

CSCE 475/875 Multiagent Systems
Handout 23: Collusion and Winner's Curse

March 30, 2020

(Based on Shoham and Leyton-Brown 2011)

Collusion

Cooperation between bidders to reduce their expected payments to the auctioneer by reducing competition among themselves is called *collusion*.

Collusion is usually illegal; however, interestingly enough, it is also very difficult for agents to pull off.

An interesting question to ask about collusion, therefore, is which *collusive protocols have the property that agents will gain by colluding while being unable to gain further by deviating from the protocol*.

Bidding ring protocol in Second-price auctions

First, consider a protocol for collusion in second-price (or Japanese/English) auctions.

We assume that a set of two or more colluding agents is chosen exogenously; this set of agents is called a *cartel* or a *bidding ring*.

Assume that the agents are risk neutral and have IPV valuations.

Assume there is an agent who is not interested in the good being auctioned, but who *serves to run the bidding ring*. This agent does not behave strategically, and hence could be a simple computer program.

We will refer to this agent as the *ring center*. Observe that there may be agents who participate in the main auction and do not participate in the cartel; there may even be multiple cartels.

The protocol follows.

1. Each agent in the cartel submits a bid to the ring center.
2. The ring center
 - a. identifies the maximum bid that it received, \hat{v}_1^r ;
 - b. submits this bid in the main auction; and
 - c. drops the other bids.(Denote the highest dropped bid as \hat{v}_2^r .)
3. If the ring center's bid wins in the main auction (at the second-highest price in that auction, \hat{v}_2), the ring center awards the good to the bidder who placed the maximum bid in the cartel and requires that bidder to pay $\max(\hat{v}_2, \hat{v}_2^r)$.
4. The ring center gives every agent who participated in the bidding ring a payment of k , regardless of the amount of that agent's bid and regardless of whether or not the cartel's bid won the good in the main auction.

How should agents bid if they are faced with this bidding ring protocol?

First of all, consider the case where $k = 0$. Here it is easy to see that this protocol is strategically equivalent to a second-price auction in a world where the bidder's cartel does not exist. The high bidder always wins, and always pays the globally second-highest price (the max of the second-highest prices in the cartel and in the main auction). Thus the auction is dominant-strategy truthful, and agents have no incentive to cheat each other in the bidding ring's "preauction." At the same time, however, agents also do not gain by participating in the bidding ring: they would be just as happy if the cartel disbanded and they had to bid directly in the main auction.

But what about the ring center? What if $\hat{v}_2^r > \hat{v}_2$? It will pay \hat{v}_2 for the good in the main auction, but it will be paid \hat{v}_2^r for it by the winning bidder. Let $c > 0$ denote the ring center's expected profit. If there are n_r agents in the ring, then the ring center could pay each agent up to $k = \frac{c}{n_r}$ and still budget balance on expectation! For values of k smaller than this amount but greater than zero, the ring center will profit on expectation while still giving agents a strict preference for participation in the bidding ring.

How are agents able to gain in this setting—doesn't the revenue equivalence theorem say that their gains should be the same in all efficient auctions?

Observe that the agents' expected payments are in fact unchanged, **although not all of this amount goes to the auctioneer.**

What does change is the *unconditional* payment that every agent receives from the ring center. **The second condition of the revenue equivalence theorem states that a bidder with the lowest possible valuation must receive zero expected utility.** This condition is violated under our bidding ring protocol, in which such an agent has an expected utility of k .

Bidding ring protocol in First-price auctions

The construction of bidding ring protocols is **much more difficult** in the first-price auction setting. Why?

In order to make a lower expected payment, the winner must actually place a lower bid. In a second-price auction, a winner can instead persuade the second-highest bidder to leave the auction and make the same bid he or she would have made anyway. This difference matters because in the second-price auction the second-highest bidder has no incentive to renege on his or her offer to drop out of the auction; by doing so, it can only make the winner pay more. In the first-price auction, the second-highest bidder could trick the highest bidder into bidding lower by offering to drop out, and then could still win the good at less than its valuation. (*Note:* Treacherous!)

So this is problematic and thus for the most part the literature on collusion on first-price auction assumes that all n bidders belong to the cartel.

Common Values & Winner's Curse

In interdependent values, agents' valuations depend on *both* their own signals and other agents' signals.

In **common values**, all agents value the good at exactly the same amount.

The twist is that the agents do *not* know this amount, though they have (common) prior beliefs about its distribution. Each agent has a **private** signal about the value, which allows it to condition its prior beliefs to arrive at a posterior distribution over the good's value.

For example, consider the problem of buying the rights to drill for oil in a particular oil field. The field contains some (uncertain but fixed) amount of oil, the cost of extraction is about the same no matter who buys the contract, and the value of the oil will be determined by the price of oil when it is extracted.

- Given publicly available information about these issues, all oil drilling companies have the same prior distribution over the value of the drilling rights.
- The difference between agents is that each has different geologists who estimate the amount of oil and how easy it will be to extract, and different financial analysts who estimate the way oil markets will perform in the future.
- These signals cause agents to arrive at different posterior distributions over the value of the drilling rights, based on which, each agent i can determine an expected value v_i .

How can this value v_i be interpreted?

One way of understanding it is to note that if a single agent i was selected at random and offered a take-it-or-leave-it offer to buy the drilling contract for price p , it would achieve positive expected utility by accepting the offer if and only if $p < v_i$.

Now consider what would happen if these drilling rights were sold in a second-price auction among k risk-neutral agents. One might *expect* that each bidder i ought to bid v_i .

However, it turns out that bidders would achieve **negative expected utility** by following this strategy.

How can this be—didn't we previously claim that i would be happy to pay any amount up to v_i for the rights?

The catch is that, since the value of the good to each bidder is the same, each bidder cares as much about *other* bidders' signals as it does about its own. When it finds out that it won the second-price auction, the winning bidder also learns that it had the *most* optimistic signal!

This information causes the winning bidder to *downgrade his or her expectation about the value* of the drilling rights, which can make him or her conclude that he or she paid too much!

This phenomenon is called the **winner's curse**.

(**Note:** Of course, the winner's curse does not mean that in the CV setting the winner of a second-price auction always pays too much. Instead, it goes to show that truth telling is no longer a dominant strategy (or, indeed, an equilibrium strategy) of the second-price auction in this setting.)