Should Auctions Be Transparent?

Dirk Bergemann and Johannes Hörner
Yale University

October 2010

What Information Gets Disclosed to Bidders?

1. Google's sponsored search keyword auction:... nothing.
2. Ebay: User ID is only known to the seller (so, nothing).
3. London bus auctions: only the winner's bid and identity.
4. U.S. Procurement : all ("Freedom of Information Act").

- To avoid fraud (Ebay).
- Rights of American citizens (procurement).
- To preserve privacy (Christie's).
- To prevent collusion (FCC spectrum auctions).
- Seller's Revenue (Google).

To compare the performance of three disclosure policies (and its impact on information revelation):

1. all bids are unobservable
2. only the winner's bid (and identity) is observable.
3. all bids (and identities) are observable.

The environment is an infinitely-repeated first-price auction:

- Values are private, independent, and perfectly correlated over time.
- Payoff is additive (discounted) across periods.

To abstract from collusion *per se*, focus on Markov strategies.

- privately observable bids

Landsberger et al. (2001): consider static environment when the ranking of the valuations is common knowledge; establish existence and uniqueness of BNE with first price auction

- repeated first price auction

Hörner and Jamison (2008) consider common value environment with one informed bidder in an infinite horizon environment

- information disclosure and repeated auction with two periods:

Février (2003), Tu (2007), Yao (2007), Thomas (2010).

More Precisely

- $n+1$ bidders: $i = 1, \ldots, n+1$.
- Values are private and binary: $u_i \in \{\underline{u}, \bar{u}\}$:

$$\bar{u} > \underline{u} \geq 0.$$

- They are independently distributed: $q = \Pr[u_i = \bar{u}]$.
- The horizon is discrete and infinite: $t = 0, \ldots, \infty$.
- Values are constant over time.
- Common discount factor δ.

- bid of bidder i in period t: $b_{i,t} \in \mathbf{R}_+$.
- To resolve issues of discontinuity, allow bids $\underline{u}_+$.
- In period t, highest bidder wins the object.
- Assignment of unit to agent i in period t: $x_{i,t} \in \{0, 1\}$.
- Reward in the case of a win: $u_i - b_{i,t}$.
- Payoff:

$$\mathbf{E}\left[\sum_{t=0}^{\infty}(1-\delta)\delta^t x_{i,t}(u_i - b_{i,t})\right].$$

- History with unobservable bids:

$$h_{i,t} = \{b_{i,0}, x_{i,0}; \ldots; b_{i,t-1}, x_{i,t-1}\}.$$

- History with observable winner:

$$h_{i,t} = \{b_{i,0}, \max_j b_{j,0}, \{x_{j,0}\}_{j=1}^{n-1}; \ldots; b_{i,t-1}, \max_j b_{j,t-1}, \{x_{j,t-1}\}_{j=1}^{n-1}\}.$$

- History with observable bids:

$$h_{i,t} = \{\{b_{j,0}, x_{j,0}\}_{j=1}^{n-1}; \ldots; \{b_{j,t-1}, x_{j,t-1}\}_{j=1}^{n-1}\}.$$

- (Behavior) strategy $\beta_i = \{\beta_{i,t}\}_{t=0}^{\infty}$ of bidder i:

$$\beta_{i,t} : \{\underline{u}_i, \bar{u}_i\} \times H_{i,t} \rightarrow \triangle \mathbf{R}_+$$

- Repeated games typically admit a plethora of equilibria.
- We shall restrict attention to equilibria in which strategies are measurable with respect to the players' beliefs.
- in this spirit, we assume that if two bidders commonly know that their valuations are high, their bids are $\overline{u}$ thereafter.
- similarly, we assume that bids are always at least $\underline{u}$.
- This is tricky with unobservable bids, as posterior beliefs are no longer common knowledge.
- But if low-value bidders bid $\underline{u}$, and high-value bidders bid strictly more, then if a high-value bidder who always lost wins, two bidders commonly know that their value is high.

Preview of the Results

- When all bids are observable, all equilibria are inefficient for a range of parameters.
- When the winner's bid is observable, some equilibria are inefficient for a range of parameters.
- When all bids are unobservable, the equilibrium is efficient (up to a possibly useless and mild additional refinement).

The Unobservable Case

- in each period, and with unobservable bids, each bidder only learns whether he lost or won the current auction
- the binary outcome of the auction (lose vs win) generates a binary information structure
- the (binary) ranking of past bids is indeed common knowledge among the bidders ...
- ... but the posterior beliefs of bidder is not common knowledge anymore
- ... and construction and verification of equilibrium is conceptually challenging

Is Pooling Possible?

- an equilibrium is pooling if bidders with different valuations use the same bidding strategy
- in contrast, if strategies eventually separate types, valuation is revealed
- in a revealing equilibrium, high valuation bidders may eventually compete against each other...
- ...making a pooling equilibrium look rather desirable for the bidders

- consider a pooling equilibrium at the low valuation $\underline{u}$, but remember bids are not observable...
- ...and hence a loss or win does not lead to a revision of the prior
- in a pooling equilibrium, a deviation slightly above $\underline{u}$ guarantees a present win without any future implications...

Theorem

For all q, n, δ, a pooling Markov sequential equilibrium does not exists with observable bids.

The Separating Equilibrium

Let us consider the case of two bidders only.

Call the bidder who lost (resp. won) at $t = 0$ the *loser* (*winner*).

If the loser ever wins with $b > \underline{u}$, the game is over: bids jump to $\overline{u}$.

The loser can get a positive reward only once.

The loser's trade-off: winning early vs. bidding low.

With monotone strategies (later bids nondecreasing in earlier ones), the greater his last bid, the more "pessimistic" the loser.

The winner's trade-off: winning a long time vs bidding low.

Equilibrium Conjecture

- Current bis are monotone nondecreasing in past bids .
- The players' last equilibrium bid summarizes their belief.
- The high-value bidder always bids at least $\underline{u}_+$.

The (High-Value) Winner's Problem ($n = 2$)

Given last bid b in period $t-1$ by the winner,

$$V_t(b) = \max_{\beta} \left\{ \frac{F_t(\beta)}{F_{t-1}(b)} \left((1-\delta)(\overline{u}-\beta) + \delta V_{t+1}(\beta)\right) \right\},$$

where F_t is the loser's c.d.f. Let expected continuation value of bid b be:

$$Y_t(b) := F_{t-1}(b) V_t(b),$$

then a version of the above Bellman equation is:

$$Y_t(b) = \max_{\beta} \left\{ \left((1-\delta)(\overline{u}-\beta) F_t(\beta) + \delta Y_{t+1}(\beta)\right) \right\}.$$

- now the continuation value of the winner is given by

$$Y_t(b)/(1-\delta) = \max_{\beta}\left\{F_t(\beta)(\overline{u}-\beta) + \delta Y_{t+1}(\beta)/(1-\delta)\right\}$$

- observe that it follows that $Y_t(b)$ is independent of b and hence $Y_t(b) = \varphi_t$ for some constant $\varphi_t \geq 0$
- but this allows us to determine F_t as

$$F_t(b) = \frac{(1-q_t)(\overline{u}-\underline{u})}{\overline{u}-b}$$

for all t, so the loser's bid is constant over time (from $t = 1$ onward).

The (High-Value) Loser's Strategy

With $n+1$ bidders:

$$F_t(b) = (1-q)\left(\frac{\overline{u}-\underline{u}}{\overline{u}-b}\right)^{1/n}.$$

All the losers' bids are constant over time (from $t=1$ onward).

In fact, the distribution is the same as in the static auction.

low value bidder: $\beta_i(\underline{u}) = \underline{u}$
high value bidder randomizes between $[\underline{b}, \overline{b}]$ with:

$$[\underline{b}, \overline{b}] = [\underline{u}, (1-q)\,\underline{u} + q\overline{u}]$$

and unique equilibrium distribution given by:

$$F(b) = \frac{1-q}{q}\frac{b-\underline{u}}{\overline{u}-b}$$

graphically for $q \in \{\frac{1}{4}, \frac{1}{2}, \frac{3}{4}\}$, $\underline{u} = 1$, $\overline{u} = 2$:

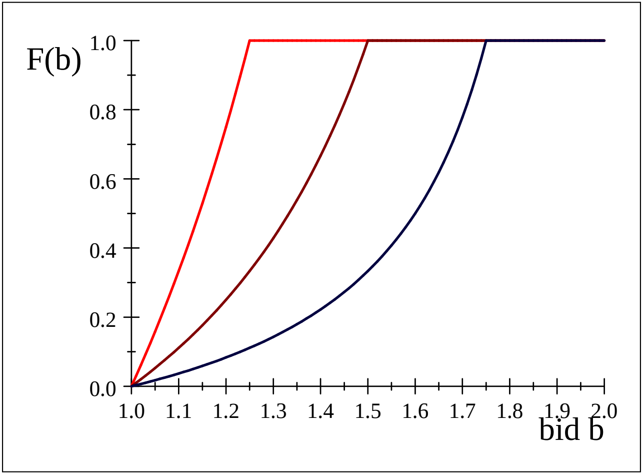

expected net utility for high valuation bidder: $(1-q)(\overline{u} - \underline{u})$

Given last bid b in period $t-1$,

$$W_t(b) = \max_{\beta}\left\{\frac{G_t(\beta)-G_t(b)}{1-G_{t-1}(b)}(1-\delta)(\overline{u}-\beta)+\delta\frac{1-G_t(\beta)}{1-G_{t-1}(b)}W_{t+1}(\beta)\right\},$$

where G_t is the winner's c.d.f.

- define component of continuation value from winning

$$X_t(b) \triangleq (1-G_{t-1}(b))\,W_t(b)$$

- allows us to rewrite above to

$$X_t(b) = \max_{\beta}\left\{(G_t(\beta)-G_{t-1}(b))(1-\delta)(\overline{u}-\beta)+\delta X_{t+1}(\beta)\right\}\ \ (t$$

while the previous trick does now longer work, first order and envelope conditions gives us:

$$\left(1-\delta\right)G_t'\left(b\right)\left(\overline{u}-b\right)=G_t\left(b\right)-G_{t-1}\left(b\right),$$

a difference-differential equation we can solve for $n+1$ bidders:

$$G_t(b)=F\left(b\right)\frac{1}{\delta^t}+F\left(b\right)^{\frac{1}{1-\delta}}\sum_{\tau=0}^{t}\frac{\left(1-\delta\right)^{\tau-t}}{\tau!}\left(\ln F\left(b\right)^{-\frac{1}{1-\delta}}\right)^{\tau}$$

with support $\left[\underline{u}_{+}\overline{u}-\left(1-q\right)^{n}\left(\overline{u}-\underline{u}\right)\right]$

The (High-Value) Winner's Strategy

The winner of $t = 0$ makes decreasing bids from $t = 1$ onward, until he loses.

The support of the bid distribution does not shrink.

As time passes, the distribution puts increasing weight on $\underline{u}_+$.

For fixed t, G_t converges to

$$(1-q)\left(\frac{\bar{u}-\underline{u}}{\bar{u}-b}\right)^{1/n},$$

the distribution of bids in a static first-price auction, as $\delta \longrightarrow 1$.

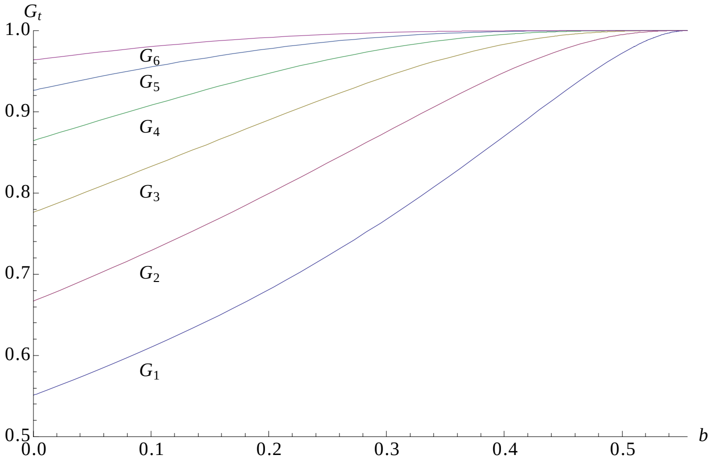

Figure: Bid distribution for $n = 2$ in periods $t = 1, \ldots, 6$, $q = 1/3$, $\delta = 9/10$, $\bar{u} = 1 = 1 - \underline{u}$ (bottom $t = 1$, top $t = 6$)

Illustration as $\delta \to 1$

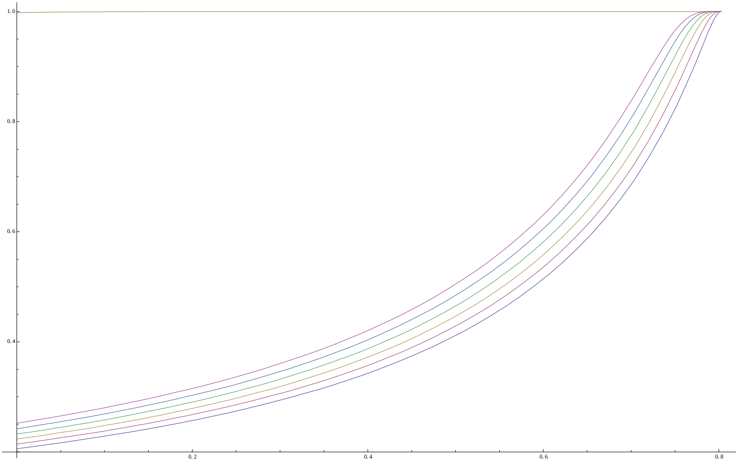

Figure: Bid distribution for $n = 4$ in periods $t = 1, \ldots, 6$ and 54, $q = 1/3$, $\delta = 99/100$, $\bar{u} = 1 = 1 - \underline{u}$ (bottom $t = 1$, top $t = 54$)

Bid Distribution in the First Period

Incentives at $t = 0$ determined by:

- Immediate reward $(1-\delta)F_0(b)(\overline{u}-b)$;
- Continuation payoff from winning $Y_t(b)$: independent of b;
- Continuation payoff from losing $W_t(b)$: decreasing in b.

As a result, bid in the first period lower than in the static auction.

With two bidders, it solves:

$$\frac{F_0(b)(\bar{u}-b)}{(1-q)(\overline{u}-\underline{u})} = \delta \ln F_0(b) + (1-\delta \ln(1-q)),$$

which leads to after transformation of variables into a Wishart function.

Illustration: Initial Distribution $n = 2$

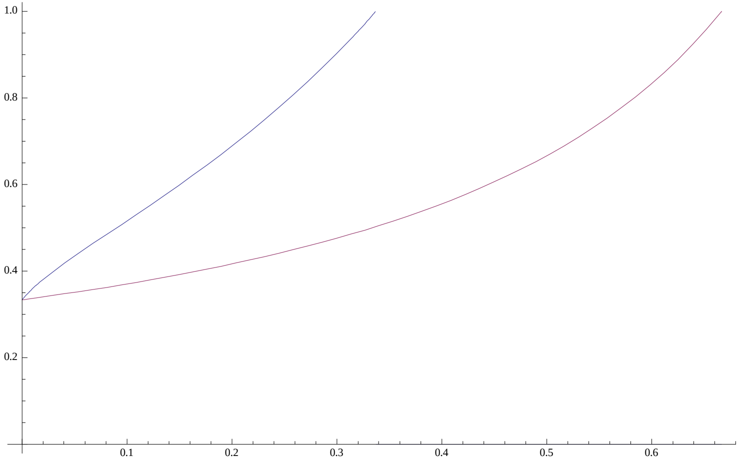

Figure: Initial distribution (blue), $q = 2/3$, $\delta = 9/10$, $\bar{u} = 1 = 1 - \underline{u}$, and bid distribution in the static auction (red)

- The allocation is efficient
(a higher-value bidder never loses to a lower-value bidder);
- The lowest type gets zero profits;
- Average expected revenue converge to the optimal auction (without reserve price) as $\delta \longrightarrow 1$.

- publicly observable bids
- information revelation by bidder i in period t:
- bid $\beta_{i,t} > \underline{u} \implies$ valuation has to be $u_i > \underline{u}$
- bid $\beta_{i,t} = \underline{u} \implies$ valuation can be $u_i = \underline{u}$ or $u_i = \overline{u}$
- bidding game begins with two-sided incomplete information, may turn into one-sided incomplete information and/or complete information game

Suppose a high-valuation bidder does not assign positive probability to the bid $\underline{u}$ in the initial period.
By bidding more, say $\underline{u}_+$, he gets

$$(1-q)^n(\bar{u}-\underline{u})$$

By deviating and bidding $\underline{u}$ in this period, followed by $\underline{u}_+$, a high-valuation bidder gets

$$(1-\delta)(1-q)^n\frac{\bar{u}-\underline{u}}{n+1}+\delta((1-q)^n+(1-\delta)nq(1-q)^{n-1})(\bar{u}-\underline{u})$$

On the other hand, separation yields the same payoff as the static auction,

$$(1-q)^n(\bar{u}-\underline{u})$$

Theorem

For all positive q and δ, there exists $\overline{n}$ such that for all $n > \overline{n}$, a separating Markov sequential equilibrium does not exist with observable bids.

In fact, comparing the above payoffs, we find that separation is *not* an equilibrium if, and only if

$$q \geq \underline{q}^o := \frac{1}{1 + (n+1)\delta}.$$

This condition, expressed in terms of the prior probability of a high valuation is satisfied if there are sufficiently many bidders and/or if the discount factor is sufficiently high.
Contrast this with the earlier result that showed with unobservable bids pooling is never an equilibrium

From Two-Sided to One-Sided Incomplete Information

- What then is the equilibrium of the game?
- It might occur that only one bidder reveals himself of high valuation, while all the other n bidders submit a bid $\underline{u}$
- hence we must understand the continuation game of one-sided incomplete information
- Informed bidder knows the valuation of his opponent
- Uninformed bidder doesn't know the valuation of his opponent

One-Sided Incomplete Information

- probability q_t that informed bidder has high valuation
- informed bidder's bid distribution F_t
- uninformed bidder's bid distribution G_t
- belief of uninformed bidder in terms of low valuation probability

$$1 - q_{t+1} = \frac{1 - q_t}{F_t\left(\underline{u}\right)}$$

- indifference condition of informed bidder

$$F_t\left(\underline{u}\right)\left(\overline{u} - \underline{u}\right) = \delta F_{t+1}^n\left(\underline{u}\right)\left(\overline{u} - \underline{u}\right)$$

yield

$$1 - q_0 = \prod_{t=0,\ldots,T} F_t(\underline{u})$$

Randomized Solution to the One-sided Case

- In all periods up to $T - 1$, (unknown) high-value bidders randomize between $\underline{u}$ and some distribution over $[\underline{u}, \bar{b}_t]$.

- In period T, the probability assigned to $\underline{u}$ is 0.
- The known high-value bidder randomizes on $[\underline{u}_+, \bar{b}_t]$ (up to T),
- So as $\delta \longrightarrow 1$, $V_U(q) \longrightarrow (1 - q)^n(\overline{u} - \underline{u})$.

- What then is the equilibrium of the overall game?
- A pooling equilibrium must involve all bidders submitting the bid $\underline{u}$.
- The payoff to a high-valuation bidder is then

$$(\bar{u} - \underline{u})/(n+1)$$

- The best deviation for a high-valuation bidder involves bidding $\underline{u}_+$, which garners

$$(1-\delta)(\bar{u} - \underline{u}) + \delta V^U(q)$$

Theorem

For all positive q, there exists $\left(\overline{\delta}, \overline{n}\right)$ such that for all $\delta > \overline{\delta}$ and $n > \overline{n}$, a pooling Markov sequential equilibrium does exist with observable bids.

A Preliminary Summary

Let us now focus on $\delta \to 1$. Because $V^U \to (1-q)^n(\bar{u} - \underline{u})$, we then get

$$q \geq \bar{q}^o \to 1 - (n+1)^{-1/n}.$$

Note that the left-hand side tends to the lowest payoff that a high-valuation bidder can guarantee, provided low-valuation bidders do not bid more than $\underline{u}$. As $\delta \to 1$:

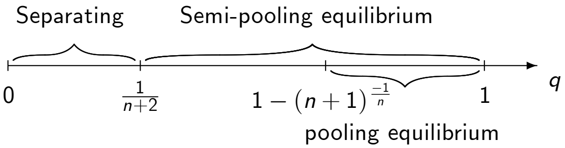

Similar to the case of observable bids, but:

If one known high-value, and n unknown bidders, no ratcheting.
For high-value bidder, $\underline{u}_+$ always dominates $\underline{u}$.

Learning only depends on the known bidder's bid (till he loses).

Known bidder willing to bid $\underline{u}_+$; lose against unknown high type:

$$V_U(q) = (1-q)^n(\overline{u} - \underline{u}).$$

No last period of separation.

Same condition as before the existence of a pooling equilibrium:

$$1 - q \leq (n+1)^{-1/n}.$$

However, a *separating* (and efficient) equilibrium always exists (if no other high type bids $\underline{u}$, why do so?)

To summarize:

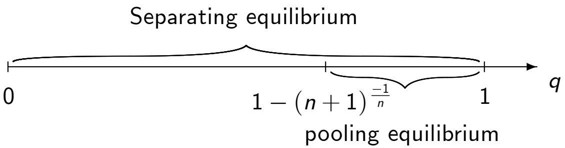

Conclusion

- repeated bidding and information disclosure
- construction of equilibrium strategies relied on dynamic programming, but not infinite horizon
- equilibrium strategies can be viewed as limits of finite horizon analysis
- privateness of bids enhanced competition and revenue
- revenue ranking of different disclosure policies with respect to past bids
- rationale for non-disclosure of past bids