CS364B: Frontiers in Mechanism Design
Lecture #19: Interim Rules and Border's Theorem*

Tim Roughgarden†

March 12, 2014

1 The Big Picture

In this lecture we continue our study of revenue-maximization in multi-parameter problems. Unlike Lectures #1–17, where we focused entirely on welfare maximization, here we strive to maximize the sum of the payments from the bidders to the mechanism. Since there is no "always optimal" mechanism, akin to the VCG mechanism for welfare-maximization, we compare the performance of different auctions using a prior distribution over valuations. Last lecture, we recalled Myerson's well-understood and satisfying single-parameter theory: maximizing expected revenue reduces to maximizing virtual welfare, where the virtual valuation of a bidder is a relatively simple formula of its valuation and the prior distribution. This reduction is interesting both conceptually and computationally. First, it tells us what optimal auctions looks like — they are virtual welfare maximizers. They are DSIC — even though we optimize over the richer space of BIC mechanisms — and with regular distributions, they are deterministic. Second, it implies that in every setting where welfare-maximization is computational tractable, revenue-maximization with respect to a prior is also tractable.

The goal of this lecture and the next is to develop a multi-parameter analog of Myerson's theorem. Even though Myerson's paper is almost 35 years old [2], some of the most interesting progress on this question is from just the last year or two.

Last lecture, we say that revenue-maximizing auctions are more complex in multi-parameter settings than in single-parameter ones. This is true even with just one buyer — where with one good, the optimal selling procedure is a take-it-or-leave-it offer at a monopoly price. With only two items and a buyer with an additive valuation drawn from extremely simple prior distributions, the optimal auction format varies significantly with the details of the prior and need not be deterministic. This means the optimal auctions need not be (deterministic) virtual welfare maximizers. It also seems unrealistic to expect tractable closed-formulas for

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

the revenue-maximizing allocation rule. This suggests tacking the problem with heavier machinery than in the single-parameter case, and we're turning to linear programming theory for this purpose.

2 An Exponential-Size Linear Program

Recall the current setup:

- 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 $v_i(S) = \sum_{j \in S} v_{ij}$.[1] 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.

From here on out, we'll assume that the F_i's are independent.[2] For a fixed bidder i, we'll never need to assume that the valuations v_{ij} for different items j are independent.

We concluded last lecture with the following linear programming formulation of the revenue-maximizing BIC and IIR mechanism:

$$\sum_{\mathbf{v} \in \mathcal{V}} \mathbf{F}(\mathbf{v}) \sum_{i=1}^{n} p_i(\mathbf{v}) \tag{1}$$

subject to

$$\sum_{\mathbf{v}_{-i} \in \mathcal{V}_{-i}} \mathbf{F}_{-i}(\mathbf{v}_{-i}) \left(\sum_{j \in U} v_{ij} x_{ij}(\mathbf{v}) - p_i(\mathbf{v}) \right) \geq \sum_{\mathbf{v}_{-i} \in \mathcal{V}_{-i}} \mathbf{F}_{-i}(\mathbf{v}_{-i}) \left(\sum_{j \in U} v_{ij} x_{ij}(v_i', \mathbf{v}_{-i}) - p_i(v_i', \mathbf{v}_{-i}) \right) \tag{2}$$

for every bidder i, true valuation v_i, and reported valuation v_i';

$$\sum_{\mathbf{v}_{-i} \in \mathcal{V}_{-i}} \mathbf{F}_{-i}(\mathbf{v}_{-i}) \left(\sum_{j \in U} v_{ij} x_{ij}(\mathbf{v}) - p_i(\mathbf{v}) \right) \geq 0 \tag{3}$$

for every bidder i and $v_i \in V_i$;

$$\sum_{i=1}^{n} x_{ij}(\mathbf{v}) \leq 1 \tag{4}$$

for every $j \in U$ and $\mathbf{v} \in \mathcal{V}$; and nonnegativity constraints. In this linear program, the variable $x_{ij}(\mathbf{v})$ indicates the probability (over the mechanism's coin flips) that bidder i receives the

[1] We'll go beyond the very simple class of additive valuations in the next lecture, but as we saw last lecture these valuations already pose interesting challenges.

[2] Some of what we'll see holds also for correlated distributions — a good exercise is to think through what extends easily to correlated distributions and what does not.

item j in the valuation profile $\mathbf{v}$. The variable $p_i(\mathbf{v})$ indicates bidder i's expected payment (over the mechanism's coin flips) in the valuation profile $\mathbf{v}$. Because bidders and the seller are risk-neutral, and bidders have additive valuations, these quantities are enough to express everyone's objective function. There is a natural correspondence between feasible solutions to this linear program and direct-revelation mechanisms that are BIC and IIR.

But what use is this linear programming formulation? First, observe that the number of decision variables and the number of constraints of the form (4) are polynomial in the number of valuation profiles $\mathbf{v}$, which is exponential in the number of players n. This renders the linear program (1)–(4) useless from a computational perspective for all but the tiniest problems. The second issue is that it's not clear what we've learned, conceptually or structurally, about revenue-maximizing auctions by characterizing them as optimal solutions to certain massive linear programs. In this lecture and the next we address both of these issues. Guided by the first criticism, we next strive for a linear programming formulation of the revenue-maximizing auction with far fewer constraints — scaling polynomially in the number n of bidders, rather than exponentially.

3 Interim Rules and Feasible Reduced Forms

The next definition is extremely important. Let $(\mathbf{x}, \mathbf{p})$ be a direct-revelation mechanism. The induced *interim allocation rule* y is defined by

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

for every bidder i, item j, and reported valuation v_i. The interim allocation rule is probability that bidder i receives item j with the report v_i, over the randomness in the allocation rule $\mathbf{x}$ *and in the valuations* $\mathbf{v}_{-i}$ *of the other bidders*. Even if $\mathbf{x}$ is a deterministic (i.e., 0-1) allocation rule, the corresponding interim allocation rule is generally not 0-1.

Similarly, the induced *interim payment rule* is given by

$$q_i(v_i) := \mathbf{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}}[p_i(v_i, \mathbf{v}_{-i})] .$$

The pair $(\mathbf{y}, \mathbf{q})$ is called the *reduced form* of $(\mathbf{x}, \mathbf{p})$.[3] We sometimes call $\mathbf{x}$ and $\mathbf{p}$ *ex post* rules for emphasis.[4]

Two observations motivate concentrating on reduced forms $(\mathbf{y}, \mathbf{q})$ rather than direct-revelation mechanisms $(\mathbf{x}, \mathbf{p})$. The first is that interim rules take as input only a single reported valuation, as opposed to a full valuation profile. Thus, only $(m+1)\sum_{i=1}^{n} |V_i|$

[3]We first encountered interim allocation and payment rules back in Lecture #12, when we discussed the characterization of BIC mechanisms in single-parameter environments (the BIC version of Myerson's Lemma).

[4]In this context, "ex post" means after the resolution of all uncertainty; "ex interim" means that, from bidder i's perspective, v_i is known with certainty while $\mathbf{v}_{-i}$ remains uncertain.

numbers are required to specify a reduced form, and this is growing linearly with the number of players. Second, despite its compressed size, the reduced form of a mechanism is enough to specify the seller's revenue and bidders' utilities, and hence the BIC and IIR constraints.

More precisely, consider the following linear program, with decision variables corresponding to the interim allocation and payment rules:

$$\max \sum_{i=1}^{n} f(v_i) q_i(v_i) \tag{5}$$

subject to

$$\sum_{j \in U} v_{ij} y_{ij}(v_i) - q_i(v_i) \geq \sum_{j \in U} v_{ij} y_{ij}({v_i}') - q_i({v_i}') \tag{6}$$

for all bidders i, valuations v_i, and reported valuations v_i; and

$$\sum_{j \in U} v_{ij} y_{ij}(v_i) - q_i(v_i) \geq 0 \tag{7}$$

for all bidders i and valuations v_i. The objective function (5) and the constraints (6)–(7) are simply (1)–(3), re-expressed in the vocabulary of the reduced form. The number of variables in (5)–(6) scales linearly with the number n of players, rather than exponentially as in (1)–(4).

How can we translate the feasibility constraints (4) of the original linear program into our new, more economical vocabulary? Call an alleged interim allocation rule $\mathbf{y}$ *feasible* if there exists an (ex post) allocation rule $\mathbf{x}$ satisfying the feasibility constraints (4). Since each item j is allocated to at most one bidder in every valuation profile, we certainly have the following necessary condition for $\mathbf{y}$ to be feasible:

$$\sum_{i=1}^{n} \underbrace{\sum_{v_i \in \mathcal{V}_i} f_i(v_i) y_{ij}(v_i)}_{\mathbf{Pr}[i \text{ wins } j]} \leq 1. \tag{8}$$

Could this also be a sufficient condition? That is, is every alleged interim allocation rule $\mathbf{y}$ that satisfies (8) induced by a bone fide (ex post) allocation rule $\mathbf{x}$?

To get a better feel for the issue, let's consider a couple of examples.

Example 3.1 Suppose there are $n = 2$ bidders, and there is $m = 1$ item. Assume that v_1, v_2 are independent and each is equally likely to be 1 or 2.

Consider the alleged interim allocation rule given by

$$y_1(1) = \tfrac{1}{2}, y_1(2) = \tfrac{7}{8}, y_2(1) = \tfrac{1}{8}, \text{ and } y_2(2) = \tfrac{1}{2}. \tag{9}$$

Since $f_i(v) = \frac{1}{2}$ for all $i = 1, 2$ and $v = 1, 2$, $\mathbf{y}$ satisfies the necessary condition 8. Can you find an (ex post) allocation rule $\mathbf{x}$ that induces the interim rule $\mathbf{y}$? Note that trying to find an $\mathbf{x}$ is much like solving a Sudoko or KenKen puzzle — the goal is to fill in the table entries

(v_1, v_2)	$x_1(v_1, v_2)$	$x_2(v_1, v_2)$
$(1,1)$		
$(1,2)$		
$(2,1)$		
$(2,2)$		

Table 1

in Table 1 so that each row sums to at most 1 (for feasibility) and that the constraints (9) are satisfied. For example, the average of the top two entries in the first column of Table 1 should be $y_1(1) = \frac{1}{2}$. In this example, there are a number of such solutions; one is shown in Table 2. Thus, $\mathbf{y}$ is feasible.

Example 3.2 Suppose we change the alleged interim allocation rule to

$$y_1(1) = \tfrac{1}{4}, y_1(2) = \tfrac{7}{8}, y_2(1) = \tfrac{1}{8}, \text{ and } y_2(2) = \tfrac{3}{4}.$$

The necessary condition (8) remains satisfied. Now, however, $\mathbf{y}$ is not feasible. One way to see this is to note that since $y_1(2) = \frac{7}{8}$ and $x_1(2,2) \geq \frac{3}{4}$ and hence $x_2(2,2) \leq \frac{1}{4}$. Similarly, $y_2(2) = \frac{3}{4}$ implies that $x_2(2,2) \geq \frac{1}{2}$, a contradictory constraint.

The first point of Examples 3.1 and 3.2 is that it is not trivial to check whether or not a given $\mathbf{y}$ is feasible — it corresponds to solving a big system of linear equations. The second point is that (8) is not a sufficient condition for feasibility. In hindsight, trying to summarize the exponentially many ex post feasibility constraints (4) with a single interim constraint (8) seems naive. Are there some additional constraints — possibly an exponential number — that we can add to (5)–(7) so that the feasible solutions $(\mathbf{y}, \mathbf{q})$ correspond precisely the reduced forms of feasible (and BIC and IIR) mechanisms $(\mathbf{x}, \mathbf{p})$?

4 Border's Theorem

The last goal of this lecture is to give an explicit system of linear constraints on the variables $\mathbf{y}$ so that the feasible solutions to this system correspond precisely to the feasible interim allocation rules. This result is a special case of *Border's Theorem* [1]. Next lecture, we derive

(v_1, v_2)	$x_1(v_1, v_2)$	$x_2(v_1, v_2)$
$(1,1)$	1	0
$(1,2)$	0	1
$(2,1)$	3/4	1/4
$(2,2)$	1	0

Table 2: One solution for the allocation rule.

interesting conceptual computational consequences of this result, and also extend it to more general settings.

For the rest of this lecture, we assume for notational convenience that the valuation sets $V_1, \ldots, V_n$ are disjoint. This is without loss of generality, since we can simply "tag" each valuation $v_i \in V_i$ with the "name" i (i.e., view each $v_i \in \mathcal{V}_i$ as the set $\{v_i, i\}$).

We next present a condition that is obviously satisfied by every feasible interim allocation rule $\mathbf{y}$. Let $\mathbf{x}$ be a feasible (ex post) allocation rule and $\mathbf{y}$ the induced interim rule. Fix an item j, and for each bidder i a set $S_i \subseteq V_i$ of valuations. Call the valuations $\cup_{i=1}^n S_i$ the *distinguished* valuations. Consider first the probability, over the random valuation profile $\mathbf{v} \sim \mathbf{F}$ and any coin flips of the allocation rule $\mathbf{x}$, that the winner of item j has a distinguished valuation. By linearity of expectations, this probability can be expressed in terms of the interim allocation rule $\mathbf{y}$:

$$\sum_{i=1}^n \sum_{v_i \in \mathcal{V}_i} F_i(v_i) y_{ij}(v_i). \tag{10}$$

The expression (10) is linear in the $y_{ij}(v_i)$'s.

The second quantity we study is the probability, over $\mathbf{v} \sim \mathbf{F}$, that there is a bidder with a distinguished type. This has nothing to do with the allocation rule, and is a function of the prior $\mathbf{F}$ only:

$$1 - \prod_{i=1}^n \left(1 - \sum_{v_i \in \mathcal{V}_i} F_i(v_i)\right). \tag{11}$$

Since there can only be a winner with a distinguished type is there is a bidder with a distinguished type, the quantity in (10) can only be less than (11). Border's theorem asserts that these conditions, ranging over all choices of $S_1 \subseteq V_1, \ldots, S_n \subseteq V_n$, are also sufficient for the feasibility of an interim allocation rule $\mathbf{y}$.

Theorem 4.1 (Border's theorem [1]) *A vector $\mathbf{y}$ is feasible if and only if for every item $j \in U$ and every choice $S_1 \subseteq V_1, \ldots, S_n \subseteq V_n$ of distinguished types,*

$$\sum_{i=1}^n \sum_{v_i \in \mathcal{V}_i} F_i(v_i) y_{ij}(v_i) \leq 1 - \prod_{i=1}^n \left(1 - \sum_{v_i \in \mathcal{V}_i} F_i(v_i)\right). \tag{12}$$

Proof: We have already argued the "only if" direction, and now prove the converse. The proof is by the max-flow/min-cut theorem — given the statement of the theorem and this hint, the proof writes itself.

Suppose the vector $\mathbf{y}$ satisfies (12) for every $j \in U$ and $S_1 \subseteq V_1, \ldots, S_n \subseteq V_n$. Fix an item $j \in U$. Form a four-layer s-t directed flow network G as follows (Figure 1(a)). The first layer is the source s, the last the sink t. In the second layer X, vertices correspond to valuations profiles $\mathbf{v}$. We abuse notation and refer to nodes as X by the corresponding valuation profile. There is an arc $(s, \mathbf{v})$ for every $\mathbf{v} \in A$, with capacity $\mathbf{F}(\mathbf{v})$. Because $\mathbf{F}$ is a probability distribution, the total capacity of these edges is 1.

In the third layer Y, vertices correspond to winner-valuation pairs; there is also one additional "no winner" vertex. We use (i, v_i) to denote the vertex representing the event

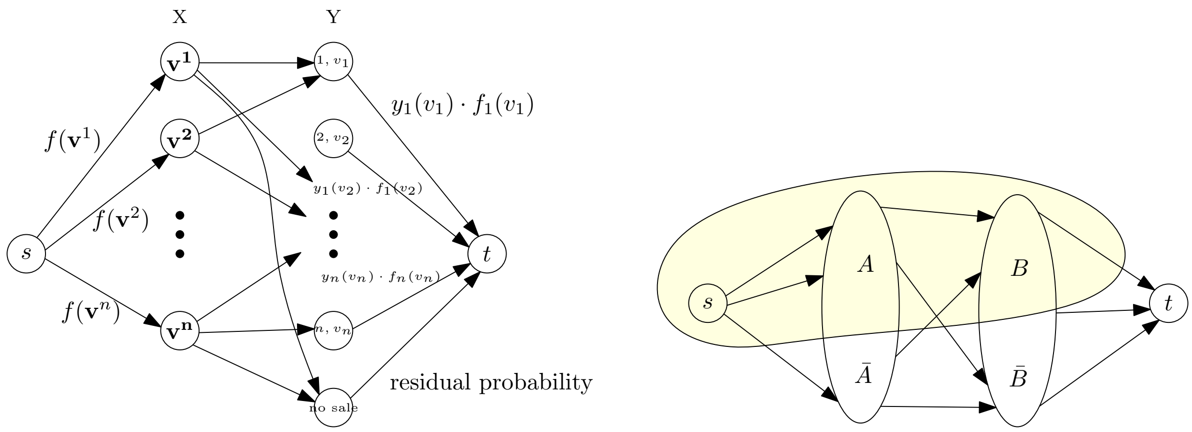

(a) (b)

Figure 1: The max-flow/min-cut proof of Border's theorem.

that bidder i wins the item j and also has type v_i. For each i and $v_i \in V_i$, there is an arc $((i, v_i), t)$ with capacity $F_i(v_i)y_{ij}(v_i)$. There is also an arc from the "no winner" vertex to t, with capacity $1 - \sum_{i=1}^{n} \sum_{v_i \in V_i} F_i(v_i)y_{ij}(v_i)\}$.[5]

Finally, each vertex $\mathbf{v} \in X$ has $n + 1$ outgoing arcs, all with infinite capacity, to the vertices $(1, v_1), (2, v_2), \ldots, (n, v_n)$ of Y and also to the "no winner vertex."

By construction, s-t flows of G with value 1 correspond to ex post allocation rules $\mathbf{x}$ with induced interim allocation rule $\mathbf{y}$, with $x_i(\mathbf{v})$ equal to the amount of flow on the arc $(\mathbf{v}, (i, v_i))$ times $1/\mathbf{F}(\mathbf{v})$. Flows of value 1 must saturate all the arcs incident to t, which is equivalent to having induced interim allocation rule $\mathbf{y}$.

To show that there exists a flow with value 1, it suffices to show that every s-t cut has value at least 1 (by the max-flow/min-cut theorem). So fix an s-t cut. Let this cut include the vertices A from X and B from Y (Figure 1(b)). For each bidder i, define $S_i \subseteq V_i$ as the possible valuations of i that are *not* represented

To show that there exists a flow with value 1, it suffices to show that every s-t cut has value at least 1 (by the max-flow/min-cut theorem). So fix an s-t cut. Let this cut include the vertices A from X and B from Y. Note that all arcs from s to $X \setminus A$ and from B to t are cut (Figure 1(b)). For each bidder i, define $S_i \subseteq V_i$ as the possible valuations of i that are *not* represented amongst the valuation profiles in A. Then, for every valuation profile $\mathbf{v}$ containing at least one distinguished type, the arc $(s, \mathbf{v})$ is cut. The total capacity of these arcs is the right-hand side (11) of Border's condition.

Next, we can assume that every vertex of the form (i, v_i) with $v_i \notin S_i$ is in B, since otherwise an (infinite-capacity) arc from A to $Y \setminus B$ is cut. Similarly, unless $A = \emptyset$ — in which case the cut has value at least 1 and we're done — we can assume that the "no winner" vertex lies in B. Thus, the only edges of the form $((i, v_i), t)$ that are not cut involve a distinguished type $v_i \in S_i$. It follows the total capacity of the cut edges incident to t is at

[5]If $\sum_{i=1}^{n} \sum_{v_i \in V_i} F_i(v_i)y_{ij}(v_i) > 1$, then $\mathbf{y}$ is clearly infeasible (recall (8)). It would also violated Border's condition for the choice $S_i = V_i$ for all i.

least 1 minus the left-hand size (10) of Border's condition. Given our assumption that (10) is at most (11), this s-t cut has value at least 1. This completes the proof of Border's theorem. ■

Border's theorem yields an explict description as a linear program of the reduced forms of BIC, IIR, and feasible mechanisms. To review, this linear program is

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

subject to

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

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

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

In the next lecture we study the conceptual and computational consequences of this linear program, and also generalize it beyond additive valutions.

References

[1] K. C. Border. Implementation of reduced form auctions: A geometric approach. *Econometrica*, 59(4):1175–1187, 1991.

[2] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.