The Clinching Auction

CSCI 1440/2440

2024-03-20

We introduce a new multiparameter setting, for which we can find an approximately welfare-maximizing EPIC auction. We prove the EPIC property by making use of our recipe for doing so: 1. we prove sincere bidding in the auction yields an approximate VCG outcome, and 2. we show consistent bidding strategies dominate inconsistent ones. Not only does this mechanism satisfy desired performance and incentive guarantees (up to some additive error), it is also tractable.

1 *Diminishing Marginal Valuations*

We introduce a new instance of the multiparameter setting, so-called **diminishing marginal valuations** for homogeneous goods. Rather than additive or unit-demand valuations, here each bidder's marginal value for an additional copy of the good is weakly decreasing.

- We assume n bidders and m homogeneous (i.e., identical) goods, with bidders indexed by i, and goods, by j.
- Each bidder i has a marginal value $\mu_i(j)$ for its jth copy of the good, meaning its value for acquiring a jth copy of the good, given they already have $j-1$ copies in its possession.
- Each bidder i's marginal values are weakly decreasing: $\mu_i(1) \geq \mu_i(2) \geq \cdots \geq \mu_i(m)$.

Our goal is to construct an approximately welfare-maximizing EPIC multiunit ascending auction for this scenario.

2 *A Direct Mechanism: A Sanity Check*

Designing a welfare-maximizing DSIC direct mechanism "reduces to"[1] designing a welfare-maximizing EPIC indirect mechanism, in the sense that a polynomial-time solution to the latter can be used to construct a polynomial-time solution to the former via the revelation principle. Therefore, solving for a welfare-maximizing EPIC indirect mechanism is at least as hard as solving for a welfare-maximizing DSIC direct mechanism. As a result, if no such DSIC direct mechanism exists (one that is welfare-maximizing in polynomial time), neither can such an EPIC indirect mechanism.

[1] See Appendix A

As a result of this argument, before we embark upon the design of an EPIC indirect mechanism that maximizes welfare in polynomial time, we best do a quick sanity check: can we design a DSIC direct mechanism that maximizes welfare in polynomial time?

Fortunately, we can solve this problem in the affirmative in the direct setting. In particular, the welfare-maximizing allocation can be computed via a simple greedy allocation algorithm:

- Collect a vector of bids $\mathbf{b}$ from all bidders $i \in [n]$, with bid $b_i(j)$ representing i's bid on the jth copy of the good.
- Sort the bids, and then allocate the goods to the bidders who submitted the highest m bids, breaking ties arbitrarily.

For example, if $b_i(4)$ is among the highest m bids, but $b_i(5)$ is not, then bidder i is allocated four copies, which we denote as $x_i = 4$.

As usual, to achieve the VCG outcome, we combine this allocation algorithm with payments that charge bidders their externalities. For each bidder i, we sort the bids by *bidders other than bidder i* from greatest to least, and then establish the following groupings.

$$\underbrace{\beta_1 \quad \beta_2 \quad \cdots \quad \beta_{m-x_i}}_{A} \quad \underbrace{\beta_{m-x_i+1} \quad \beta_{m-x_i+2} \quad \cdots \quad \beta_m}_{B} \quad \underbrace{\beta_{m+1} \quad \beta_{m+2} \quad \cdots \quad \beta_{mn}}_{C}$$

The bids in group A are those that are allocated regardless of i's presence. The bids in group C are those that are *not* allocated regardless of i's presence. The bids in group B are those whose allocation depends on i's presence. These bids comprise bidder i's externality. We therefore charge bidder i, in total, for all x_i copies in group B, the sum of these x_i bids: i.e.,

$$p_i(x_i) = \sum_{j=1}^{x_i} \beta_{m-x_i+j}.$$

By charging each bidder its externality, we charge the VCG payments, thereby guaranteeing the DSIC property.

Example 2.1. Here is an example from the paper that introduced the clinching auction,[2] the subject of this lecture. It is loosely based on the first US Nationwide Narrowband spectrum auction in where there were five bidders and five licenses.

[2] Lawrence M. Ausubel. An efficient ascending-bid auction for multiple objects. *American Economic Review*, 94(5):1452–1475, December 2004

The bidders' marginal values are as follows:

License	A	B	C	D	E
First	123	74	125	84	44
Second	113	5	125	64	24
Third	103	3	49	7	5

If we interpret the values in this table as bids, and sort them from high to low, we arrive at 125, 125, 123, 113, 103, etc., which implies that C wins two licenses, and A wins three. It remains to determine the VCG payments associated with this allocation.

Sorting the bids *sans* C yields: 123, 113, 103, 84, 74, etc. As C's two bids of 125 displace D, who bid 84, and B who bid 74, C's VCG payments are 84 and 74.

Similarly, sorting the bids *sans* A yields: 125, 125, 84, 74, 64, etc. As A's three bids of 123, 113, and 103 displace D, who bid 84, B who bid 74, and D again who bid 64, A's VCG payments are 84, 74, and 64.

Once again, when bidder i is allocated x_i copies of the good, its VCG payment is the sum of x_i bids, one per copy of the good. If we imagine adding i's bids to the sorted list of all other bidders' bids, *one at a time*, each additional winning bid placed by i displaces another lower bid. The smallest bid that i displaces can be viewed as i's payment for its first copy of the good; the second-smallest bid is then i's payment for its second copy of the good; and so on. Payments for additional copies of a good are thus *weakly increasing* (even though values are—by assumption—weakly decreasing).

Building on these observations, we can express bidder i's VCG payment for good j in terms of the other bidders' demands. Define bidder k's **demand set** at price q, $D_k(q) = \max_j \{ j \le m \mid \mu_k(j) \ge q \}$. Now, bidder i's payment for the jth copy is given by:

$$p_i(j) = \inf \left\{ q \;\middle|\; \sum_{k \neq i} D_k(q) \le m - j \right\}. \tag{1}$$

This payment is the price at which the demand of all bidders other than i falls below $m - j$. As expected, these payments are weakly increasing in j. Each additional copy of the good costs *no less* than the previous, as other bidders' demands fall as the price rises.

3 *An Indirect Mechanism: The Clinching Auction*

Having satisfied the precondition for potential success, we now set our sights on a welfare-maximizing EPIC ascending auction. To this end, we present the **clinching auction**:

- Initialize $q = 0$.
- Collect demand sets from all bidders. (Initially, when $q = 0$, it should be that $D_i(q) = m$, for all bidders i.)
- Alternate between incrementing q by ϵ and collecting demand sets from bidders until $\sum_{i=1}^{n} D_i(q) \le m$.
- Activity rule: Ensure that no bidder's demand increases over time: i.e., bidder's demands can only decrease as prices increase.

- Let q^* denote the final price. At this price, $\sum_{i\in[n]} D_i(q^*) \leq m < \sum_{i\in[n]} D_i(q^* - \epsilon)]$. (If $\sum_{i\in[n]} D_i(q^* - \epsilon)]$ were equal to m, then the auction would have terminated at price $q^* - \epsilon$.) We can thus allocate y_i copies to bidder i s.t. $y_i \in [D_i(q^*), D_i(q^* - \epsilon)]$ and $\sum_i y_i = m$.

  One way to achieve such an allocation is to allocate to bidder i her final demand, namely $D_i(q^*)$ copies. Then, if any unallocated goods remain, they can be allocated at random to any bidders with leftover demand at price $q^* - \epsilon$: i.e., bidders i for whom $D_i(q^* - \epsilon) - D_i(q^*) > 0$.

- Charge bidder i (within ϵ of) its externality. Specifically, charge bidder i for its jth copy of the good:

$$q_i(j) = -\epsilon + \min_{t\in\mathbb{Z}^+} \left\{ \epsilon t \;\middle|\; \sum_{k\neq i} D_k(\epsilon t) \leq m - j \right\}. \qquad (2)$$

  As intended, this price is (near) the price at which the demand of all other bidders falls below $m - j$.

  N.B. When $\sum_{k\in[n]} D_k(\epsilon t) = m$, it is not necessary to subtract ϵ from ϵt. But when $\sum_{k\in[n]} D_k(\epsilon t) < m$, the situation is analogous to the last remaining bidders all dropping out at the same time in an English or a Japanese auction, in which case the (one) good is sold at the final price less ϵ to ensure individual rationality.

Example 3.1. Continuing the setup in Example 2.1, the clinching auction then proceeds as follows, with demands depicted only at the most relevant prices:

Price	A	B	C	D	E
10	3	1	3	2	2
25	3	1	3	2	1
45	3	1	3	2	0
50	3	1	2	2	0
65	3	1	2	1	0
75	3	0	2	1	0
85	3	0	2	0	0

At price 65, the aggregate demand of all bidders other than bidder A falls below the total supply of 5. Hence, bidder A "clinches" its first license at this price.[3] The license is "clinched," because the fact that the other bidders' demands have fallen below 5 guarantees that bidder A must win this license.

[3] Technically, the price should be 65 less ϵ, but we ignore this adjustment factor.

At price 75, the aggregate demand of all bidders other than bidder A falls below 4, so A clinches its second license at this price. In addition, the aggregate demand of all bidders other than bidder C falls below 5, so C clinches its first license at this price.

Note that the auction maintains separate counters for each bidder: i.e., for A, the aggregate demand of others falls below 4, while for C, it falls below 5. Bidder A displaced its first bidder, namely D, at price 65, and is displacing its second bidder, namely B, at price 75, while bidder C is displacing its first bidder, again B, at price 75.

The auction terminates at price 85, when bidder A clinches its third license, and bidder C, its second, both displacing bidder D, and total demand meets total supply. In sum, bidder A pays $65 + 75 + 85$ for its three licenses, and bidder C pays $75 + 85$ for its two licenses. As expected, prices on additional licenses are weakly increasing. (In fact, in this example, they are strictly increasing.)

The outcome of the clinching auction in this example (and always; see Proposition 4.1) is efficient (up to $m\epsilon$). In contrast, in this example, a uniform-price auction[4] would *not* have yielded an efficient outcome, as it would have been in bidder A's best interest to decrease its demand to two licenses when the price reached \$75 rather than win all three licenses for \$85 each. Winning only two licenses, A's utility would have been $236 - 2(75) = 86$, whereas winning all three, A's utility would have been $339 - 3(85) = 84$. A uniform-price auction is thus susceptible to **strategic demand reduction**.

[4] A uniform-price auction for multiple copies of a homogeneous good is one that charges the same price for all copies of the good.

4 *The Clinching Auction is EPIC*

To prove the clinching auction is approximately EPIC, we follow the design recipe for EPIC auctions. That is, we first show that the outcome of sincere bidding in the clinching auction is approximately VCG (both allocation and payments; Steps 1 and 2 of the design recipe, respectively). We then show that no inconsistent deviations are preferable to sincere bidding (Step 3).

We first show that any allocation the clinching auction attains is welfare maximizing up to $m\epsilon$ (Step 1).

Proposition 4.1. *Assuming sincere bidding, the clinching auction yields total welfare within $m\epsilon$ of the optimal.*

Proof. Let A denote the highest m marginal values among all bidders,[5] and let B be the marginal values corresponding to the final allocation. Note that regardless of who the precise winners are, payments do not change. Hence, it suffices to show that the difference between the total value of A and B is bounded above by $m\epsilon$.

[5] breaking ties arbitrarily

Call the final round k. Assuming sincere bidding, the set of highest marginal values S_k at round k is a subset of B, because demand may fall below m in the kth round, so that some marginal values in S_{k-1} are allocated. (In the worst case, $S_k = \emptyset$.) Additionally,

- every marginal value in $i \in A \setminus S_k$ is at most $k\epsilon$—otherwise, it would be in S_k; and
- every marginal value in $i \in A \setminus S_k$ is also at least $(k-1)\epsilon$, because bidding is sincere, the auction terminates at round k, and the marginal values in A are the highest m marginal values.

Thus, the maximal difference between the total value of A and B is:

$$\begin{aligned}\sum_{i \in A \setminus B} v_i &\leq \sum_{i \in A \setminus S_k} v_i \\ &\leq (|A| - |S_k|)\,(k\epsilon - (k-1)\epsilon) \\ &= (|A| - |S_k|)\,\epsilon \\ &\leq m\epsilon.\end{aligned}$$

□

En route to showing that the payments in the clinching auction are close to those of VCG (Step 2), we first show that the clinching auction is IR: i.e., no bidder earns negative utility by participating.

Lemma 4.2. *The clinching auction is IR.*

Proof. Assume bidder i wins j copies of the good in the clinching auction, with final price q^*. By the definition of demand sets, $j \leq D_i(q^* - \epsilon)$ iff $\mu_i(j) \geq q^* - \epsilon$. Further, by the design of the payment rule (Equation 2), $q_i(j) \leq q^* - \epsilon$, for all bidders i and all copies of the good j. Therefore, $\mu_i(j) \geq q_i(j)$, for all bidders i and goods j. □

Theorem 4.3. *Given a multiparameter auction setting in which all bidders' have diminishing marginal valuations, the difference in utility earned by a truthful bidder in the VCG auction (assuming others also bid truthfully) and the same bidder bidding sincerely in the clinching auction (assuming others also bid sincerely) is at most $m\epsilon$.*

Proof. By Equation 1, the utility of bidder i in the VCG auction, assuming truthful bidding, is given by:

$$\sum_{j=1}^{x_i} \left(\mu_i(j) - p_i(j)\right). \tag{3}$$

Similarly, by Equation 2, the utility of bidder i in the clinching auction, assuming sincere bidding, is given by:

$$\sum_{j=1}^{y_i} \left(\mu_i(j) - q_i(j)\right). \tag{4}$$

In these expressions, x_i (y_i) is the number of copies of the good that bidder i wins in the VCG (clinching) auction. The proof concentrates on the cases where x_i and y_i differ, i.e., $x_i < y_i$ or $y_i < x_i$.

In particular, we show the following:

1. To the extent that $y_i < x_i$ (i.e., bidder i wins fewer copies in the clinching auction than in VCG), the value added by any term that appears in Equation 3 but not in Equation 4 is bounded above by ϵ, so that in total any missing terms (i.e., missing copies of the good) can forego at most $m\epsilon > 0$ utility for bidder i.

2. To the extent that $y_i > x_i$ (i.e., bidder i wins more copies in the clinching auction than in VCG), the value added by any term that appears in Equation 4 but not in Equation 3 is also bounded above by ϵ, so that in total any missing terms (i.e., missing copies of the good) can forego at most $m\epsilon > 0$ utility for bidder i.

The result then follows from the fact that the clinching auction and VCG are IR, so that all terms in Equations 3 and 4 are non-negative.

Case 1 ($y_i < x_i$) Assume j is s.t. $\mu_i(j) > p_i(j) + \epsilon$. By the definition of $p_i(j)$ (Equation 1), there exists a time t s.t. $\mu_i(j) > \epsilon t$ (i.e., $D_i(\epsilon t) \geq j$) and $\sum_{k \neq i} D_k(\epsilon t) \leq m - j$. At this time t, the jth copy is clinched. Therefore, any terms that appear in Equation 3 but not in Equation 4 can add value at most ϵ.

Case 2 ($x_i < y_i$) Assume j is s.t. $\mu_i(j) - q_i(j) > \epsilon$. By the design of the payment rule for the clinching auction (Equation 2), $q_i(j) + \epsilon \geq p_i(j)$. It follows that $\mu_i(j) > p_i(j)$. But then i wins j in VCG. Therefore, any terms that appear in Equation 4 but not in Equation 3 can add value at most ϵ. □

Having completed Steps 1 and 2 of the EPIC auction design recipe, we have established that sincere bidding in the clinching auction is an EPNE up to $m\epsilon$, among consistent strategies. The remaining piece of this puzzle, then, is to further show that sincere bidding in the clinching auction is an EPNE up to $m\epsilon$, among consistent *and inconsistent* strategies: i.e., that no inconsistent deviations would yield substantially greater utility than sincere bidding. This claim is established in the following theorem.

Theorem 4.4. *The clinching auction is EPIC, up to $m\epsilon$.*

Proof. Assume all bidders except bidder i bid sincerely. Consequently, other bidders' behaviors are not impacted by i's strategy. Moreover, i's payments are dictated entirely by the other bidders' demands, which again, i cannot influence.

Under these circumstances, we argue that i cannot benefit from bidding inconsistently. To bid inconsistently in the clinching auction would be to report false demand sets. But, *given the activity rule*,[6] which ensures that bidders' demands can never increase, all such reports are in fact consistent with some valuation or another. So it is not actually possible to bid inconsistently in the clinching auction.

[6] At long last, we discover the purpose of the activity rule. It rules out inconsistent bidding.

We have already established that bidding sincerely in the clinching auction is an EPNE up to $m\epsilon$, among consistent strategies. As there are no inconsistent strategies, it is likewise EPIC up to $m\epsilon$. □

A Reductions: A Primer

There are three steps in the reduction process, from a known hard (e.g., NP-hard) problem O (for old) to a new problem N (for new), to show N is also hard.

- Pick a known hard problem O.
- Assume a polynomial-time algorithm for the new problem N.
- Derive a polynomial-time algorithm for solving O using the polynomial-time algorithm for solving N as a subroutine.

These three steps leads to a contradiction, as O is known to be hard. Therefore, there can be no polynomial-time algorithm for N (e.g., unless P = NP).

We apply these three steps to reduce the design of a multiparameter DSIC direct auction (i.e., VCG) to the design of a multiparameter EPIC indirect auction. We assume general valuations: i.e., we assume only monotonicity and free disposal.

- Designing a multiparameter DSIC direct auction (i.e., computing a VCG outcome) is a known hard problem, O.
- Assume a polynomial-time algorithm for a multiparameter EPIC indirect auction, N.
- The relevation principle provides a polynomial-time algorithm for solving O, using the polynomial-time algorithm for solving N as a subroutine, because the outcome of the EPIC indirect auction, in which sincere bidding is an EPNE, would yield a DSIC direct auction, in which truthful bidding is a DSE (i.e., a VCG outcome).
- We can thus compute a VCG outcome in polynomial time.

This reasoning lead to a contradiction, as designing a multiparameter DSIC direct auction (i.e., computing a VCG outcome) is known to be NP-hard. Therefore, there can be no polynomial-time algorithm for the design of a multiparameter EPIC indirect auction (unless P = NP) in the general case.

In the special case of diminishing marginal valuations, or unit-demand valuations, where computing the VCG outcome is *not* NP-hard, this reduction merely implies that designing a multiparameter EPIC indirect auction is at least as hard as (i.e., is no easier than)—measured on a complexity theory scale—designing a multiparameter DSIC direct auction, because *if there were a more efficient solution to the indirect auction design problem, it would just as well apply to the direct auction design problem via the revelation principle.*

References

[1] Lawrence M. Ausubel. An efficient ascending-bid auction for multiple objects. *American Economic Review*, 94(5):1452–1475, December 2004.