to the auction itself.
- On average there are 7 bidders participating in any individual auction, and 59 bids are placed during the 2 minute auction.
- However due to the auction rules, the bidders in any single auction do not know how many other bidders are present, what their bids are, and they only learn the winning price in the auction if they hold the highest bid at the end of the two minute auction.

Key Findings

- We show that there is a substantial amount of early bidding in these auctions, even though a game-theoretic analysis suggests that the informational restrictions should create strong incentives for *bid sniping* — i..e. waiting to submit a bid only in the last instant of the auction.
- We define an *informative equilibrium* of the dynamic auction to be one in which there are bids placed before the final instant of the auction using strictly monotonic bid functions.
- In a two bidder, two period example, we show there is no informative PBE and thus no early bidding.
- This creates a challenge: can the early bidding we observe in these auctions be explained as a PBE outcome?

Model

- Due to the difficulty of computing PBE and because it is not clear that there is an informative PBE that would be consistent with the bidding behavior we observe in these auctions, we adopt an alternative modeling approach.
- We develop a *behavioral DP bidding model* that assumes *bidders have rational beliefs* about the stochastic process for the high bid price in the auction.
- Using these beliefs, we solve a dynamic program to determine the optimal bidding strategy implied by these beliefs.
- We define an *rational expectations equilibrium* as a self-confirming system of beliefs, i.e. where the actual stochastic process for the highest bid during the auction is approximately equal to bidders' beliefs about this stochastic process.

Accounting for learning

- Our approach involves *learning* but employs a simpler model of *experiential learning* rather than full Bayesian updating.
- Our model is appropriate for *experienced bidders* who have participated in many auctions, and thus have well-defined and fixed beliefs about the stochastic process for the high bid in the auction.
- If a bidder has the highest bid at t, then they know it. But if the bidder does not have the highest bid at t, they must predict it based on a rational belief of the probability distribution of the high bid.
- We show that there is a significant reduction in uncertainty by learning that one has the high bid prior to the end of the auction.
- Thus the motivation for early bidding is to gather information about the current high bid, helping bidders to win the auction without paying more than necessary.

Behavioral Market Design

- A major rental car company experimented with alternative mechanisms for selling its cars.
- It originally used *open outcry auctions* held at each rental car location, but the owners suspected bidding collusion that lowered their bids.
- It developed its own unique online bidding system (via the Internet) to try to defeat possible collusion.
- Around 2012 it abandoned its own online auction system and sold cars via an auction house which used a sequence of English auctions.

Is the linkage principle valid?

- *Linkage principle* if the bidders' valuations are affiliated, auctions that release more information over the course of the auction will result in higher average prices compared to auctions that reveal less information.
- *Revenue Equivalence principle* if the bidders' valuations are *IID* then auctions with the same probability of assigning a winner generate the same expected revenue.
- Our 2014 *JINDEC* paper, "Is the Linkage Principle Valid? Evidence from the Field" compared the alternative auction formats the rental company used with respect to mean revenue, and found evidence consistent with the linkage principle — the prices from the auction house were 10% higher than the company's online auction system.
- The dynamic rental auction releases more information than a static first price auction, yet we find that the dynamic auction results in *lower* expected revenue.

Auction 1, January 26, 2005

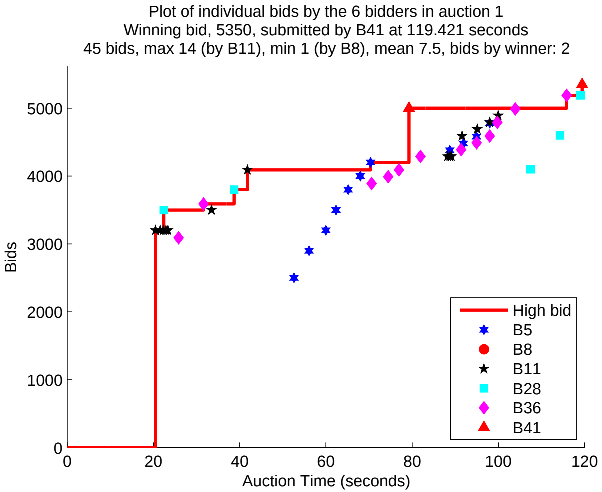

Frequent use of "probing/testing" strategies

- Notice that bidder B5 makes frequent bids, each slightly higher than the previous one.
- It seems evident that B5 was trying to "probe" or "test" the market to learn what the current high bid was.
- However B5 never succeeded in placing a highest bid, and only learned that the high bid was higher than each of its sucessive bids.
- B5's last bid was $4500 placed less than 30 seconds remaining in the auction, after which B5 gave up and declined to submit any further bids.

Auction 3, January 26, 2005

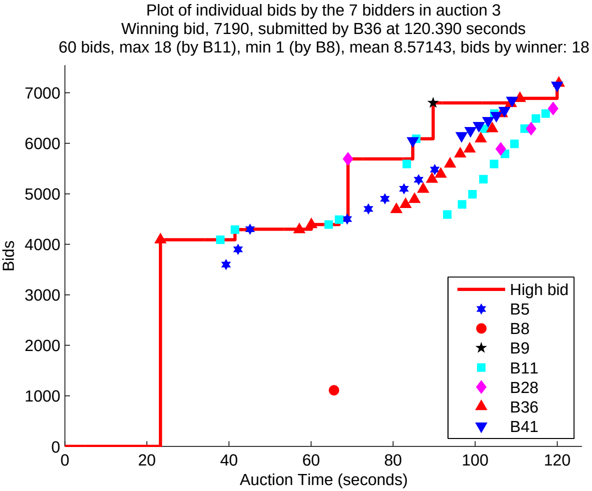

Auction 394 — bid sniping

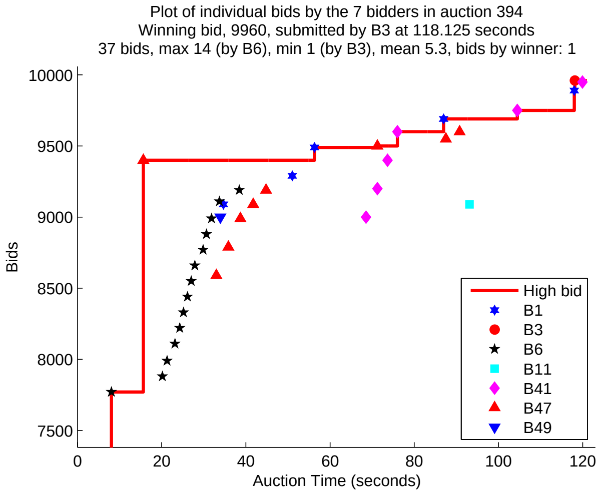

Auction 29 — early high bidder

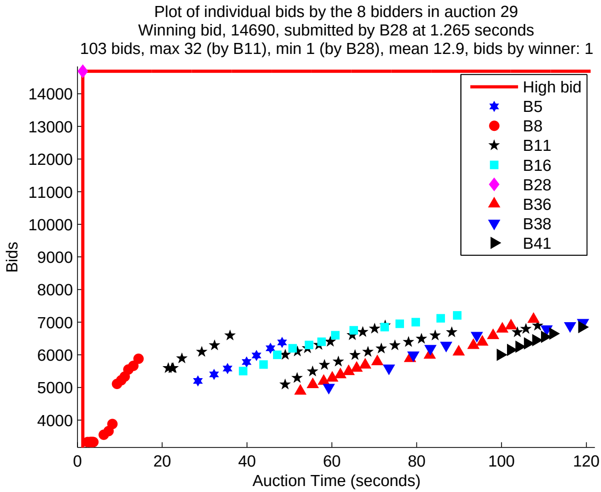

Auction 32 — a “crazy” bidder B11

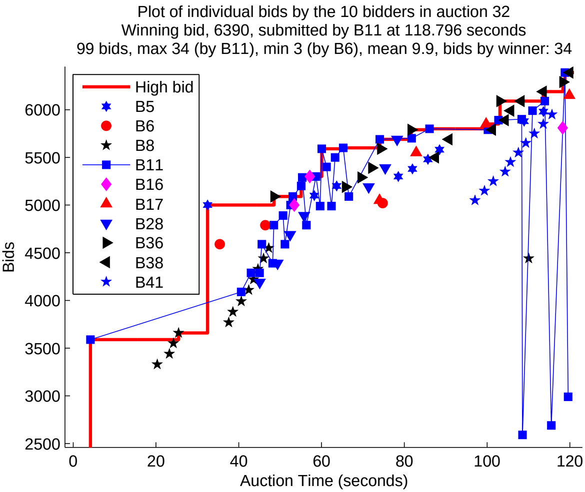

Rescaled bid trajectories

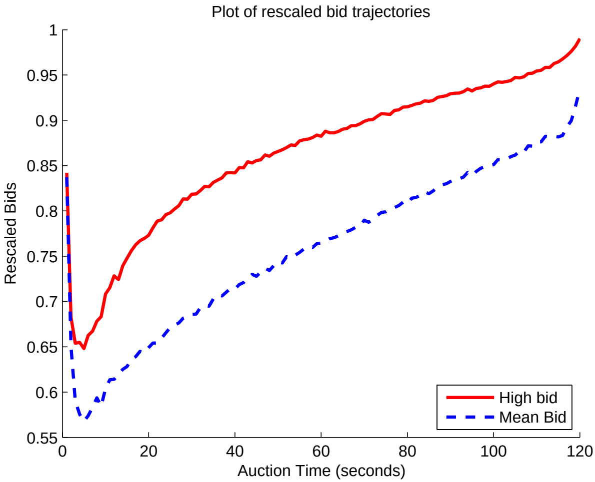

Rescaled bid trajectories — high early bids removed

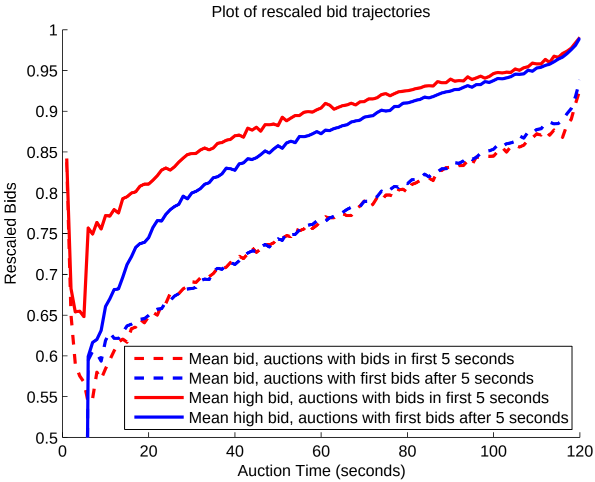

CDFs of first bid and winning bid by elapsed time

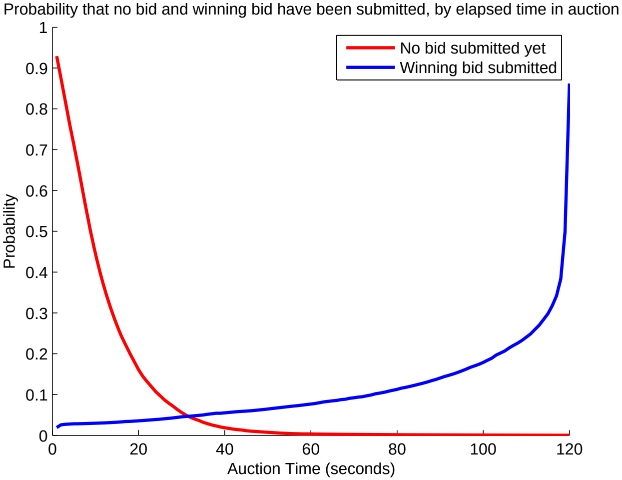

Distribution of time winning bid was submitted

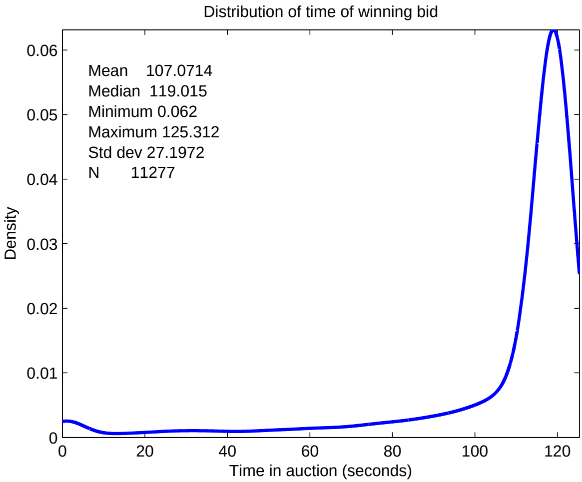

Distribution of percentile order of winning bid

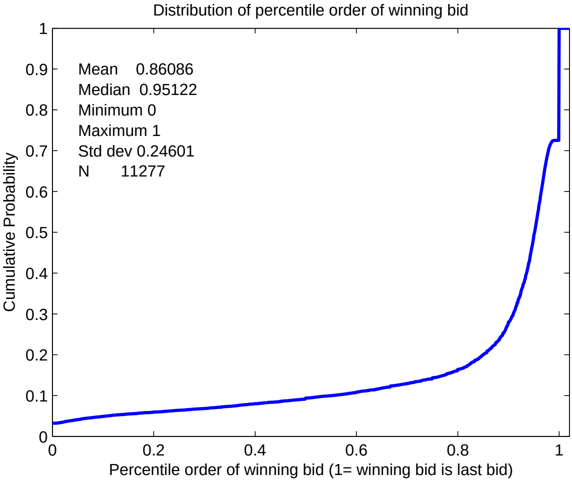

Bids submitted per second in the auction

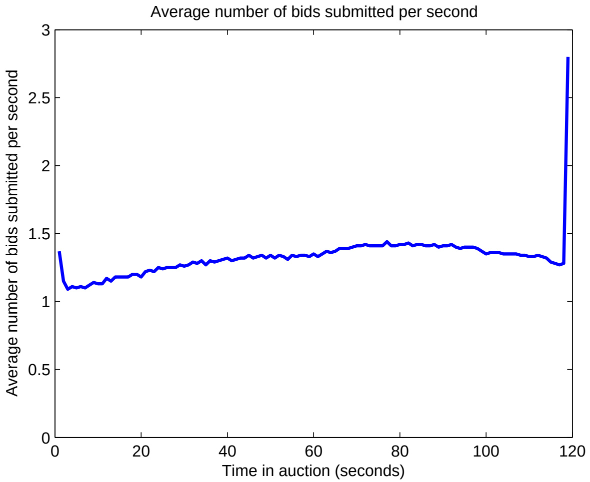

CDFs of rescaled bids by time in the auction

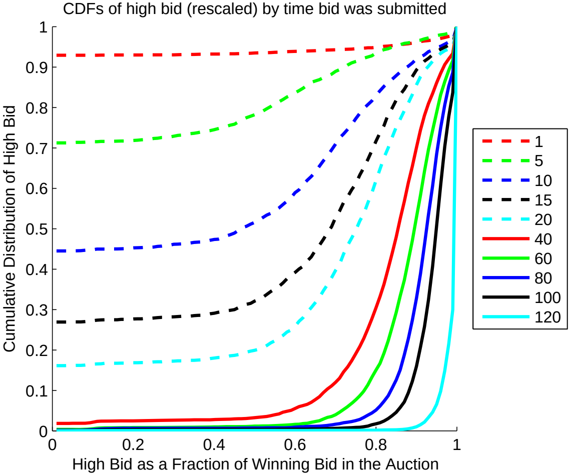

Distribution of number of bidders per auction

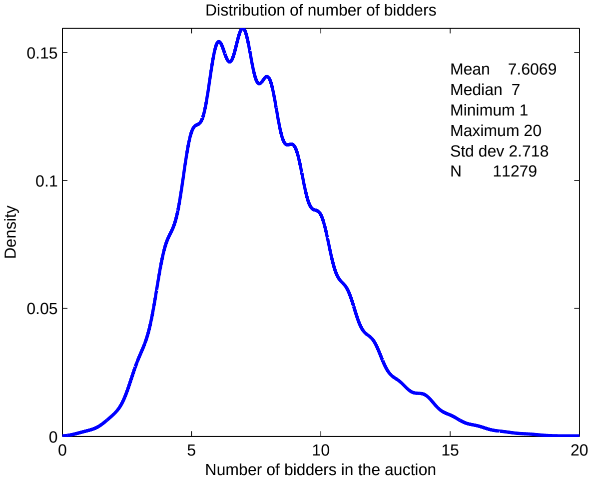

Distribution of number of bids per auction

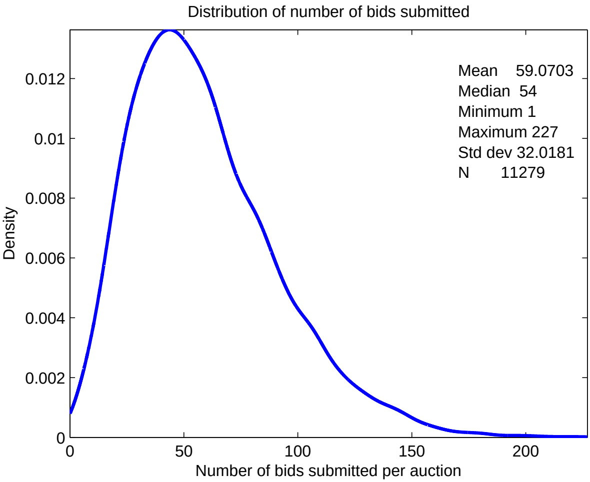

Graph of the simulated auction 998 we played yesterday

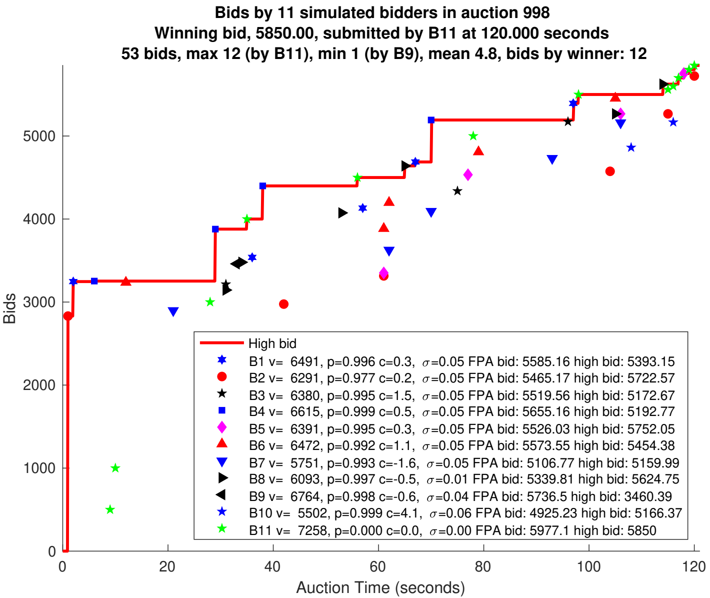

Graph of the actual auction 998

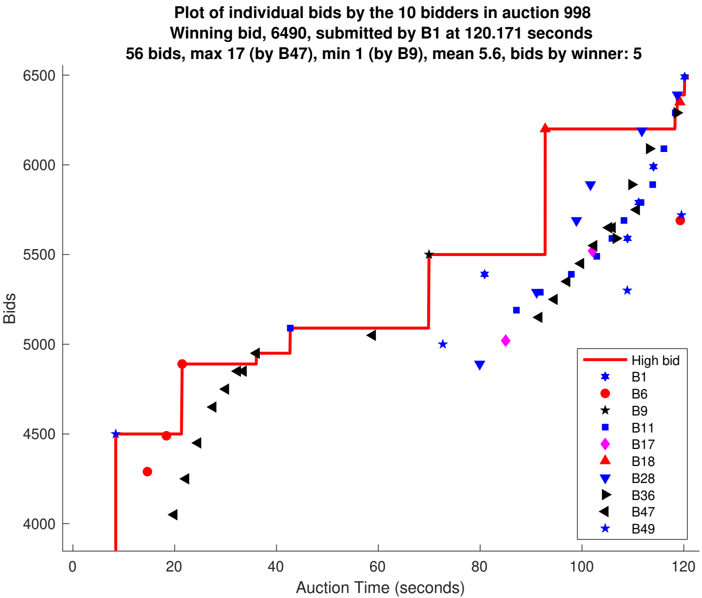

A dynamic structural approach to bidding

- Is the bidding behavior we observe consistent with a Perfect Bayesian Equilibrium (PBE) of the auction, formulated as a dynamic game of incomplete information?
- **Definition** A *Perfect Bayesian Equilibrium* of a dynamic game of incomplete information is a subgame perfect equilibrium, where players' beliefs are updated using Bayes rule wherever possible.
- "wherever possible" means that it may not be possible to use Bayes Rule for certain off the equilibrium path deviations, where certain events may occur that have zero probability under one or more of the players' *equilibrium beliefs.*

The uninformative equilibrium

- **Proposition** In a discrete time approximation to the dynamic auction, the *uninformative equilibrium* is always a PBE of the dynamic auction game.
- In an uninformative equilibrium, players do not bid, or submit bids of 0 in times $t = 1, 2, \ldots, T - 1$ of the auction, and in time T they submit bids equal to those that they would submit in a single shot first price sealed bid auction.
- In an uniformative equilibrium, the players do not bother trying to test/probe in the early stages of the auction, so the value of learning is zero since there is *no learning in this equilibrium.*

Bid Sniping in Dynamic Second Price Auctions

Ockenfels and Roth "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet" *American Economic Review* 2002.

> *In fact, sniping in an auction with a fixed deadline, in which very late bids have some probability of not being successfully transmitted, need not depend on the presence of irrational bidders. There can be equilibria even in purely private-value auctions in which bidders have an incentive to bid late, even though this risks failing to bid at all. This kind of equilibrium can be interpreted as a kind of implicit collusion against the seller, because it has the effect of probabilistically suppressing some bids, and hence giving higher profits to the successful bidders. The probability that some late bids will not be successfully transmitted is a risk for each bidder, but a benefit for his opponents, and it is this 'public good' aspect of the risk of bidding late that creates the possibility of a profitable collusive late-bidding equilibrium in eBay (but not in Amazon).*

Logic for Bid Sniping: Avoiding Bidding Wars

Ockenfels and Roth also quote from a website `Esnipe.com`, a site that offers to automatically place a predetermined bid a few seconds before the end of the eBay auction, that nicely summarizes some of these reasons but also speaks to the risks involved:

> *There are many reasons to snipe. A lot of people that bid on an item will actually bid again if they find they have been outbid, which can quickly lead to a bidding war. End result? Someone probably paid more than they had to for that item. By sniping, you can avoid bid wars. That's not all. Experienced collectors often find that other bidders watch to see what the experts are bidding on, and then bid on those items themselves. The expert can snipe to prevent other bidders from cashing in on their expertise . . . Will esnipe guarantee that my bids are placed? We certainly wish we could, but there are too many factors beyond our control to guarantee that bids always get placed.*

Do informative equilibria exist?

- The bidding data allow us to easily reject the hypothesis that all bidders are playing uninformative PBE bidding strategies.
- The significant frequency of *bid sniping* could indicate that *some* bidders are trying to play the uninformative equilbrium.
- However if some players are deviating from the uninformative equilibrium, playing the uninformative equilibrium (i.e. bid sniping) may no longer be a best response.
- Can the bidding behavior we observe be rationalized as *some* PBE of this game?

Difficulties of computing nontrivial PBEs

- The main difficulty is that all bidders must be endowed with *priors* over a) the number of bidders in the auction, and b) their valuations. These beliefs must be updated at each instant based on the history of bids made so far.
- Even if the history is very limited due to the informational restrictions of this auction, the history for each player includes at least, a) the current time t, b) the player's own history of bids, and c) whether the player's bid is the highest or not.
- It is extremely challenging to compute a posterior distribution over these quantities, and the dimensionality of the posterior is effectively infinite-dimensional (unless the posterior could be shown to be a member of a conjugate prior class, which seems unlikely).

Costs and benefits of informative bidding

- The main gain to placing "serious" bids early in the auction is to gather significant information on what the high bid is, and to use this to try to win without overpaying.
- However there are at least two costs of placing a serious bid: a) the bidder could mistakenly overbid and the auction rules commit the bidder to pay the *highest* bid submitted during the two minute auction.
- and b) by bidding, the bidder provides information to other bidders that could affect their subsequent bidding behavior to the detriment of the bidder in question.

Non-existence of an informative PBE, a 2x2 example

- Consider a symmetric equilibrium where in period $t = 1$ both bidders submit bids according to a single bid function $b_1(v)$, where $f(v)$ is the density of the bidders' valuations.
- Even if the bid function in period $t = 1$ is symmetric and given by $b_i = b_1(v_i), = 1, 2$ where v_i is the valuation of bidder i (a realization of the random variable $\tilde{v}$ with density $f(v)$), the relevation of information about which bid is the highest in period $t = 1$ results in *endogenous asymmetry* in the bid functions at time $t = 2$.
- if $b_1(v_1) > b_1(v_2)$ and b_1 is a strictly monotonic bid function (a necessary condition for an informative equilibrium), this implies that $v_1 > v_2$.

Non-existence of an informative PBE, a 2x2 example

- The information from period 1 bids about which of the two bidders has the highest valuation is the source of informational asymmetry in period 2.
- Suppose bidder 1 learns that he has the higher valuation. Then bidder 1's posterior belief of bidder 2's valuation in period 2 is $F(v)/F(v_1)$ where $F(v)$ is the prior belief of the CDF of valuations in period 1.
- For bidder 2, their posterior belief of bidder 1's valuation in period 2 is given by $[F(v) - F(v_2)]/[1 - F(v_2)]$.
- Let $b_{2,h}(v)$ be the period two bid function for a bidder who learns they had the high bid in period 1, and let $b_{2,l}(v)$ be the bid function of a bidder who learns their bid was the low bid in period 1.

Period 2 equilibrium bid functions

Solve the game by backward induction, In period two the equilibrium bid functions solve

$$
\begin{aligned}
b_{2,h}(v,b) &= \underset{b' \geq b}{\operatorname{argmax}} (v-b') \times \\
&\quad \int_0^v I\{b_{2,l}(v', b_1(v')) \leq b'\} f(v') dv' / F(v) \\
b_{2,l}(v,b) &= \underset{b' \geq b}{\operatorname{argmax}} (v-b') \times \\
&\quad \int_v^\infty I\{b_{2,h}(v', b_1(v')) \leq b'\} f(v') dv' / [1-F(v)]
\end{aligned}
$$

Period 1 equilibrium bid function

$$
\begin{aligned}
b_1(v) \quad = \quad & \underset{b}{\operatorname{argmax}}(v - b_{2,h}(v, b)) \times \\
& \left[\int_0^v I\{b_{2,l}(v', b_1(v')) \leq b_{2,h}(v, b)\} f(v') dv' \right] + \\
& (v - b_{2,l}(v, b)) \times \\
& \left[\int_v^\infty I\{b_{2,h}(v', b_1(v')) \leq b_{2,l}(v, b)\} f(v') dv' \right].
\end{aligned}
$$

Equilibrium bid functions in period 2

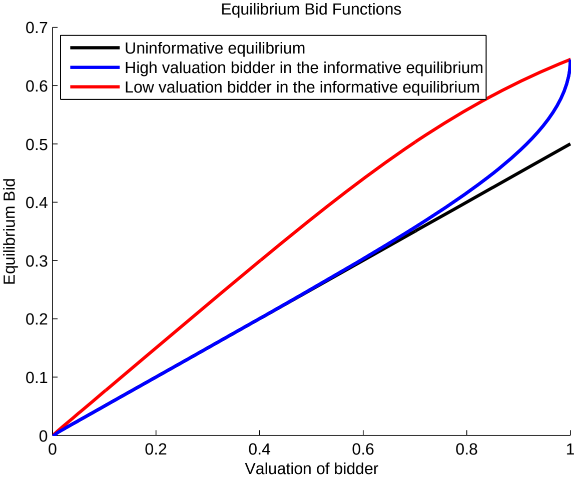

Gain from deviating from the informative equilibrium

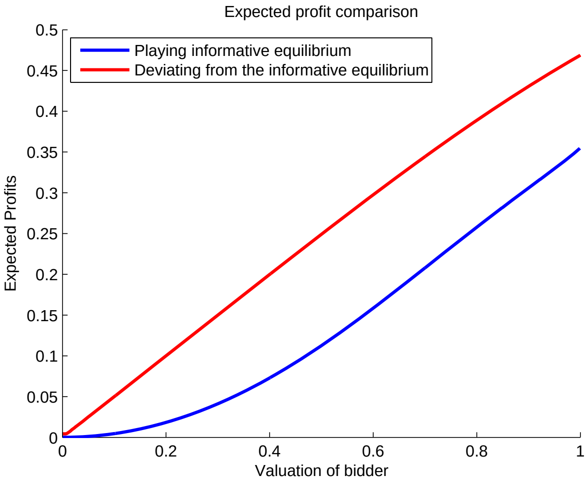

Feasibility of taking game theory "seriously"

- A preponderance of experimental evidence suggest that bidders in auction do not behave like rational, risk-neutral expected payoff maximizers who use Perfect Bayesian Equilibrium bidding strategies.
- For example even in a Second Price sealed bid auction where bidding truthfully is a dominant strategy, we frequently observe *overbidding* (i.e. bidding more than one's valuation for the item) by experimental subjects.
- Given the practical consideration that the computational complexity in solving for even one PBE, it may be wise to consider other computationally simpler and potentially more empirically plausible approaches to modeling bidding in auctions.

Behavioral DP model of bidding with fixed beliefs

- We discretize the two minute auction into $T = 120$ one second time steps.
- Let $\tau = (v, c)$ denote the *type* of the bidder, where v is the bidder's valuation of the car being auctioned, and c is the bidder's psychic *cost of submitting a bid.*
- We assume that bidders are experienced and have *fixed, rational beliefs* about the stochastic process for the high bid in auctions for homogeneous types of cars.
- Bidder beliefs are captured by a family of conditional probability distributions $\{\lambda_t(b|b_t, h_t)\}$ where $\lambda_t(b|b_t, h_t)$ is a CDF for the high bid at during the interval $(t, t+1]$ of the auction, conditioned on b_t, the bidder's highest bid up to time t (or 0 if the bidder has not bid yet), and h_t is an indicator of whether the bidder holds the high at t.

Justification for fixed beliefs

- We focus on auctions of a homogeneous class of rental cars (Hyundai Avante Elanta XD with 1.6L engines) which are unlikely to have unique characteristics that make individual auctions to be "unique" (as opposed to an auction for a Picasso or Rembrandt).
- As a result it is plausible that experienced bidders will assume there is a common stochastic process describing the evolution of the high bid in these auction, and we assume all bidders know this stochastic process.
- Thus, participating in additional auctions is unlikely to change the bidder's beliefs about this stochastic process — learning has "converged" to a rational expectation of this stochastic process.
- What a bidder does learn during an individual auction is whether he/she holds the high bid based their history of their own bids during the auction.
- Thus early bidding can be regarded as means of learning what the high bid is in order to avoid overpaying to win the auction.

2005 Hyundai Avante Elantra XD 1.6L

Unconditional beliefs about high bid in each second of auction

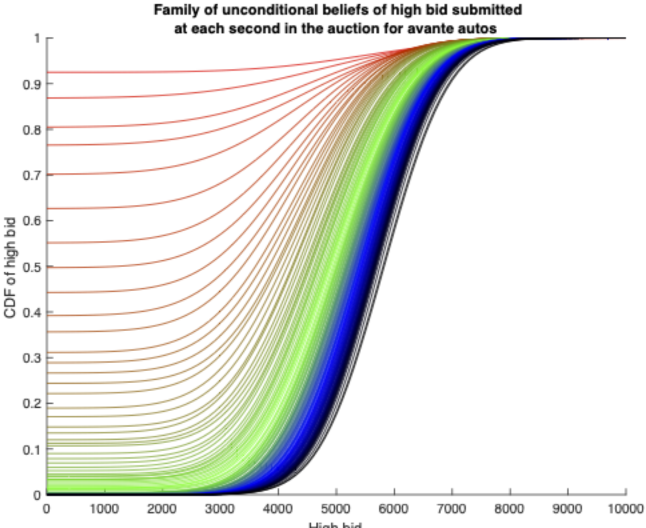

Evolution of beliefs of B23 in auction 9810

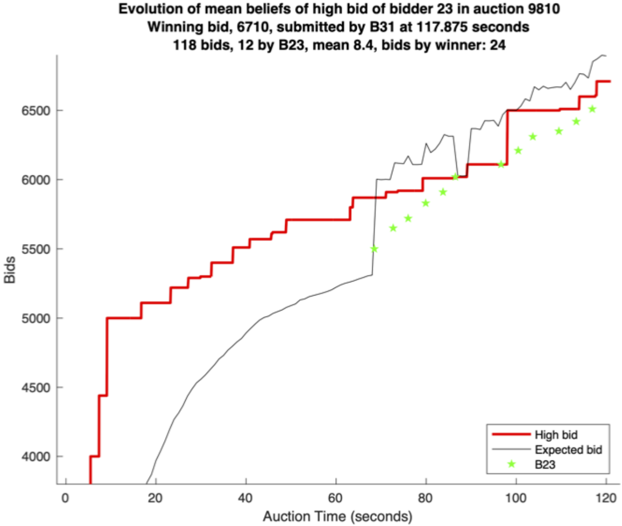

Evolution of beliefs of B23 in auction 10007

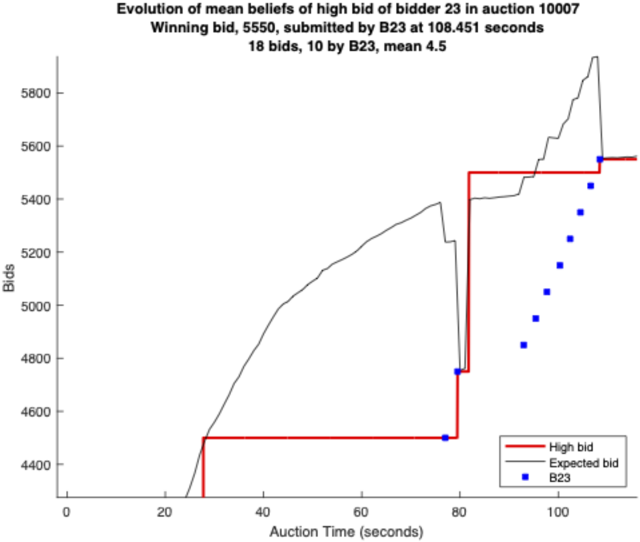

Two Step Estimation Strategy

- **Step 1** Using data on 533 auctions of Avante cars, we estimate the family of beliefs $\{\lambda_t\}$ about the stochastic process for the high bids in the auction.
- **Step 2** Using $\{\hat{\lambda}_t\}$ we solve a discrete dynamic programming problem determining the optimal bidding of the bidder at each second of the auction, resulting in a family of bid functions $\{\beta_t\}$ where $\beta_t(b_t, h_t)$ is the optimal bid at the start of second t in the auction when the bidder's high bid so far is b_t and h_t is an indicator of whether the bidder has the high bid or not.
- Note the only unknown parameters in step 2 are $\tau = (v, c, p, \sigma)$. Thus, we form a likelihood $L(\tau)$ for *each bidder in each auction* resulting in auction-by-auction bidder-specific estimates of valuations, v and other parameters τ characterizing the bidder's type.
- Our goal is to see if such a model is capable of explaining the early bidding we observe in these auctions.

Bidder's DP problem

- Let $h_t = 1$ if the bidder has the highest bid up to second t in the auction, $h_t = 0$ otherwise. Let b_t be the highest bid submitted by the bidder up to second t in the auction.
- The timing is as follows. At the start of each "bidding instant" t ($t = 0, 1, \ldots, 120$), the bidder observes (b_t, h_t) and decides whether to submit a bid $b > b_t$ or not bid, which is equivalent to a non-improved bid of $b = b_t$. At $t = 0$, the auction is initialized with $b_0 = 0$ and $h_0 = 0$ for all bidders.
- The transition rule for bids by a given bidder is as follows: $b_{t+1} = b$, where b is the bid decision at time t. Thus, if the decision is not to bid, then $b_{t+1} = b_t$, otherwise if the bidder submits a bid of $b > b_t$, then $b_{t+1} = b$. Also $h_{t+1} = 1$ if b_{t+1} is the highest bid outstanding at start of second $t + 1$, otherwise $h_{t+1} = 0$.
- We assume that there is a "distraction probability" p that prevents a bidder from focusing on the auction and deciding whether to update their bid at each second t of the auction. Thus, with probability at least p, no bid is submitted and $b_{t+1} = b_t$.

Bidder's DP problem, continued

- The terminal payoff of the bidder at the conclusion of the auction at $T+1 = 121$ (after the final bids have been submitted so the high bid can be determined) is

$$W_{T+1}(b_{T+1}, h_{T+1}) = (v - b_{T+1}) I\{h_{T+1} = 1\},$$

  where b_{T+1} is the bid the bidder submitted at the last possible bidding instant $T = 120$ and $h_{T+1} = 1$ if this was the highest bid in the auction, or 0 otherwise.
- Define $\lambda_T(b|b_T, h_T) = E\{I\{h_{T+1} = 1\}|b, b_T, h_T\}$, i.e. this is the probability that the bidder will win the auction by placing a bid of b at the last possible instant $T = 120$, conditioning on their information (b_T, h_T) at this instant.

Bidder's DP problem, continued

- Define the *bid-specific value function* $w_T(b, b_T, h_T)$ by

$$w_T(b, b_T, h_T) = E\{W_{T+1}(b_{T+1}, h_{T+1})|b, b_T, h_T\} = (v - b)\lambda_T(b|b_T, h_T).$$

  Thus, $w_T(b, b_T, h_T)$ is the expected payoff to the bidder from placing a final bid of b at the last possible bidding instant T in the auction, assuming the bidder is not distracted and thus able to bid.

- Define the value function $W_T(b_T, h_T, \epsilon_T)$ by

$$W_T(b_T, h_T, \epsilon_T) = \max\left[w_T(b_T, b_T, h_T) + \epsilon_T(0), \max_{b \geq b_T}[-c + \epsilon_T(1) + w_T(b, b_T, h_T)]\right]$$

  where $\epsilon_T = (\epsilon_T(0), \epsilon_T(1))$ is a bivariate Type-1 extreme value distribution that reflects idiosyncratic "noise" affecting the bidder's calculation of an optimal bid. Parameter c is the cost of "mental effort" to calculate an improved bid. We assume that passing on bidding involves zero additional mental effort.

Bidder's DP problem, continued

- If the bidder is not distracted from bidding at T their expected value is $EW_T(b_T, h_T)$, given by

$$\begin{aligned} & EW_T(b_T, h_T) \\ & = \int_{\epsilon_T} W_T(b_T, h_T, \epsilon_T) q(\epsilon_T) \\ & = \sigma \log \left(\exp\{w_T(b_T, b_T, h_T)/\sigma\} + \exp\{\max_{b \geq b_T} [w_T(b, b_T, h_T) - c]/\sigma\} \right) \end{aligned}$$

- However if the bidder is distracted at T and does not bid, their value is $w_T(b_T, b_T, h_T)$.
- Then at time $T-1$, the bid-specific value function is $w_{T-1}(b, b_{T-1}, h_{T-1})$ is

$$\begin{aligned} & w_{T-1}(b, b_{T-1}, h_{T-1}) = \\ & [p w_T(b, b, 1) + (1-p) EW_T(b, 1)] \, \lambda_{T-1}(b | b_{T-1}, h_{T-1}) + \\ & [p w_T(b, b, 0) + (1-p) EW_T(b, 0)] \, [1 - \lambda_{T-1}(b | b_{T-1}, h_{T-1})]. \end{aligned}$$

Bidder's DP problem, continued

- Continuing the backward induction from $t = T, T-1, \ldots, 0$ we have solved for the optimal dynamic bidding strategy in the auction.
- The formulas for the expected value of bidding for bidders who are not distracted the same as given above, so we recursively calculate $EW_t(b_t, h_t)$, and the value of being distracted is $w_t(b_b, b_t, h_t)$, recursively for $t = T-1, T-2, \ldots, 1, 0$.
- Then at each time t the bid-specific value function is $w_t(b, b_t, h_t)$ given by

$$
\begin{aligned}
& w_t(b, b_t, h_t) = \\
& [pw_{t+1}(b, b, 1) + (1-p)EW_{t+1}(b, 1)]\, \lambda_t(b|b_t, h_t) + \\
& [pw_{t+1}(b, b, 0) + (1-p)EW_{t+1}(b, 0)]\, [1 - \lambda_t(b|b_t, h_t)].
\end{aligned}
$$

Maximum Likelihood Estimation

- We are able to estimate the parameters $\tau = (v, c, p, \sigma)$ for each bidder in each auction they participate in by maximum likelihood.
- For a given auction, we observe $\{(b_t, h_t), t = 0, \ldots, 120\}$ where b_0 is the first bid made at bidding instant $t = 0$ and b_{120} is the final bid made at $T = 120$. Let the initial conditions be $b_{-1} = h_{-1} = 0$.
- Let $L(\tau)$ be the likelihood of bids by a given bidder in a given auction

$$L(\tau) = \prod_{t=0}^{120} P_t(b_t | b_{t-1}, h_{t-1}, \tau),$$

where the probability $P_t(b' | b, h, \tau)$ is given by the MNL formula

$$P_t(b' | b, h, \tau) = \frac{\exp\{-cI\{b' \geq b\} + w_t(b', b, h)/\sigma\}}{\exp\{w_t(b, b, h)/\sigma\} + \sum_{b' \geq b} \exp\{-c + w_t(b', b, h)/\sigma\}}.$$

Maximum Likelihood Estimation

- The maximum likelihood estimator presumes that for each t and integer bid $b' > b$ there is a corresponding extreme value distributed idiosyncratic shock $\epsilon_t(b')$ associated with choosing b'. The bid-specific value function for this version of the model is

$$W_t(b_t, h_t, \epsilon_t) =$$
$$\max\left[w_t(b_t, b_t, h_t) + \epsilon_t(0), \max_{b' \geq b_t}[-c + \epsilon_t(b') + w_t(b', b_t, h_t)]\right]$$

  and the expected value is

$$EW_t(b_t, h_t) = \int_{\epsilon_t} W_t(b_t, h_t, \epsilon_t) q(\epsilon_t)$$
$$= \sigma \log\left(\exp\{w_t(b_t, b_t, h_t)/\sigma\} + \sum_{b' \geq b} \exp\{[w_t(b', b_t, h_t) - c]/\sigma\}\right).$$

- This model predicts a positive probability for any integer bid $b' \geq b$ given by the logit probability above.

Quasi Maximum Likelihood Estimation

- However evaluation of the sum of exponentiated bid-specific value functions for all integer bids $b' \geq b_t$ is computationally expensive. So we propose an alternative *quasi-maximum likelihood estimator* based on an incomplete model of bidding that does not a formal theory (i.e. positive probability of) any potential bid $b' \geq b_t$.
- Under this alternative model, there are only two idiosyncratic shocks $(\epsilon_t(0), \epsilon_t(1))$ per bidding instant and the value function is

$$W_t(b_t, h_t, \epsilon_t) =$$
$$\max \left[w_t(b_t, b_t, h_t) + \epsilon_t(0), \max_{b' \geq b_t} [w_t(b', b_t, h_t) - c] + \epsilon_t(1) \right]$$

and expected value is

$$EW_t(b_t, h_t)$$
$$= \int_{\epsilon_t} W_t(b_t, h_t, \epsilon_t) q(\epsilon_t)$$
$$= \sigma \log \left(\exp\{w_t(b_t, b_t, h_t)/\sigma\} + \exp\{\max_{b' \geq b_t} [w_t(b', b_t, h_t) - c]/\sigma\} \right).$$

Quasi Maximum Likelihood Estimation

- Suppose we observe a bid of b_{t+1} at bidding instant t in bidding state (b_t, h_t).
- If $b_{t+1} > b_t$ (i.e. the bidder improved their bid), the model with only two idiosyncratic shocks per bidding instant cannot formally "explain" this bid, i.e. there is zero probability of observing "suboptimal bids" $b_{t+1} \neq \beta_t(b_t, h_t)$.
- But the QMLE assigns the following probability to a bid $b_{t+1} > b_t$

$$\Pi_t(b_{t+1}|b_t, h_t) =$$
$$\frac{\exp\{w_t(b_{t+1}, b_t, h_t)/\gamma\}}{\exp\{w_t(b_{t+1}, b_t, h_t)/\gamma\} + \exp\{w_t(\beta_t(b_t, h_t), b_t, h_t)/\gamma\}},$$

where $\beta_t(b_t, h_t)$ is the *optimal bid function* at instant t given by

$$\beta_t(b_t, h_t) = \underset{b' \geq b_t}{\operatorname{argmax}}\, w_t(b', b_t, h_t).$$

and $\gamma \geq 0$ is a *smoothing parameter* or *penalty parameter* for observations $b_{t+1} \neq \beta_t(b_t, h_t)$.

Aside on the Optimal Bid Function

- In our "partial" model of bidding in the Korean auction, the optimal bid function is actually also a function of the unobserved shocks $\epsilon_t = (\epsilon_t(0), \epsilon_t(1))$. We denote this bid function by $\beta_t(b_t, h_t, \epsilon_t)$ and it is given by

$$\beta_t(b_t, h_t, \epsilon_t) =$$

$$\mathrm{argmax}\left[w_t(b_t, b_t, h_t) + \epsilon_t(0), \underset{b' \geq b_t}{\mathrm{argmax}}[w_t(b', b_t, h_t) - c] + \epsilon_t(1)\right].$$

- The relationship between $\beta_t(b, h, \epsilon)$ and $\beta_t(b, h)$ is as follows

$$\beta_t(b, h, \epsilon) = \begin{cases} \beta_t(b, h) & \text{if } w_t(b, b, h) + \epsilon(0) \leq w_t(\beta_t(b, h), b, h) + \epsilon(1) \\ b & \text{if } w_t(b, b, h) + \epsilon(0) > w_t(\beta_t(b, h), b, h) + \epsilon(1) \end{cases}$$

- Thus the optimal bid is $\beta_t(b, h)$ for any combination of private bidding shocks that makes it optimal for the bidder to improve their existing bid b, otherwise it is optimal not to improve the current bid (i.e. not bid).

Quasi Maximum Likelihood Estimation

- Thus, under the QMLE, $\Pi_t(b_{t+1}|b_t, h_t)$ is maximized at the value $\Pi_t(b_{t+1}|b_t, h_t) = 1/2$ when $b_{t+1} = \beta_t(b_t, h_t)$.
- To maximize the QMLE, parameters τ are found that make the optimal bidding function $\beta_t(b_t, h_t, \tau)$ to be as close as possible to the observed bid b_{t+1} since this maximizes $\Pi_t(b_{t+1}|b_t, h_t)$.
- The QMLE then is defined by

$$\hat{\tau} = \operatorname*{argmax}_{\tau} QL(\tau) \equiv \operatorname*{argmax}_{\tau} \prod_{t=0}^{120} P_t(b_t|b_{t-1}, h_{t-1}, \tau),$$

where $P_t(b_t|b_{t-1}, h_{t-1}, \tau)$ is given by

$$P_t(b'|b, h) = \begin{cases} 1 - \pi_t(b|b, h, \tau) & \text{if } b' = b \\ \pi_t(b'|b, h, \tau)\Pi_t(b'|b, h, \tau) & \text{if } b' \geq b \end{cases} \tag{1}$$

Quasi Maximum Likelihood Estimation

- Where $\pi_t(b'|b,h,\tau)$ is the probability of bidding given by

$$\pi_t(b'|b,h,\tau) = \frac{\exp\{[w_t(\beta_t(b,h),b,h,\tau)-c]/\sigma\}}{\exp\{w_t(b,b,h,\tau)/\sigma\}+\exp\{[w_t(\beta_t(b,h),b,h,\tau)-c]/\sigma\}},$$

  and $\Pi_t(b'|b,h,\tau)$ is the probability of observing a potentially suboptimal bid b'

$$\Pi_t(b'|b,h,\tau) = \frac{\exp\{w_t(b',b,h,\tau)/\gamma\}}{\exp\{w_t(b',b,h,\tau)/\gamma\}+\exp\{w_t(\beta_t(b,h),b,h,\tau)/\gamma\}}.$$

- Thus, the QMLE $\hat{\tau}$ is a value of τ that maximizes the probability of the observed sequence of bids in the auction by a given bidder, even when we have an *incomplete model of bidding* — i.e. our behavioral model does not assign a positive probability to every possible bid b'.

MLE and QMLE Gradients

- Except for the parameter σ the other parameters in $\tau = (v, c, p, \sigma)$ enter the MLE and QMLE only via the bid-specific value functions $w_t(\beta_t(b, h, \tau), b, h, \tau)$
- So now we consider recursive calculation of gradients $\nabla_\tau w_t(\beta_t(b, h, \tau), b, h, \tau)$
- Note via Envelope Theorem we only have consider the direct effect of changes in τ on $w_t(b', b, h, \tau)$ at $b' = \beta_t(b, h, \tau)$ and not the indirect effect of τ on $\beta_t(b, h, \tau)$ since we have
- Case A: if $\beta_t(b, h, \tau) > b$ (interior optimum) the first order condition for an optimal bid holds

$$\frac{\partial}{\partial b'} w_t(b', b, h, \tau) = 0 \quad \text{at } b' = \beta_t(b, h, \tau),$$

- Case B: if $\beta_t(b, h, \tau) = b$ (boundary solution) we see directly that in this case we have

$$\nabla_\tau \beta_t(b, h, \tau) = \nabla_\tau b = 0.$$

Recursion for $\nabla_v w_T(b', b, h, \tau)$

- Consider first the recursion formula for the gradients of the bid-specific values with respect to v.
- In the final period, $T = 120$ we have

$$w_T(\beta_T(b,h,\tau), b, h, \tau) = [v - \beta_T(b,h,\tau)]\lambda_T(\beta_T(b,h,\tau)|b,h).$$

- Thus, by the Envelope Theorem we have

$$\begin{aligned} \nabla_v w_T(\beta_T(b,h,\tau), b, h, \tau) &= \lambda_T(\beta_T(b,h,\tau)|b,h), \\ \nabla_v w_T(b', b, h, \tau) &= \lambda_T(b'|b,h). \end{aligned}$$

$$\begin{aligned} \nabla_v EW_T(b,h,\tau) = & [1 - P_T(\beta_T(b,h,\tau)|b,h,\tau)]\lambda_T(b|b,h) + \\ & P_T(\beta_T(b,h,\tau)|b,h,\tau)\lambda_T(\beta_T(b,h,\tau)|b,h). \end{aligned}$$

- Thus $\nabla_v EW_T(b,h,\tau)$ is the *ex ante expected win probability* at time T, i.e. before the bidder realizes the values of the two bidding shocks $(\epsilon_T(0), \epsilon_T(1))$.

Recursions for $\nabla_v w_{T-1}(b', b, h, \tau)$ and $\nabla_v w_t(b', b, h, \tau)$

- Now go back to bidding instant $T-1$ and take gradient with respect to v using the recursion for $w_{T-1}(b', b, h, \tau)$ to get

$$\nabla_v w_{T-1}(b', b, h) =$$
$$[p\nabla_v w_T(b', b', 1) + (1-p)\nabla_v EW_T(b', 1)]\, \lambda_{T-1}(b'|b, h) +$$
$$[p\nabla_v w_T(b', b', 0) + (1-p)\nabla_v EW_T(b', 0)]\, [1 - \lambda_{T-1}(b'|b, h)].$$

- Substituting earlier formulas, we can that $\nabla_v w_{T-1}(b', b, h)$ is the expected probability of winning the auction at instant $T-1$ when a bid of b' is submitted.
- Continuing the backward induction, at instant t we have

$$\nabla_v w_t(b', b, h) =$$
$$[p\nabla_v w_{t+1}(b', b', 1) + (1-p)\nabla_v EW_{t+1}(b', 1)]\, \lambda_t(b'|b, h) +$$
$$[p\nabla_v w_{t+1}(b', b', 0) + (1-p)\nabla_v EW_{t+1}(b', 0)]\, [1 - \lambda_t(b'|b, h)].$$

- Thus, at all bidding instants t, $\nabla_v w_t(b', b, h, \tau)$ equals the expected probability of winning the auction as of time t, given the information available (b, h) and conditional on submitting a bid equal to $b' \geq b$.

Martingale property for $\{\nabla_v w_t(\beta_t(b_t, h_t), b_t, h_t, \tau)\}$

- **Theorem** *Consider the stochastic process* $\{X_t\}$ *defined by*

  $$X_t \equiv \nabla_v w_t(\beta_t(b_t, h_t, \epsilon_t), b_t, h_t, \tau)$$

  i.e. $\{X_t\}$ *is the the gradient of the value function evaluated along the optimal bidding strategy for the controlled stochastic process defining optimal bidding in the Korean auction. Then* $\{X_t\}$ *is a martingale, i.e. with probability 1 we have*

  $$X_t = E\{\{h_T = 1\}|\mathcal{I}_t\} = E\{\tilde{X}_{t+1}|\mathcal{I}_t\},$$

  where $\mathcal{I}_t$ *is the information available at instant t in the auction, which given the Markovian nature of the process is* $\mathcal{I}_t = (b_t, h_t, \epsilon_t)$.
- **Corollary** $\nabla_v w_t(b', b, h) \in [0, 1]$ *for all* $b' \geq b$, $b \geq 0$ *and* $h \in \{0, 1\}$ *and* $t \in \{0, 1, \ldots, 120\}$.
- The proof of the Corollary follows easily from the definition of $\nabla_v w_t(b', b, h)$, which we have shown is the conditional expectation of the probability of winning the auction, and hence must be in the unit interval for all values of its arguments.

Recursion for $\nabla_c w_T(b', b, h, \tau)$

- Next, consider the recursion formula for the gradients of the bid-specific values with respect to c.
- In the final period, $T = 120$ we have

$$w_T(\beta_T(b, h, \tau), b, h, \tau) = [v - \beta_T(b, h, \tau)]\lambda_T(\beta_T(b, h, \tau)|b, h).$$

- So clearly, we have

$$\nabla_c w_T(b', b, h, \tau) = 0.$$

$$\nabla_c EW_T(b, h, \tau) = -P_T(\beta_T(b, h, \tau)|b, h, \tau).$$

- $\nabla_c EW_T(b, h, \tau)$ is the negative of the *ex ante probability of improving the bid* at time T, i.e. before the bidder realizes the values of the two bidding shocks $(\epsilon_T(0), \epsilon_T(1))$.

Recursion for $\nabla_c w_{T-1}(b', b, h, \tau)$

- Now go to bidding instant $T-1$. We have

$$\nabla_c w_{T-1}(b', b, h) = -(1-p)P_T(\beta_T(b',1)|b',1)\lambda_{T-1}(b'|b,h)$$
$$-(1-p)P_T(\beta_T(b',0)|b',0)[1-\lambda_{T-1}(b'|b,h)].$$

- This equals the *ex ante* reduction (as of instant $T-1$) in the probability that the bidder will place a bid at the last bidding instant T. Now we calculate

$$\nabla_c EW_{T-1}(b,h) = -P_{T-1}(\beta_{T-1}(b,h)|b,h)+$$
$$P_{T-1}(\beta_{T-1}(b,h)|b,h)\nabla_c w_{T-1}(\beta_{T-1}(b,h),b,h)+$$
$$[1-P_{T-1}(\beta_{T-1}(b,h)|b,h)]\nabla_c w_{T-1}(b,b,h).$$

- This gradient equals the direct effect of an increase in c on bidding at instant $T-1$, $-P_{T-1}(\beta_{T-1}(b,h)|b,h)$, plus the *ex ante* expected effect of reduced likelihood of bidding at the final instant T.

Recursion for $\nabla_c w_t(b', b, h, \tau)$

- At generic bidding instant $t < T - 1$ we have

$$\nabla_c w_t(b', b, h) = \left[p\nabla_c w_{t+1}(b', b', 1) + (1-p)\nabla EW_{t+1}(b', 1)\right]\lambda_t(b'|b, h)$$
$$+ \left[p\nabla_c w_{t+1}(b', b', 0) + (1-p)\nabla_c EW_{t+1}(b', 0)\right]\left[1 - \lambda_t(b'|b, h)\right].$$

- This equals the *ex ante* reduction (as of instant $T - 1$) in the payoff to the bidder from an increase in c, which is the negative of the recursively calculated probability that the bidder will place a bid at at any instant between t and T. Now we calculate

$$\nabla_c EW_t(b, h) = -P_t(\beta_t(b, h)|b, h)+$$
$$P_t(\beta_t(b, h)|b, h)\nabla_c w_t(\beta_t(b, h), b, h)+$$
$$[1 - P_t(\beta_t(b, h)|b, h)]\nabla_c w_t(b, b, h).$$

- Thus, the gradients cumulate as we move backward to earlier bidding instants in the auction, since the effect of an increase in bidding costs c affects both the current and future propensity to bid in the auction.

Recursion for $\nabla_p w_T(b', b, h, \tau)$

- Next, consider the recursion formula for the gradients of the bid-specific values with respect to p. It should act similarly to an increase in c in that it is also a "bidding friction" that reduces the propensity to bid in the auction, and hence can be expected to reduce expected winnings.
- In the final period, $T = 120$ we have

$$\begin{aligned} \nabla_p w_T(b', b, h, \tau) &= 0 \\ \nabla_p EW_T(b, h, \tau) &= 0. \end{aligned}$$

- These are zero since at the last bidding instant, there are no future bidding instants left where the bidder can be "distracted" from bidding. If the bidder was distracted at T then their value would be just $w_T(b, b, h)$, i.e. their previous bid b would also be their final bid. If the bidder was not distracted, then they would have the potential to adjust their final bid and submit a bid $b' > b$.

Recursion for $\nabla_p w_{T-1}(b', b, h, \tau)$

- But at $T - 1$, an increase in p does affect future payoffs in the auction. We have

$$\nabla_p w_{T-1}(b', b, h) =$$
$$[w_T(b', b', 1) - EW_T(b', 1)]\, \lambda_{T-1}(b'|b, h) +$$
$$[w_T(b', b', 0) - EW_T(b', 0)]\, [1 - \lambda_{T-1}(b'|b, h)].$$

- And this implies that

$$\nabla_p EW_{T-1}(b, h) = P_{T-1}(\beta_{T-1}(b, h)|b, h) \nabla_p w_{T-1}(\beta_{T-1}(b, h), b, h) +$$
$$[1 - P_{T-1}(\beta_{T-1}(b, h)|b, h)] \nabla_p w_{T-1}(b, b, h).$$

- It is easy to see that $\nabla_p w_{T-1}(b', b, h) \leq 0$ since $w_T(b', b, h) \leq EW_T(b', h)$ for any bid $b' \geq b$ and all $h \in \{0, 1\}$.

Recursion for $\nabla_p w_t(b', b, h, \tau)$

- At generic bidding instant $t < T$ the future bid-specific value functions $w_t(b', b, h)$ depend on p so we have

$$
\begin{aligned}
&\nabla_p w_t(b', b, h) = \\
&[w_{t+1}(b', b', 1) - EW_{t+1}(b', 1)]\,\lambda_t(b|b, h) + \\
&[p\nabla_p w_{t+1}(b', b', 1) + (1-p)\nabla_p EW_{t+1}(b', 1)]\,\lambda_t(b|b, h) + \\
&[w_{t+1}(b', b', 0) - EW_{t+1}(b', 0)]\,[1 - \lambda_t(b'|b, h)] + \\
&[p\nabla_p w_{t+1}(b', b', 0) + (1-p)\nabla_p EW_{t+1}(b', 0)]\,[1 - \lambda_t(b'|b, h)].
\end{aligned}
$$

- And we have

$$
\begin{aligned}
&\nabla_p EW_t(b, h) = P_t(\beta_t(b, h)|b, h)\nabla_p w_t(\beta_t(b, h), b, h) + \\
&[1 - P_t(\beta_t(b, h)|b, h)]\nabla_p w_t(b, b, h).
\end{aligned}
$$

- By induction, it is easy to see that $\nabla_p w_t(b', b, h) \leq 0$ for all t, $b' \geq b$ and all $h \in \{0, 1\}$.

Recursion for $\nabla_\sigma w_T(b', b, h, \tau)$

- Finally, consider the recursion formula for the gradients of the bid-specific values with respect to σ. Similar to the case of parameters c and p we have

$$\nabla_\sigma w_T(b', b, h, \tau) = 0.$$

- However we have

$$\begin{aligned}\nabla_\sigma EW_T(b, h) = &\frac{1}{\sigma} EW_T(b, h) \\ &- \frac{1}{\sigma} P_T(\beta_T(b, h)|b, h)[w_T(\beta_T(b, h), b, h) - c] \\ &- \frac{1}{\sigma}[1 - P_T(\beta_T(b, h)|b, h)] w_T(b, b, h).\end{aligned}$$

- It is not hard to show that $\nabla_\sigma EW_T(b, h) \geq 0$ using the definition of $EW_T(b, h)$ and the formula above.

Recursion for $\nabla_\sigma w_{T-1}(b', b, h, \tau)$

- We have

$$\nabla_\sigma w_{T-1}(b', b, h) = (1-p)\nabla_\sigma EW_T(b', 1)\lambda_{T-1}(b'|b, h)+$$
$$(1-p)\nabla_\sigma EW_T(b', 0)[1-\lambda_{T-1}(b'|b, h)].$$

$$\begin{aligned}
\nabla_\sigma EW_{T-1}(b, h) &= \frac{1}{\sigma} EW_{T-1}(b, h) \\
&- \frac{1}{\sigma} P_{T-1}(\beta_{T-1}(b, h)|b, h)[w_{T-1}(\beta_{T-1}(b, h), b, h) - c] \\
&- \frac{1}{\sigma}[1 - P_{T-1}(\beta_{T-1}(b, h)|b, h)] w_{T-1}(b, b, h) \\
&+ P_{T-1}(\beta_{T-1}(b, h)|b, h)\nabla_\sigma w_{T-1}(\beta_{T-1}(b, h), b, h) \\
&+ [1 - P_{T-1}(\beta_{T-1}(b, h)|b, h)]\nabla_\sigma w_{T-1}(b, b, h).
\end{aligned}$$

- It is not hard to show that $\nabla_\sigma w_{T-1}(b', b, h) \geq 0$ and $\nabla_\sigma EW_{T-1}(b, h) \geq 0$ using the definition of $EW_{T-1}(b, h)$ and the formulas above.

Recursion for $\nabla_\sigma w_t(b', b, h, \tau)$

- At generic bidding instant $t < T - 1$ we have

$$\nabla_\sigma w_t(b', b, h) =$$
$$[p\nabla_\sigma w_{t+1}(b', b', 1) + (1-p)\nabla_\sigma EW_{t+1}(b', 1)]\,\lambda_t(b'|b, h) +$$
$$[p\nabla_\sigma w_{t+1}(b', b', 0) + (1-p)\nabla_\sigma EW_{t+1}(b', 0)]\,[1 - \lambda_t(b'|b, h)].$$

$$\begin{aligned}\nabla_\sigma EW_t(b, h) &= \frac{1}{\sigma} EW_t(b, h) \\ &- \frac{1}{\sigma} P_t(\beta_t(b, h)|b, h)[w_t(\beta_t(b, h), b, h) - c] \\ &- \frac{1}{\sigma}[1 - P_t(\beta_t(b, h)|b, h)]w_t(b, b, h) \\ &+ P_t(\beta_t(b, h)|b, h)\nabla_\sigma w_t(\beta_t(b, h), b, h) \\ &+ [1 - P_t(\beta_t(b, h)|b, h)]\nabla_\sigma w_t(b, b, h).\end{aligned}$$

- It is not hard to show that $\nabla_\sigma w_t(b', b, h) \geq 0$ and $\nabla_\sigma EW_t(b, h) \geq 0$ using the definition of $EW_t(b, h)$ and the formulas above.

Simulated Auction 15

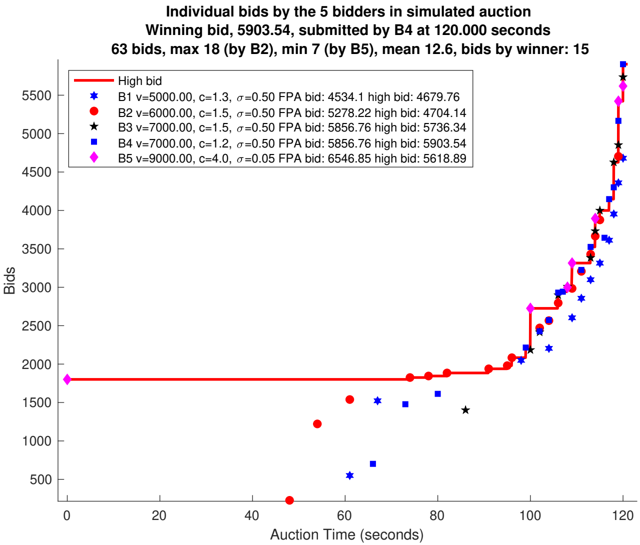

Expected Win Probability Martingales Simulated Auction 15

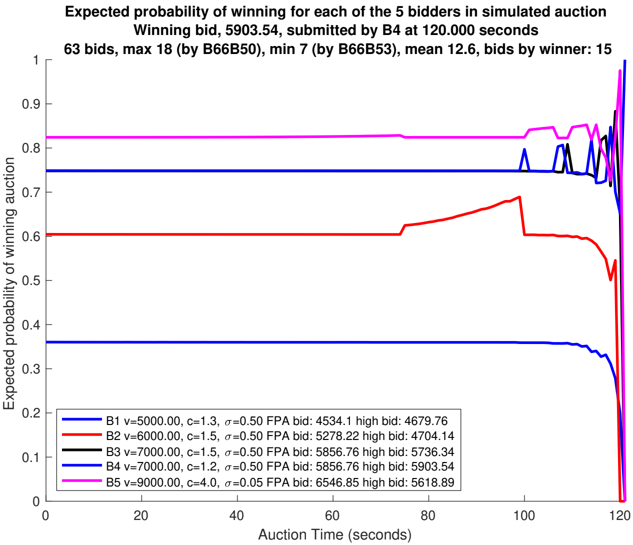

Simulated Auction, $c = 0$

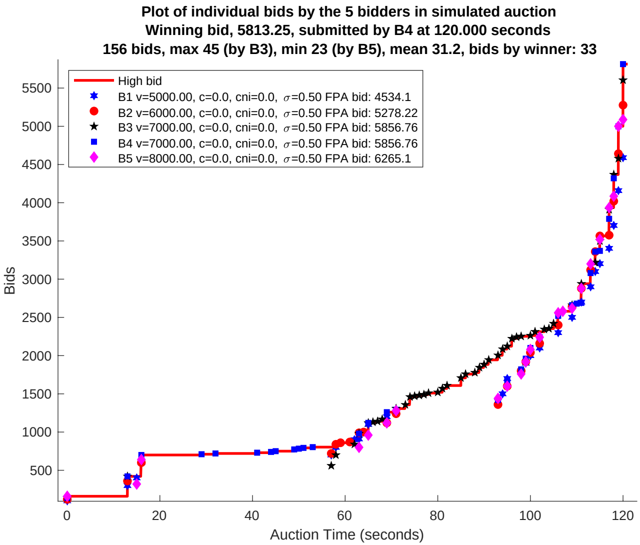

Simulated Auction, $c = 5$

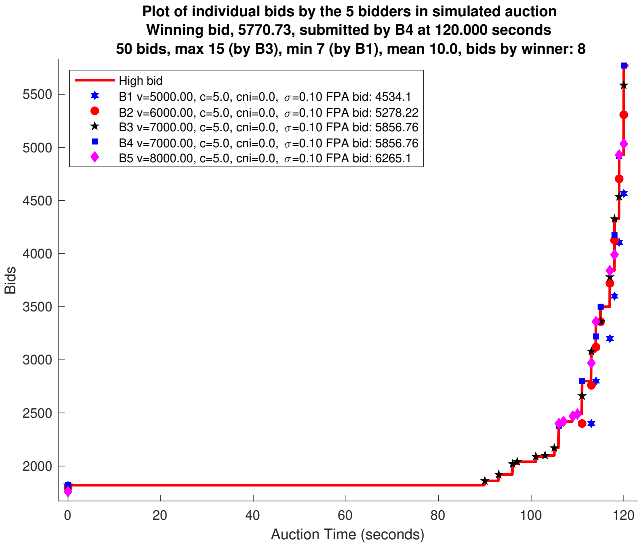

Simulation of Auction 9925

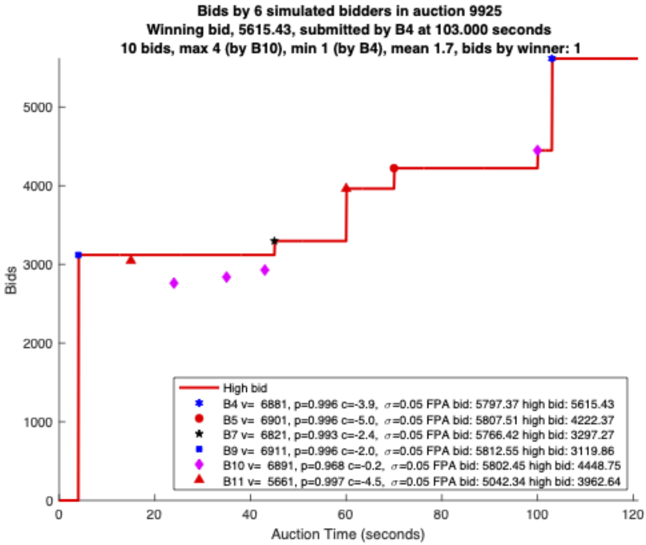

Estimation results: bidder valuations v

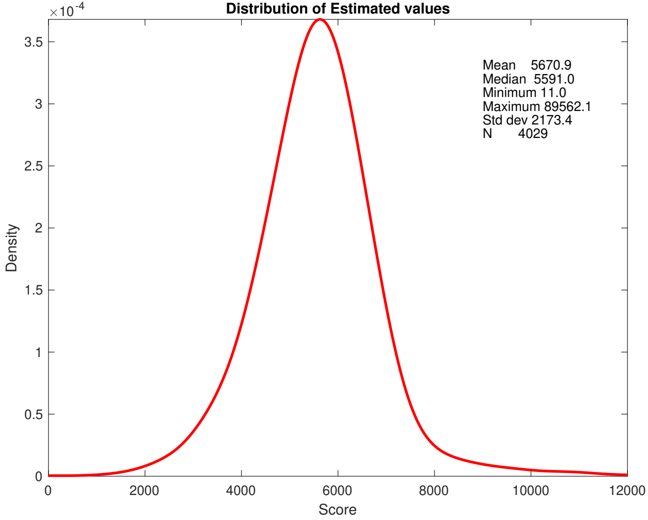

Estimation results: cost of bidding c

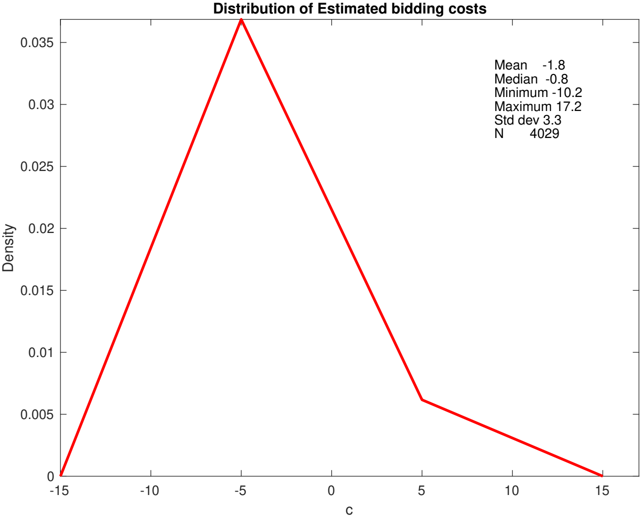

Estimation results: inattention probability p

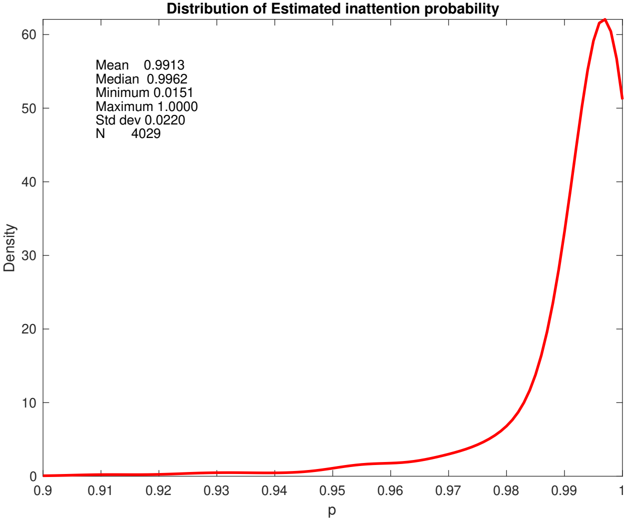

Estimation results: extreme value scale parameter σ

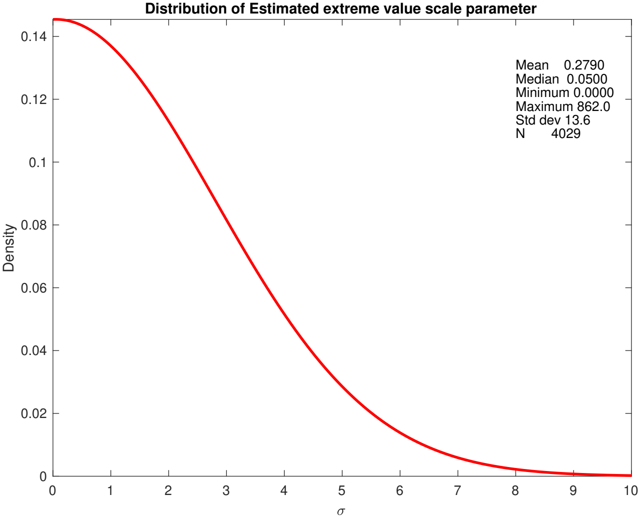

Actual vs predicted bids: bidder 23 in auction 10086

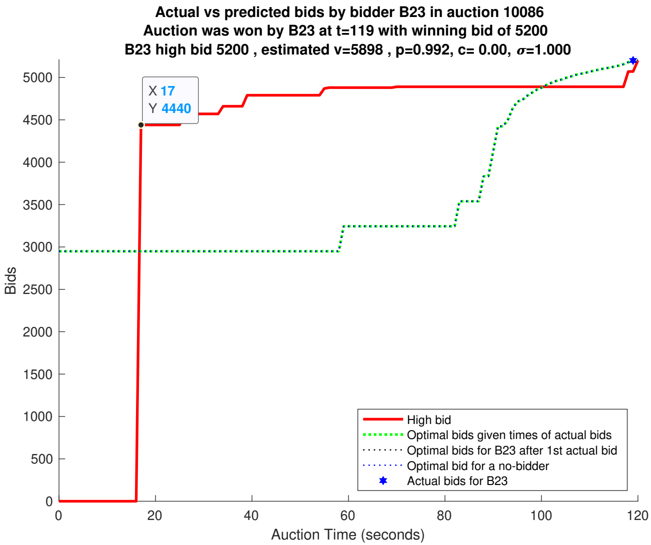

Actual vs predicted bids: bidder 10 in auction 10007

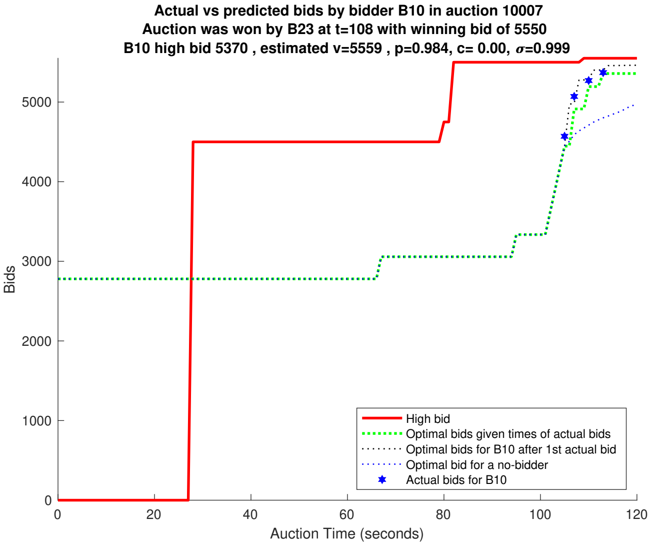

Probability of bidding and winning: B10 auction 10007

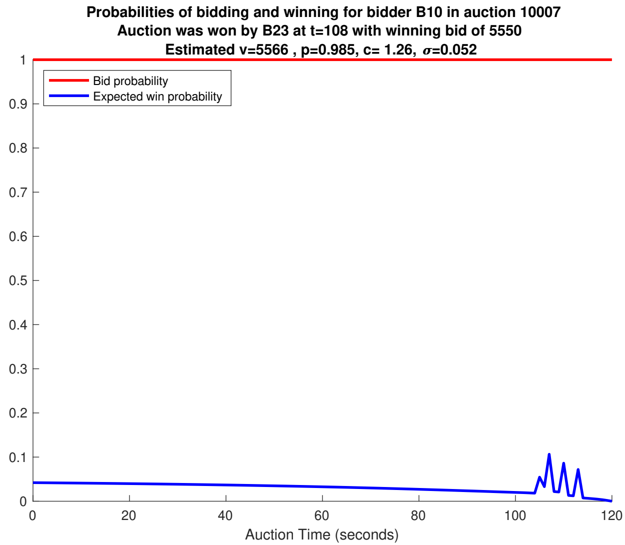

Payoffs from bidding and not bidding: B10 auction 10007

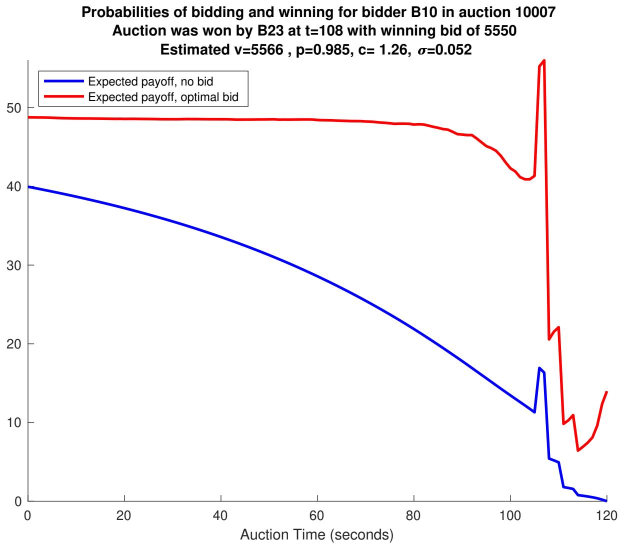

Actual vs predicted bids: bidder 32 in auction 10313

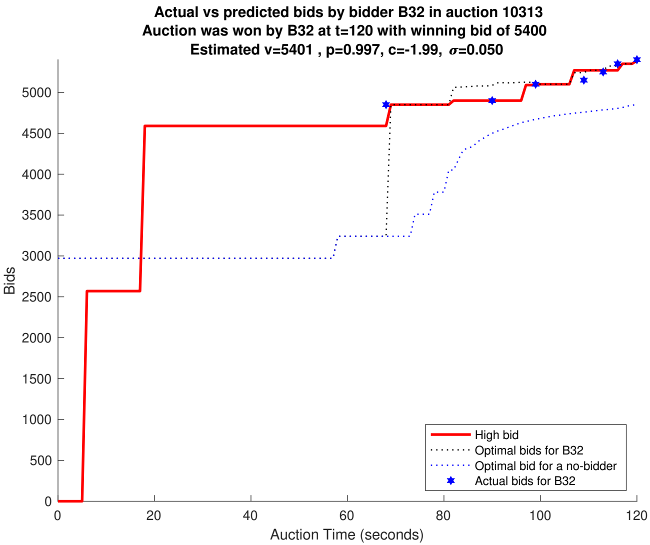

Estimated profits of auction winner

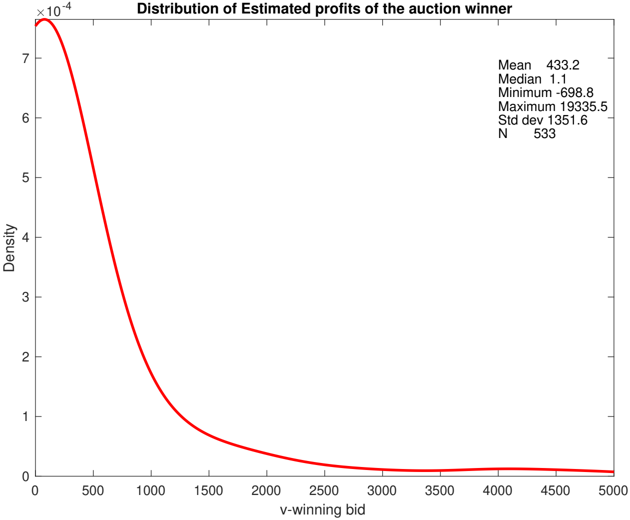

Distribution of winning bid less predicted sniping bid

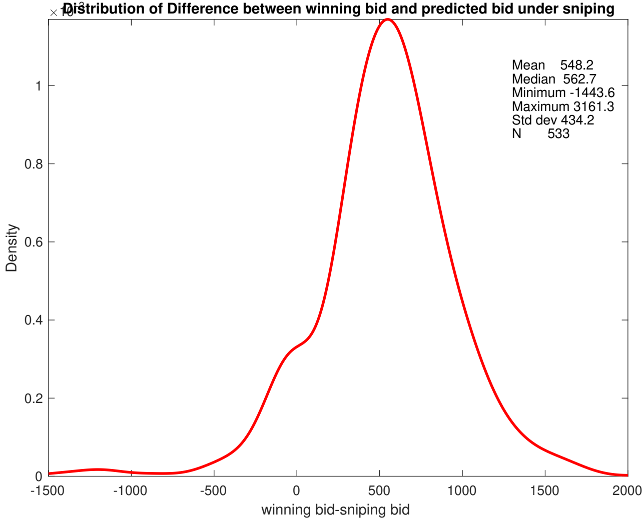

Preliminary Conclusions

- We have analyzed a unique new data set on dynamic informationally restricted auctions invented by a Korean rental car company.
- We have shown that early bidding in these auctions is very prevalent and appears to reflect an attempt by bidders to learn the value of the high bid in the auction in order to win without overpaying.
- However we have suggested that this behavior may be inconsistent with the predictions of a perfect Bayesian equilibrium model of bidding in these auctions.
- In a 2 bidder, 2 period example, we showed there is no informative PBE: the only PBE is an uninformative equilibrium in which both bidders wait to the last period to submit their bids.
- In an *uninformative equilibrium* (which always exists) there is no early bidding and the outcome is the same as the equilibrium in a static first price sealed bid auction.

Preliminary Conclusions

- In order to explain the bidding behavior we observe we developed a behavioral bidding model that relaxes the assumption that bidders use PBE strategies.
- Instead we assume that experienced bidders have rational expectations of the stochatic process governing the high bid in the auction.
- Bidders solve dynamic programs to maximize their expected payoff from the auction, given their beliefs about the stochastic process governing the highest bid during the auction.
- We have solved these dynamic programs and shown that they can produce the early bidding behavior we observe.
- Early bidding enables bidders to learn the value of the high bid and to minimize the amount they ned to pay to win the auction.
- Our initial results indicate that the dynamic Korean auctions result in lower expected revenues than the rental company would earn had they used static first price sealed bid auctions to sell their used cars.