**CIS 700 Differential Privacy in Game Theory and Mechanism Design** April 4, 2014

Lecture 10

*Lecturer: Aaron Roth* *Scribe: Aaron Roth*

### Asymptotic Dominant Strategy Truthfulness in Ascending Price Auctions

# 1 Introduction

At the very beginning of the semester, we saw an incredibly powerful tool in mechanism design: the VCG mechanism. You remember: the VCG mechanism lets us find the welfare-optimal solution for any social choice problem, and pair it with payments that makes truthful reporting of one's preferences a *dominant strategy* for every player. Pretty impressive!

Yet, one-shot direct revelation mechanisms are rarely used in practice. There are a couple of reasons for this that are commonly cited:

1. It can be a non-trivial task for an agent to report their type. In general, they have to pin down a value for every possible bundle they might receive, and quote this value to the mechanism. Even when this is a concise set (for example, with unit demand bidders in an auction with $m$ goods, they only need to report $m$ values), it can be hard for people to decide exactly how much they value each good when asked.

2. A direct revelation mechanism requires that agents provide (possibly) more information to the mechanism than it needs to solve the allocation problem. For example, in a single item auction, the VCG mechanism is just a second price auction. Nevertheless, the winning bidder is forced to reveal to the mechanism his value for the good, even though all that was necessary to reveal was that his value was higher than that of the 2nd highest bidders. Hence if the bidder is concerned about privacy, he might prefer not to reveal more than necessary.

Instead, people often prefer to run *iterative ascending price auctions*. These are just (generalizations of) the kind of auctions you see on TV. At any moment in time, each good being sold has a price. When it is a bidder's turn to bid, she only need decide which good she wants to bid on (or if she wants to drop out), which is equivalent to asking her "What is your favorite good at the current prices?" They resolve some of the difficulties of the VCG mechanism listed above because:

1. "What is your favorite good at the current prices?" is an easier question to answer than to quantify your entire valuation function, so such auctions can be less demanding on the bidders, and

2. Ascending auctions halt at the final price, and so are more parsimonious in what they require bidders to report. For example, the high bidder in a single item auction never has to reveal whether his value is $\epsilon$ more than the 2nd highest bidder, or 1000 times higher.

However, there are two problems that we might hope to resolve. First, the "privacy" guaranteed by an ascending price auction is qualitative and informal. If we could implement such an auction such that the prices were differentially private in the actions of the players (and the allocation jointly differentially private), then we could make the privacy guarantee formal.

The second has to do with the strength of implementation. The VCG mechanism is dominant strategy truthful, but truth-telling (called *sincere bidding* in ascending price auctions) is usually *not* a dominant strategy in ascending price auctions. The reason is easily captured in the following simple example, and is informally because in an iterative setting, other bidders can threaten you.

**Example:** Suppose we have two unit demand bidders 1 and 2, and two goods for sale $a$ and $b$. We have $v_{1,a} = 1, v_{1,b} = \epsilon$ and $v_{2,a} = 1/2, v_{2,b} = 1/2 - \epsilon$. Suppose moreover that bidder 2 takes the following strategy: "Bid on good $a$. If bidder 1 bids on good $a$, then outbid him on whatever he bids on until the price is $\geq 1$." Against this strategy, bidder 1 cannot obtain non-negative utility if he bids on his favorite good ($a$), and so his best response is to place an insincere bid on good 2. Moreover, bidder 2 has a clear motivation to take this threatening position – he obtains substantially higher payoff than if players followed sincere bidding, since he gets his most preferred good without any competition.

Because of strategies like in the above example, classical implementations of ascending price auctions do not implement sincere bidding as a dominant strategy, but instead as a Nash equilibrium. This is the second improvement we might hope to make: if we can implement an ascending price auction such that the prices are differentially private, then price-taking behavior will be approximately optimal, or in other words, sincere bidding will be an asymptotic dominant strategy for every player.

# 2 An Ascending Price Auction for Unit Demand Bidders

Here we will describe a classical ascending price auction for unit demand bidders due to Crawford and Knoer [CK81]. (This auction generalizes to bidders with *gross substitutes valuations*, as shown by Kelso and Crawford [KJC82]. Our private implementation, and its corresponding incentive properties also generalize to gross substitute valuations, but for simplicity we'll stick to unit demand bidders here). Formally, suppose there are $n$ bidders $B$ and $n$ goods $G$. A bidder $i$ is *unit-demand* if for every $S \subseteq G$, $v_i(S) = \max_{j\in S} v_i(\{j\})$ – i.e. if he only wants *one* good. Such a bidder's valuation function can be expressed as just $n$ numbers $v_{i,1}, \ldots, v_{i,n}$ representing his value for each of the $m$ goods: $v_{i,j} \equiv v_i(\{j\})$. For this lecture, lets assume $v_{i,j} \in [0, 1]$ for all $i, j$.

We describe the Crawford/Knoer Auction which builds a matching $\mu : B \to G$ between bidders and goods. (i.e. a function $\mu$ such that for every good $j$, $|\mu^{-1}(j)| \leq 1$: every bidder is matched to at most one good, and every good is matched to at most one bidder). We allow bidders to be unmatched, which we write as $\mu(i) = \emptyset$. We define $v_{i,\emptyset} = 0$ for each bidder $i$.

---

**Algorithm 1** The Crawford-Knoer Ascending Price Auction for Unit Demand Bidders. It takes as input a bid-increment $\alpha$.

---

**Ascend**($\alpha$):
  **For each** bidder $i$ **Let** $\mu(i) = \emptyset$.
  **For each** good $j \in [m]$, **Let** $p_j \leftarrow 0$
  **while** There exist unmatched bidders **do**
    **for** $i = 1$ to $n$ **do**
      **if** Bidder $i$ is unmatched and $\exists j$ such that $v_{i,j} - p_j > 0$ **then**
        **Ask** bidder $i$ for any $j^* \in \arg\max_{j\in G}(v_{i,j} - p_j)$.
        **Let** $\mu(\mu^{-1}(j^*)) \leftarrow \emptyset$ (high bidder of $j^*$ is now unmatched) and $\mu(i) \leftarrow j^*$ (Good $j^*$ is matched to $i$).
        **Let** $p_{j^*} \leftarrow p_{j^*} + \alpha$
      **end if**
    **end for**
  **end while**
  **Allocate** each bidder $i$ $\mu(i)$ and charge them $p_{\mu(i)}$.

---

Lets note a couple of things about this auction. First, it always halts and outputs a matching:

**Lemma 1** *Ascend halts after at most $n/\alpha$ bids are made.*

**Proof** No bidder ever bids on an item with price $\geq 1$, since this would result in negative utility. Hence at the end of the auction, we have $p_j < 1$ for all $j$ and $\sum_{j=1}^{n} p_j < n$. But each bid increases $\sum_{j=1}^{n} p_j$ by $\alpha$, and so the claim follows. ■

Next, when it ends, everyone is getting (approximately) their most preferred good given the final prices. Such an outcome is called an approximate Walrasian equilibrium:

**Definition 2** *Given a set of bidder valuations $v$, a matching $\mu$ together with a vector of prices $p$ forms an $\alpha$-approximate Walrasian equilibrium if for every agent $i$:*

$$v_{i,\mu(i)} - p_{\mu(i)} \geq \max_j \left(v_{i,j} - p_j\right) - \alpha$$

**Lemma 3** *Assuming sincere bidding, at completion, the output of Ascend is an $\alpha$-approximate Walrasian equilibrium.*

**Proof** By inspection. At the time that bidder $i$ was last matched to $\mu(i)$, it was his most preferred good. Subsequently, the price $p_{\mu(i)}$ was incremented by at most $\alpha$, and all other prices only increased. ■

Lets consider whether the fact that the auction ends with approximate Walrasian equilibrium prices imparts any useful incentive properties to the mechanism. We can observe that if the prices were given from on-high, and fixed independent of the bidder's behavior, this Walrasian equilibrium property would mean that sincere bidding was an approximately dominant strategy – since surely no strategy can do better than giving a bidder his most preferred good! However, this is not the case, because the prices are computed as a function of the bidding behavior of the agents. It could be that by misrepresenting his demand, a bidder can gain some benefit by changing how the prices are computed. However, if we could maintain the prices while satisfying differential privacy throughout the run of the auction, we would have that sincere bidding remains an approximately dominant strategy. This will be the goal of this lecture.

Finally, lets note that any approximate Walrasian equilibrium must be near welfare optimal:

**Theorem 4** *Let $OPT = \max_\nu \sum_{i=1}^{n} v_{i,\nu(i)}$ be the welfare of the max-welfare matching. Then for any matching $\mu$ that is part of an $\alpha$-approximate Walrasian equilibrium we have:*

$$\sum_{i=1}^{n} v_{i,\mu(i)} \geq OPT - \alpha n$$

*(Note that by Lemma 3 such a $\mu$ results when players sincerely bid in Ascend.)*

**Proof** Let $\nu$ be the welfare optimal matching. By assumption, for all $i$:

$$v_{i,\mu(i)} - p_{\mu(i)} \geq \left(v_{i,\nu(i)} - p_{\nu i}\right) - \alpha$$

Summing over all $i$ gives:

$$\sum_{i=1}^{n}(v_{i,\mu(i)} - p_{\mu(i)}) \geq \sum_{i=1}^{n}\left(v_{i,\nu(i)} - p_{\nu i}\right) - \alpha n$$

Since $\mu$ and $\nu$ are matchings, we can rewrite the sums as:

$$\sum_{i=1}^{n} v_{i,\mu(i)} - \sum_{j:\mu^{-1}(j)\neq\emptyset} p_j \geq \sum_{i=1}^{n} v_{i,\nu(i)} - \sum_{j:\nu^{-1}(j)\neq\emptyset} p_j - \alpha n$$

Note however that by inspection of the auction, for every good $j$ such that $\mu^{-1}(j) = \emptyset$ (i.e. all unmatched goods), we have $p_j = 0$. Hence, we have:

$$\sum_{i=1}^{n} v_{i,\mu(i)} \geq \sum_{i=1}^{n} v_{i,\nu(i)} + \sum_{j:\nu^{-1}(j)=\emptyset} p_j - \alpha n \geq \sum_{i=1}^{n} v_{i,\nu(i)} - \alpha n$$

which proves the claim. ■

## 3 Privacy and Sincere Bidding

Let us argue that any ascending price auction that converges to an $\alpha$-approximate Walrasian equilibrium, with prices that are differentially private in the bids of the agents makes sincere bidding an approximate dominant strategy.

We'll be a little informal here. Lets say that an *iterative auction protocol* $M$ is an algorithm that in $T$ rounds, adaptively chooses a bidder to *query*. When a bidder is queried at round $t$, she is shown a set of prices $p^t \in [0,1]^n$, and may respond with the name of a good $g^t \in [n]$. Based on $g^1, \ldots, g^t$, the protocol can update the prices $p^{t+1}$ and decide who to query next. At the end, the protocol outputs the final set of prices $P^T$ together with a matching $\mu$. Each bidder $i$ receives good $\mu(i)$, and pays price $p_i$.

A *bidding strategy* for a bidder is any function $f$ that maps price histories $p^1, \ldots, p^{t-1}$ to goods $g^t$: formally, a function $f : ([0,1]^n)^* \to [n]$. We can think of an iterative auction protocol as an algorithm that takes as input $n$ bidding strategies $f_1, \ldots, f_n$ of the players, runs the auction by querying the bidding strategies, and then outputs the resulting allocation and payments.

The *sincere* bidding strategy for a player $i$ is the bidding strategy that always outputs the most preferred good for player $i$ at the current prices. In other words:

$$f_i^s(p^1, \ldots, p^t) \equiv f_i^s(p^t) \stackrel{\text{def}}{=} \arg\max_{j \in [n]} (v_{i,j} - p_j^t)$$

**Theorem 5** *Let $M$ be an iterative auction protocol such that any player who practices sincere bidding ends up being matched to a good $j^*$ such that $v_{i,j^*} - p_{j^*} \geq \max_j (v_{i,j} - p_j) - \alpha$, where $p$ are the prices computed by the mechanism. Suppose also that the computation of the prices $p$ is $\epsilon$-differentially private in the bidding strategies of the players. Then sincere bidding is an $\eta$-approximate dominant strategy for:*

$$\eta = \epsilon + \alpha$$

**Proof** Recall that the utility that player $i$ receives under a matching $\mu$ and prices $p$ is $u_i(\mu, p) = v_{i,\mu(i)} - p_{\mu(i)}$. We fix any vector of bidding strategies $f_{-i}$ for players $j \neq i$ and consider bidder $i$'s utility under sincere bidding $f_i^s$ compared to his utility if he deviates to any other bidding strategy $f_i'$.

$$\begin{aligned}
\mathrm{E}_{\mu,p\sim M(f_i^s, f_{-i})}[u_i(\mu, p)] &= \mathrm{E}_{\mu,p\sim M(f_i^s, f_{-i})}[v_{i,\mu(i)} - p_{\mu(i)}] \\
&\geq \mathrm{E}_{p\sim M(f_i^s, f_{-i})}[\max_j v_{i,j} - p_j] - \alpha \\
&\geq \exp(-\epsilon)\mathrm{E}_{p\sim M(f_i', f_{-i})}[\max_j v_{i,j} - p_j] - \alpha \\
&\geq \mathrm{E}_{\mu,p\sim M(f_i', f_{-i})}[u_i(\mu, p)] - \epsilon - \alpha
\end{aligned}$$

■

## 4 Making the Crawford Knoer Auction Private

It remains to sketch conditions under which Ascend can be implemented such that the final prices are differentially private in the strategies of the bidders. To do so, we introduce the idea of *identical goods*.

We want to consider markets in which goods are not unique, but come in at least some small supply. (For example, if you want to buy a 24 inch Samsung LCD television, there is not only one such good, but many, and you are indifferent between receiving any one of them at a given price). Two goods $g_1$ and $g_2$ are said to be *identical* if for every feasible valuation function $v_i$, $v_{i,g_1} = v_{i,g_2}$. We say that a set of $m$ goods consists of $k$ *types* of goods if there exist goods $g_1, \ldots, g_k$ such that for every other good $g \in [m]$, $g$ is identical to one of $g_1, \ldots, g_k$. We say that the *supply* corresponding to a good of type $g_i$ is the number of goods that are identical to $g_i$.

We begin by introducing a change to the Crawford/Knoer auction that has no effect on the bidding behavior of people who choose to bid sincerely. Consider what happens in the Crawford Knoer auction to the prices of $s$ copies of an identical good: because everyone values the $s$ copies identically, nobody ever bids on any copy of a good other than the one that is currently the *least expensive*. Hence, rather than reporting the prices of all $s$ copies of the good to each bidder, it suffices to report for each *type* of good, the price of the copy of that type of good that is currently the least expensive. This *minimum* price increments every time the type of good receives another $s$ bids. More precisely, the (minimum) price that a good of type $j$ is available at at time $gt$ is exactly $p_j = \lfloor b_j(t)/s \rfloor$, where $b_j(t)$ is the number of bids that have been placed up through time $t$ on goods of type $j$.

Thus, the Crawford/Knoer auction can equally well be run as follows:

- A price $p_j \leftarrow 0$ is initialized for each *type* of good $j \in [k]$.
- In rounds, unmatched bidders bid on their most preferred *type* of good at the current prices $p$.
- A count $b_j(t)$ is maintained on the number of bids each type of good has received at time $t$, and $p_j = \lfloor b_j(t)/s \rfloor$.
- Bidders become matched to the goods that they bid on, and become unmatched after $b_j(t)$ has incremented more than $s$ ticks from its value when they bid on the good. (and so at most $s$ units of any good are allocated).

Note that this is simply a different description of the same algorithm that we already analyzed (when bidding is sincere – but note that we have reduced the strategy space for non-sincere bidding a bit), and so we continue to have that the algorithm halts and outputs a matching after at most $m/\alpha$ steps (for $m = k \cdot s$, the total number of goods), that the matching has weight at least $\text{OPT} - \alpha m$, and that for every bidder $i$ who follows a sincere bidding strategy, they are matched to a good $\mu(i)$ such that $v_{i,\mu(i)} - p_{\mu(i)} \geq \max_j (v_{i,j} - p_j) - \alpha$.

Note however that with the above description, a sufficient statistic to make available to allow all bidders to coordinate their bidding behavior is a running count $b_j(t)$ on the number of bids each *type* of good has received so far. Given that bidders know these counts, they know both:

1. The current price for each type of good, which allows them to decide what to bid on, and
2. Whether they are currently matched or unmatched, which allows them to know whether they should bid at all. (They become matched on any day that they bid, and unmatched on the day that the bid-count on their good increments $s$ above where it was when they bid).

But this should seem promising: Recall, we learned about a tool to privately keep track of running sums:

**Theorem 6 ([DNPR10, CSS10])** *There is an $\epsilon$-differentially private algorithm that is simultaneously $(E, \beta)$-accurate on $k$ streams of length $T$ that jointly have sensitivity $\Delta$ for:*

$$E = O\left( \frac{\Delta \cdot \log\left(\frac{T \cdot k}{\beta}\right)^{5/2}}{\epsilon} \right)$$

Here we need to maintain $k$ separate streams, one to count bids on each of the $k$ types of goods. And what are the sensitivity $\Delta$ of our streams? Well, at each time step $t$, a stream for a good of type $j$ receives a 1 if there was a bid on that type of good, and a 0 otherwise (so that the running sums on the streams are bid counts). But each agent bids at most 1 time on each type of good at each price point: So the total number of bids a single person can make is $\Delta \leq k/\alpha$.

Hence, we can run the auction while maintaining bid counters that have error:

$$E = \tilde{O}\left(\frac{k \cdot \log{(n \cdot k)}^{5/2}}{\alpha\epsilon}\right)$$

Lets consider the implication of what happens when we run the auction with noisy bid counters, rather than exact bid counters. First note that since the price of any good only increments after $s$ bids have been made, error $E$ in the bid counters only results in error $\leq \alpha$ in the prices under the condition that $s \geq E$.

But there is another problem: when we run the auction with exact bid counters, then any over-demanded type of good has its supply constraint exactly satisfied: exactly $s$ people are matched to that type of good. Here, because of the noisy counters, the number of people who end up matched to a particular type of good may be any value in $s \pm E$. Therefore, if we don't want to violate supply constraints, we must *reserve* $E$ copies of each good outside of the auction, to satisfy demand that results from the auction's possible over-allocation of $E$ units. But if we are reserving $E$ copies of the good, we must be sure that this does not reduce welfare by more than a $(1-\alpha)$ factor (which we are already losing do to the approximation guarantee of the auction). However, this will be the case whenever $s \geq \frac{E}{\alpha}$.

Putting this all together, we have sketched the argument for the following theorem

**Theorem 7** *There exists an ascending price auction for unit demand bidders bidding on $k$ types of items, each with supply at least $s$ such that:*

1. *Sincere bidding is an $\alpha$-approximate dominant strategy,*
2. *Results in differentially private prices and a jointly differentially private allocation, and*
3. *Results in an outcome that achieves welfare at least $OPT - \alpha m$*

*whenever the supply of each good is at least:*

$$s \geq \tilde{O}\left(\frac{k \cdot \log{(n \cdot k)}^{5/2}}{\alpha^3}\right)$$

*In particular, if the supply of each type of good grows slightly superlinearly in the number of types of goods as the market grows large, then we get an auction that is asymptotically dominant strategy truthful and welfare optimal.*

# 5 Discussion

These results can be generalized to bidders who have *gross substitutes valuations*, using the generalization of the Crawford/Knoer auction due to Kelso and Crawford [KJC82].

As with many of the "large market" results that we have seen in this class, the advantage that we have over many of the large market results from the economics literature is the weakness of the assumptions that we need to make. Note that we needed to make *no* assumptions on the valuation functions of the bidders – i.e. they need not be drawn from a prior, nor generated as part of a replication economy, and bidders can be completely unique. Our only "large market" assumption is that the supply of each good is large – moderately larger than the number of distinct types of items.

**Bibliographic Information** The ascending price auction that we study in this lecture is due to Crawford and Knoer [CK81], and generalized by Kelso and Crawford [KJC82].

The main result in this lecture is based on ongoing work with Justin Hsu, Zhiyi Huang, Tim Roughgarden, and Steven Wu, as well as the results in [HHR+14] (Which give a private implementation of the Kelso/Crawford auction that achieves better bounds than those quoted here, but at the expense of incentive properties).

# References


[CK81] Vincent P Crawford and Elsie Marie Knoer. Job matching with heterogeneous firms and workers. *Econometrica*, 49(2):437–450, 1981.

[CSS10] TH Hubert Chan, Elaine Shi, and Dawn Song. Private and continual release of statistics. In *Automata, Languages and Programming*, pages 405–417. Springer, 2010.

[DNPR10] Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. Differential privacy under continual observation. In *Proceedings of the 42nd ACM symposium on Theory of computing*, pages 715–724. ACM, 2010.

[HHR+14] Justin Hsu, Zhiyi Huang, Aaron Roth, Tim Roughgarden, and Zhiwei Steven Wu. Private matchings and allocations. *Proceedings of the 41st annual symposium on the Theory of Computing (STOC)*, 2014.

[KJC82] Alexander S Kelso Jr and Vincent P Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica: Journal of the Econometric Society*, pages 1483–1504, 1982.