CS364B: Frontiers in Mechanism Design
Lecture #20: Characterization of Revenue-Maximizing Auctions*

Tim Roughgarden†

March 12, 2014

1 The Story So Far

Recall the scenario studied last lecture:

- A set U of m non-identical items.
- Each bidder $i = 1, 2, \ldots, n$ has an additive valuation drawn from a prior distribution F_i. Recall this means that the valuation v_i is an m-vector, with $v_i(S) = \sum_{j \in S} v_{ij}$. The distribution F_i has a finite support $\mathcal{V}_i$ and the probabilities $\{f_i(v_i)\}_{v_i \in \mathcal{V}_i}$ are provided explicitly as input. The v_{ij}'s can be correlated across items j for a fixed bidder i, but are independent across bidders.

Recall what we are shooting for: a multi-parameter analog of Myerson's theory of revenue-maximizing optimal auctions. Single-parameter optimal mechanisms are virtual welfare maximizers. We've seen that multi-parameter optimal mechanisms are generally randomized. It would seem that the coolest statement that could be true is: revenue-maximizing optimal auctions for multi-parameter problems are always distributions over virtual welfare maximizers. This lecture identifies conditions under which this coolest-possible statement is actually true.

Recall that the reduced form $(\mathbf{y}, \mathbf{q})$ induced by a direct-revelation mechanism $(\mathbf{x}, \mathbf{p})$ consists of an *interim allocation rule* y defined by

$$\begin{aligned} y_{ij}(v_i) &:= \mathbf{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}}[x_{ij}(v_i, \mathbf{v}_{-i})] \\ &= \sum_{\mathbf{v}_{-i} \in \mathcal{V}_{-i}} \mathbf{f}_{-i}(\mathbf{v}_{-i}) x_{ij}(v_i, \mathbf{v}_{-i}) \end{aligned} \quad (1)$$

†Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. Email: tim@cs.stanford.edu.

for every bidder i, item j, and reported valuation v_i, and an *interim payment rule* given by

$$q_i(v_i) := \mathbf{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}}[p_i(v_i, \mathbf{v}_{-i})] = \sum_{\mathbf{v}_{-i} \in \mathcal{V}_{-i}} \mathbf{f}_{-i}(\mathbf{v}_{-i}) p_i(v_i, \mathbf{v}_{-i}) \tag{2}$$

for every bidder i and reported valuation v_i.

The culmination of last lecture was an explicit description as a linear program of the reduced forms of BIC and IIR mechanisms:

$$\max \sum_{i=1}^{n} f(v_i) q_i(v_i)$$

subject to

$$\sum_{j \in U} v_{ij} y_{ij}(v_i) - q_i(v_i) \geq \sum_{j \in U} v_{ij} y_{ij}(v_i{}') - q_i(v_i{}') \qquad \forall i \text{ and } v_i, v_i{}' \in V_i \tag{3}$$

$$\sum_{j \in U} v_{ij} y_{ij}(v_i) - q_i(v_i) \geq 0 \qquad \forall i \text{ and } v_i \in V_i \tag{4}$$

$$\sum_{i=1}^{n} \sum_{v_i \in \mathcal{V}_i} F_i(v_i) y_{ij}(v_i) \leq 1 - \prod_{i=1}^{n} \left(1 - \sum_{v_i \in \mathcal{V}_i} F_i(v_i)\right) \qquad \forall j \in U \text{ and } S_1 \subseteq V_1, \ldots, S_n \subseteq V_n. \tag{5}$$

Observe that the number of BIC constraints and IIR constraints is polynomial in the total number $\sum_{i=1}^{n} |\mathcal{V}_i|$ of possible types (and n), while the number of feasibility constraints is exponential in this quantity.

It will also be helpful to have a geometric interpretation of this result. For a given prior $\mathbf{F}$, we can think of passing from a mechanism $(\mathbf{x}, \mathbf{p})$ to its reduced form $(\mathbf{y}, \mathbf{q})$ as a "projection" via the linear operators (1) and (2). This maps the original high-dimensional linear program from Lecture #18 — with decision variables corresponding to the ex post allocation and payment rules — to the relatively low-dimensional linear program above (Figure 1). A linear projection of a polytope — i.e., an intersection of halfspaces that is bounded — is again a polytope (Exercise), and hence can be described as the solutions to a finite number of linear inequalities. The linear program above is an explicit description of these inequalities. This projection step from the set of mechanisms to the set of reduced forms is not generally injective — many different ex post allocation and payment rules yield the same reduced form.

2 Next Steps

Satisfying as it may be, the explicit description above raises more problems than it solves.

1. Does this explicit description helps us understand the structure of optimal mechanisms?

2. Can we optimize over this explicit description — which still has exponentially many constraints — in time polynomial in n and $\sum_{i=1}^{n} |\mathcal{V}_i|$?

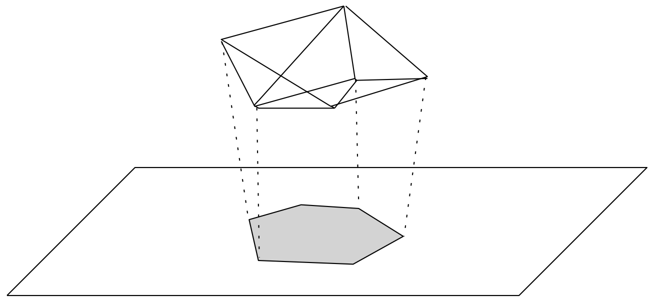

Figure 1: A geometric interpretation of the projection.

3. Given a reduced form, can we recover a direct-revelation mechanism that induces it in polynomial time?

4. Can we say anything beyond bidders with additive valuations?

In this lecture we focus on the first, structural question. The best way to do this is to tackle the fourth question simultaneously. There are corresponding tractability results for the second and third questions [2]; other than the brief discussion in Section 5, we won't have time to cover them here.

For concreteness, we focus on bidders with unit-demand valuations.[1] Recall how this compares with additive valuations: a bidder's valuation is still described fully by its valuations $v_{i1}, \ldots, v_{im}$ for singletons, but now $v_i(S) = \max_{j \in S} v_{ij}$ rather than $v_i(S) = \sum_{j \in S} v_{ij}$. That is, a bidder effectively throws away all of its items except for its favorite. For this reason, we can restrict attention to mechanisms that always give at most one item to each bidder. With this restriction, feasible allocations correspond to bipartite matchings between the bidders and the items.

With the original massive linear program from Lecture #18, it's easy to add constraints corresponding to unit-demand bidders: in addition to the original feasibility constraints

$$\sum_{i=1}^{n} x_{ij}(\mathbf{v}) \le 1 \tag{6}$$

for every $j \in U$ and $\mathbf{v} \in \mathcal{V}$, we just add the constraints

$$\sum_{j=1}^{n} x_{ij}(\mathbf{v}) \le 1 \tag{7}$$

for every bidder $i \in U$ and $\mathbf{v} \in \mathcal{V}$. Border's theorem describes what the inequalities (6) look like after projecting to the lower-dimensional space of reduced forms; what happens when we also have the constraints (7)?

We won't answer this question as explicitly as we did with Border's theorem, but we'll be able to say enough to meet our original goal of deriving a structural characterization of optimal mechanisms.

[1]The structural results in Sections 3 and 4 hold more generally, with the same proofs, for an arbitrary linear system of extra feasibility constraints.

3 Main Structural Theorem

Define a *virtual welfare maximizer (VWM)* as an (ex post) allocation rule $\mathbf{x}$ that, for some functions $\varphi_{ij} : \mathcal{V}_i \to \mathcal{R}$ for every bidder i and item j, chooses for each valuation profile $\mathbf{v}$ the feasible allocation that maximizes the virtual welfare

$$\sum_{i=1}^{n} \sum_{j \in U} \varphi_{ij}(v_i) x_{ij}(\mathbf{v}).$$

For example, for bidders with unit-demand valuations, a VWM chooses an allocation corresponding to a maximum-weight bipartite matching, where the weight of edge (i, j) is $\varphi_{ij}(v_i)$. If $\mathbf{y}$ is an interim allocation rule induced by a VWM, then we call $\mathbf{y}$ an *interim VWM (iVWM)*.

When there is only one item, Myerson's theory (CS364A, Lecture #5) implies that every optimal auction is a virtual welfare maximizer; in this sense, revenue-maximization reduces to welfare-maximization for single-parameter problems. The next results, due to Cai, Daskalakis, and Weinberg [2], outline the appropriate generalization to certain multi-parameter problems.[2] Let X denote the polytope of feasible ex post allocation rules $\mathbf{x}$ and Y the corresponding polytope of interim allocation rules $\mathbf{y}$ (cf., Figure 1).

Theorem 3.1 ([2]) *Every vertex*[3] *of Y is an iVWM.*

We give a proof of Theorem 3.1 in the next section. We next explain how it yields a multi-parameter analog of Myerson's characterization of optimal mechanisms.

Corollary 3.2 *For every allocation rule $\mathbf{x} \in X$, there is an allocation rule $\mathbf{x}' \in X$ such that $\mathbf{x}'$ is a probability distribution over VWMs and such that $\mathbf{x}$ and $\mathbf{x}'$ induce the same interim allocation rule.*

Proof: Consider any allocation rule $\mathbf{x} \in X$, and let $\mathbf{y} \in Y$ be the corresponding interim allocation rule. The vector $\mathbf{y}$ can be written as a convex combination of the vertices of Y. Moreover, the number of vertices needed is most one more than the (polynomial) number of dimensions.[4], By Theorem 3.1, this convex combination can be interpreted as a probability distribution over iVWMs — say, the iVWMs $\mathbf{y}_1, \ldots, \mathbf{y}_\ell$ with probabilities $\lambda_1, \ldots, \lambda_\ell$. Each iVWM $\mathbf{y}_h$ is the interim rule of some VWM $\mathbf{x}_h$. Since the projection from X to Y is linear, setting $\mathbf{x}'$ to be the probability distribution over the $\mathbf{x}_h$'s (with the λ_h's as probabilities) yields a distribution over VWMs with the same interim rule as $\mathbf{x}$. ■

Corollary 3.3 *There is a revenue-maximizing auction with an allocation rule that is a distribution over VWMs.*

[2]Related results appear in Alaei et al. [1].

[3]The *vertex* v of a polytope P is, intuitively, a "corner." One of several equivalent definitions is that v cannot be written as a non-trivial convex combination of other points of P. Another is that there exists a halfspace H such that $H \cap P = \{v\}$.

[4]This is an elementary property of polytopes; see e.g. [3].

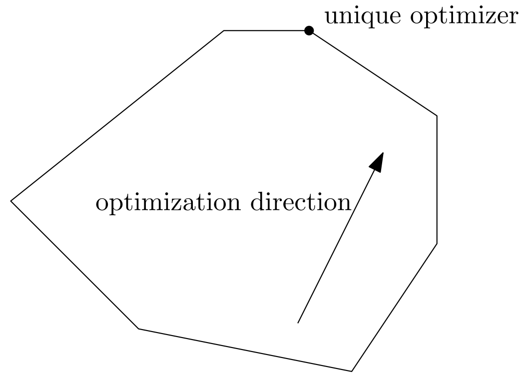

Figure 2: The marked vertex is the unique optimizer in the given direction.

For example, with unit-demand bidders, there is an optimal auction that picks random edge weights from some distribution and then allocates to a maximum-weight matching (and charges suitable prices).

Proof of Corollary 3.3: Let $(\mathbf{x}^*, \mathbf{p}^*)$ be a revenue-maximizing auction. By Corollary 3.2, there is an allocation rule $\mathbf{x}'$ that is a distribution over VWMs and that induces the same interim rule as $\mathbf{x}^*$. Since the incentive constraints and expected revenue depend only on the induced interim rules, $(\mathbf{x}', \mathbf{p}^*)$ is also a revenue-maximizing mechanism. ■

Conceptually, the only drawbacks of Corollary 3.3 relative to Myerson's single-parameter theory are: (i) the virtual valuation functions φ_{ij} are randomized rather than deterministic; and (ii) the virtual valuation functions have no closed form, and instead are the output of a linear program. The examples in Lecture #18 provided early warnings that these two drawbacks would be necessary, and in this this sense Corollary 3.3 is a best-case scenario. Conceptually, revenue-maximization continues to reduce to welfare-maximization.

4 Proof of Theorem 3.1

Consider a vertex $\mathbf{y}^*$ of the polytope Y. An intuitive and elementary fact about polytopes is that a point $\hat{\mathbf{y}}$ in a polytope is a vertex if and only if there exists an linear objective function $\mathbf{w}$ (i.e., a direction) such that $\hat{\mathbf{y}}$ is the unique maximizer of $\mathbf{w}^T\mathbf{y}$ over all feasible points $\mathbf{y}$ (Figure 2).[5] The idea is now to show that when an interim allocation rule is optimizing some linear function, then there is a corresponding ex post rule (inducing the interim rule) that is optimizing some linear function, and this linear optimization corresponds to virtual welfare maximization with respect to the appropriate virtual valuation functions.

[5] One direction is clear: if $\mathbf{y}$ is not a vertex, then $\mathbf{y}$ can be written as a non-trivial convex combination of two other feasible points, so there is no linear objective for which $\mathbf{y}$ is the unique maximum. For the converse, let H be a halfspace whose intersection with the polytope is $\{\mathbf{y}\}$, and choose $\mathbf{w}$ as -1 times the normal vector of H.

Let $\mathbf{w}$ be a linear objective function for which $\mathbf{y}^*$ is the unique maximum over Y. Recalling the decision variables that define Y, this objective function is

$$\max \sum_{i=1}^{n} \sum_{v_i \in \mathcal{V}_i} \sum_{j \in U} w_{ij}(v_i) y_i(v_i).$$

Define virtual valuation functions by

$$\varphi_{ij}(v_i) = \frac{w_{ij}(v_i)}{f_i(v_i)}. \tag{8}$$

We next observe that the expected virtual welfare obtained by an (ex post) allocation rule is a function only of its interim allocation rule. Precisely, for every allocation rule $\mathbf{x}$ with interim rule $\mathbf{y}$, its expected virtual welfare is

$$\sum_{\mathbf{v} \in \mathcal{V}} f(\mathbf{v}) \sum_{i=1}^{n} \sum_{j \in U} \varphi_{ij}(v_i) x_{ij}(\mathbf{v}) = \sum_{i=1}^{n} \sum_{v_i \in \mathcal{V}_i} f_i(v_i) \sum_{j \in U} \varphi_{ij}(v_i) y_j(v_i) = \sum_{i=1}^{n} \sum_{v_i \in \mathcal{V}_i} \sum_{j \in U} w_{ij}(v_i) y_j(v_i), \tag{9}$$

where the first equation follows from linearity of expectation and the second from the definition (8) of the virtual valuations.

It is clear which allocation rule $\mathbf{x}^*$ maximizes the left-hand side of (9): for each $\mathbf{v} \in \mathcal{V}$, set $\mathbf{x}^*(\mathbf{v})$ equal to the feasible allocation that maximizes the virtual welfare $\sum_{i=1}^{n} \sum_{j \in U} \varphi_{ij}(v_i) x_{ij}(\mathbf{v})$. By definition, $\mathbf{x}^*$ is a VWM for the virtual valuations defined in (8). Because $\mathbf{y}^*$ is the unique maximizer of the right-hand side of (9), the interim allocation rule of $\mathbf{x}^*$ must be $\mathbf{y}^*$. Since $\mathbf{y}^*$ was an arbitrary vertex of Y, the proof is complete.

5 Computational Considerations

Now that we understand the structure of optimal mechanisms, we would like two types of computational tractability results. First, given a description of the prior distributions, we would like to compute a revenue-maximizing mechanism in time polynomial in $\sum_{i=1}^{n} |\mathcal{V}_i|$. Second, we would like the mechanism itself to run in polynomial time.

Efficient solutions to the following computational questions lead to the above tractability results.

1. Given a vector $\mathbf{y}$, can we efficiently check membership in the polytope Y?
2. Can we efficiently optimize a linear function (like expected revenue) over the polytope of reduced forms?
3. Given a feasible $\mathbf{y} \in Y$ (like a revenue-maximizing reduced form), can we efficiently reconstruct an ex post allocation rule $\mathbf{x}$ with interim rule $\mathbf{y}$? The goal here is a constructive version of Corollary 3.3.

We discuss only the first question; "ellipsoid magic" effectively reduces the second and third questions to the first. In the spirit of Farkas's Lemma and strong duality, let's ask: what would convince you that a given vector $\mathbf{y}$ does *not* belong to the polytope Y of interim allocation rules? Suppose for a linear objective $\{w_{ij}(v_i)\}_{i,j,v_i}\}$ and corresponding virtual valuation functions φ_{ij} defined by (8), the right-hand side of (9) is strictly larger than the maximum-possible virtual welfare of an ex post allocation rule. Then, by the equality (9), $\mathbf{y}$ cannot be a realizable interim allocation rule. Strong linear programming duality can be shown to imply the converse: if for every $\mathbf{w}$, the right-hand side of (9) is at most the maximum virtual welfare with respect to the virtual valuation functions (8), then $\mathbf{y} \in Y$. This characterization, together with relatively standard "ellipsoid magic" (cf., Lectures #6 and #9), implies that the first computational question above reduces in polynomial time to the following subroutine:

(*) given the alleged interim allocation rule $\mathbf{y}$, is there an objective function $\mathbf{w}$ and a number W such that

$$\sum_{i=1}^{n} \sum_{v_i \in \mathcal{V}_i} \sum_{j \in U} w_{ij}(v_i) y_j(v_i) \geq W \tag{10}$$

while

$$\sum_{\mathbf{v} \in \mathcal{V}} \mathbf{f}(\mathbf{v}) \sum_{i=1}^{n} \sum_{j \in U} \varphi_{ij}(v_i) x_{ij}(\mathbf{v}) < W \tag{11}$$

for every ex post allocation rule $\mathbf{x}$, where the φ_{ij}'s are defined as in (8)?

Observe that the problem (*) is a linear program with polynomially many decision variables ($\mathbf{w}$ and W) and a huge number of constraints. Recalling the ellipsoid method (Lectures #6 and #9), we can solve this problem if we can design a polynomial-time separation oracle. The input to this oracle is a fixed $\mathbf{w}$ and W (and $\mathbf{y}$, which was fixed beforehand). The constraint (10) can be checked directly. Checking the constraints (11) reduces to checking it only for the ex post allocation rule that maximizes the left-hand side. Doing this exactly would seem to require enumerating over the exponentially many valuation profiles $\mathbf{v} \in \mathcal{V}$, and computing the virtual welfare-maximizing allocation for each. Randomly sampling a polynomial number of valuation profiles and maximizing the virtual welfare for each yields a good enough approximation to push the "ellipsoid magic" through; see [2] for the non-trivial arguments. This shows that approximate revenue-maximization reduces to welfare-maximization also in a computational sense: if welfare-maximization can be solved in polynomial time (e.g., for unit-demand bidders), then there is a polynomial-time algorithm for computing an approximately revenue-maximizating (and approximately BIC) auction for given prior distributions, and this optimal auction runs in polynomial time.

References

[1] Saeed Alaei, MohammadTaghi Hajiaghayi, and Vahid Liaghat. Online prophet-inequality matching with applications to ad allocation. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, pages 18–35, 2012.

[2] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *Proceedings of the 53rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 130–139, 2012.

[3] G. M. Ziegler. *Lectures on Polytopes*. Springer-Verlag, 1995.