

CPSC 536F: Algorithmic Game Theory

Lecture: Smooth Games and the PoA of First Price Auctions

Hu Fu

November 16, 2017

Recall that we defined the Price of Anarchy (PoA) of a game with respect to a certain objective function as the largest gap, measured as a ratio, between the best achievable objective and the objective in a worst possible equilibrium of the game. In this lecture we consider PoA of auctions, and throughout we focus on social welfare as the objective. In the process, we develop a framework for proving PoA, that of *smoothness*, first developed by Roughgarden (2015).

1 PoA for First Price Auctions

Notations: Let bidder i have value v_i for the item, and her bid is denoted by b_i . We allow randomized strategies, and strategy s_i is a distribution over bids.¹ Let (s_1^N, \dots, s_n^N) be a Nash equilibrium. For now we consider full information game, i.e., everyone knows everyone else's value. From an equilibrium point of view though, what matters is only the distribution of the other bidders' bids, and not their values. Let $u_i(s_i, s_{-i})$ denote the utility of bidder i .

Let us make a first attempt at bounding the PoA for a single item first price auction. The idea is to consider possible deviations that bidders could make. The utilities of these deviations (when the other bidders keep playing s_{-i}^N) should bear a relationship to the optimal social welfare, and we would like that the equilibrium condition $u_i(s_i^N, s_{-i}^N) \geq u_i(s_i^*, s_{-i}^N)$, where s_i^* is the deviation considered, would lead to a statement about the Nash welfare with respect to the optimal welfare.

Theorem 1. *A single item first price auction has PoA at most 2.*

*First proof.*² It is not hard to see that, if a pure Nash exists in the first price auction, it actually has to be efficient (i.e., welfare optimal): if the bidder with the highest value v_{i^*} is not winning, the highest bid on the item must be more than v_{i^*} , which means the winner incurs negative utility, a contradiction. Therefore only the mixed equilibria are interesting.

Let i^* be the bidder with the highest value, and consider a deviation of bidding $b_{i^*}^* = \frac{v_{i^*}}{2}$. Let x_1, \dots, x_n be the probability that bidder i wins the auction in the Nash equilibrium. The equilibrium condition gives us

$$u_i(s_i^N, s_{-i}^N) = v_{i^*} x_{i^*} - \mathbf{E} \left[b_{i^*} \cdot \mathbb{1}_{b_{i^*} \geq \max b_j} \right] \geq u_i(b_{i^*}^*, s_{-i}^N) = \frac{v_{i^*}}{2} \cdot \Pr_{b_j \sim s_j^N} \left[\max_{j \neq i^*} b_j < \frac{v_{i^*}}{2} \right].$$

¹With a slight abuse of notation, we use b_i to denote the pure strategy of bidding b_i .

²The language of this proof aims to explain the thought process. Its style is not recommended for formal mathematical writing.

Note that v_{i^*} is the optimal social welfare. So the RHS already looks like a fraction of the optimal social welfare. If the probability in it is 1, then we immediately get a PoA of 2 (because the Nash social welfare is lower bounded by LHS, the utility of bidder i^*). But what happens when some b_j is higher than $\frac{v_{i^*}}{2}$? That means someone else must have a value larger than $\frac{v_{i^*}}{2}$, and whenever that happens, the social welfare is no lower than $\frac{v_{i^*}}{2}$ is generated, because in a first price auction, bidders never bid more than their values. To formalize this idea, note that for all other bidders,

$$u_j(s_j^N, s_{-j}^N) = v_j x_j - \mathbf{E} \left[b_j \cdot \mathbb{1}_{b_j > \max_k b_k} \right] \geq 0.$$

Summing everything up, we have

$$\sum_i v_i x_i \geq \frac{v_{i^*}}{2} \cdot \Pr_{b_j \sim s_j^N} \left[\max_{j \neq i^*} b_j < \frac{v_{i^*}}{2} \right] + \mathbf{E}_{b_i \sim s_i^N} \left[\max_i b_i \right] \geq \frac{v_{i^*}}{2},$$

where the second inequality comes from the fact that, whenever the event in the first term does not happen, $\max_i b_i$ is at least $\frac{v_{i^*}}{2}$. \square

Second proof. Looking at the proof more carefully, we see that as we sum up $u_i(s_i^N, s_{-i}^N)$, we get $\sum_i v_i x_i - \sum_i p_i$, where p_i is the expected payment made by bidder i in the Nash equilibrium. To bound the social welfare $\sum_i x_i v_i$, it suffices to construct deviations b_i^* such that

$$\sum_i u_i(b_i^*, s_{-i}^N) \geq \frac{1}{2} \text{OPT} - \sum_i p_i = \frac{1}{2} \text{OPT} - \mathbf{E}_{b_i \sim s_i^N} \left[\max_i b_i \right]. \quad (1)$$

Now it is easy to see that setting $b_{i^*}^*$ to be $\frac{v_{i^*}}{2}$ suffices to guarantee $u_{i^*}(b_{i^*}^*, s_{-i^*}^N)$ to be at least the RHS of (1): when $b_{i^*}^*$ is the highest bid, the utility of i^* is $\frac{v_{i^*}}{2}$, which is trivially at least the RHS; when $b_{i^*}^*$ is not the highest, bidder i^* gets utility 0 while the RHS of (1) is negative. Therefore for all other $i \neq i^*$, we need just b_i^* (and hence $u_i(b_i^*, s_{-i}^N)$) to be 0. This completes the proof. \square

2 Smooth Mechanisms

The second proof of Theorem 1 has the advantage that the proof simply aims at finding deviations whose utilities satisfy (1); what remains is mechanical manipulation to derive a PoA bound from (1). The following definition abstracts this mechanical part and streamlines it, allowing us to focus on finding deviation strategies that lead to PoA bounds.

Definition 1. A mechanism is (λ, μ) -smooth, $\lambda, \mu \geq 0$, if for any value profile $\vec{v} = (v_1, \dots, v_n)$, any strategy s_1, \dots, s_n , there exists for each bidder i a deviation $s_i^*(\vec{v}, s_i(v_i))$, such that

$$\sum_i u_i^{v_i}(s_i^*(\vec{v}, s_i(v_i)), s_{-i}(v_{-i})) \geq \lambda \text{OPT} - \mu \sum_i p_i(s_i(v_i), s_{-i}(v_{-i})), \quad (2)$$

where $u_i^{v_i}$ is the utility of bidder i with value v_i , and $p_i(s_i, s_{-i})$ is the payment made by bidder i .³

³There are several definitions of smooth games/mechanisms in the literature. They are similar in spirit but not altogether comparable. We adhere here to the definition in Syrgkanis and Tardos (2013). For our illustrations we will not make use of the deviation's dependence on the player's own strategy.

Note that we allow randomized strategies, and $u_i^{v_i}$ and p_i both have expectations implicit. The next theorem formalizes the “mechanical” part of the proof for Theorem 1.

Theorem 2. *A (λ, μ) -smooth mechanism has PoA at most $\frac{\max\{1, \mu\}}{\lambda}$, assuming each bidder has an action that guarantees utility 0.*

The action guaranteeing zero utility can be bidding 0 or not participating the auction.

Proof. Recall that (s_1^N, \dots, s_n^N) is an equilibrium, hence for each i , $u_i^{v_i}(s_i^N, s_{-i}^N) \geq u_i^{v_i}(s_i^*(\vec{v}, s_i^N), s_{-i}^N)$, where s_i^* is the strategy in Definition 1. Summing up the inequalities, we have

$$\sum_i u_i^{v_i}(s_i^N, s_{-i}^N) = \text{SW}^N - \sum_i p_i(s_i^N, s_{-i}^N) \geq \sum_i u_i^{v_i}(s_i^*(\vec{v}, s_i^N), s_{-i}^N) \geq \lambda \text{OPT} - \mu \sum_i p_i(s_i^N, s_{-i}^N),$$

where we use SW^N to denote the social welfare of the Nash equilibrium. Moving the term $\sum_i p_i(s_i^N, s_{-i}^N)$, we have

$$\text{SW}^N \geq \lambda \text{OPT} + (1 - \mu) \sum_i p_i(s_i^N, s_{-i}^N).$$

For $\mu \leq 1$, obviously $\text{SW}^N \geq \lambda \text{OPT}$. For $\mu > 1$, note that $\text{SW}^N \geq \sum_i p_i(s_i^N, s_{-i}^N)$ (otherwise some bidder has negative utility and can play the quitting action), and so

$$\mu \text{SW}^N = \text{SW}^N + (\mu - 1) \text{SW}^N \geq \text{SW}^N + (\mu - 1) \sum_i p_i(s_i^N, s_{-i}^N) \geq \lambda \text{OPT},$$

and $\text{SW}^N \geq \frac{\lambda}{\mu} \text{OPT}$. □

3 A tighter bound for the first price auction

Equipped with a definition for smooth games and Theorem 2, let us try to improve the PoA bound for first price auctions. Looking at inequality (1), we see that the deviation $b_{i^*}^*$ we devise is rather loose: when $\max b_i$ is $\frac{v_{i^*}}{2} - \epsilon$, the RHS is almost 0, whereas the LHS is $\frac{v_{i^*}}{2}$. Intuitively, we would like to have a deviation $b_{i^*}^*$ such that, no matter how much $\max b_i$ is, the inequality (2) is about tight for some fixed λ and μ . It is not hard to see that, for a deterministic deviation $b_{i^*}^*$, (1) is the best one can do. (Convince yourself of this.) But we can have a randomized $b_{i^*}^*$ instead. Let’s fix μ to be 1, and we would like to have that, for any value of $\max_i b_i$, the expected utility of deviating to $b_{i^*}^*$ is $\lambda v_{i^*} - \max b_i$, for some λ . Let the density of $b_{i^*}^*$ at v be $g(v)$, then

$$u_{i^*}^{v_{i^*}}(b_{i^*}^*, \max_{i \neq i^*} b_i = v) = \int_v^{v_{i^*}^*} (v_{i^*} - s)g(s) ds.$$

We hope this to be equal to $\lambda \text{OPT} - v$. Taking derivative with respect to v on both sides, we get $(v_{i^*} - v)g(v) = 1$, which suggests that we should have $g(v) = \frac{1}{v_{i^*} - v}$. Let λ be $1 - \frac{1}{e}$. Since $\int_0^{\lambda v_{i^*}} \frac{1}{v_{i^*} - v} dv = 1$, we’ll let the deviation $b_{i^*}^*$ be distributed on $[0, \lambda v_{i^*}]$ with density $\frac{1}{v_{i^*} - v}$. As before, for all other bidder i , b_i^* is to bid simply 0. It is not hard to verify

$$\sum_i u_i^{v_i}(b_i^*, s_{-i}^N) \geq \lambda v_{i^*} - \sum_i p_i(s_i^N, s_{-i}^N),$$

which shows that the first price auction is a $(1 - \frac{1}{e}, 1)$ -smooth game.

Theorem 3 (Syrngkanis and Tardos, 2013). *The first price auction is a $(1 - \frac{1}{e}, 1)$ -smooth game, and hence has PoA at most $\frac{e}{e-1}$.*

Exercise 1. Prove Theorem 3 using the deviation we have just derived.

4 Extension theorem

One of the major advantages of showing PoA with smoothness is that the PoA bound we prove often extends to more general equilibrium concepts for free. For example, let's consider an incomplete information game, where each bidder's type is drawn from the product distribution $F_1 \times F_2 \cdots \times F_n$, and Bayesian Nash Equilibrium is the proper solution concept.

Theorem 4 (Roughgarden, 2015; Syrgkanis, 2012). *At any Bayesian Nash equilibrium of a (λ, μ) -smooth game, the expected social welfare is at least $\frac{\lambda}{\max\{\mu, 1\}}$ fraction of the optimal expected social welfare.*

In other words, the *Bayesian Price of Anarchy* (BPoA) is at most $\frac{\max\{\mu, 1\}}{\lambda}$. As an immediate corollary of Theorem 3, we have

Corollary 5. *A first price auction where bidders' values are drawn from a product distribution has BPoA at most $\frac{e}{e-1}$.*

A naïve attempt at proving the theorem is to use the deviation $s_i^*(\vec{v}, s_i^N)$ in the definition of smooth games, and argue that $\mathbf{E}[u_i^{v_i}(s_i^N(v_i), s_{-i}^N(v_{-i}))] \geq \mathbf{E}[u_i^{v_i}(s_i^*(\vec{v}, s_i^N(v_i)), s_{-i}^N(v_{-i}))]$. This, however, is not true. The BNE condition only says that the utility of strategy s_i^N is no less than that of any other strategy played against $s_{-i}^N(v_{-i})$; this other strategy, however, cannot depend on any knowledge on the *realization* of the other bidders' types or strategies. In some sense, the deviation s_i^* in Definition 1 allows more power than can be used in a naïve application of the BNE condition. For this reason, we need a “Doppelgänger” trick, which simulates s_i^* without resorting to the actual realization of v_{-i} , but by resampling the types (and simulating the equilibrium strategies using these samples). Linearity of expectation and independence will be crucial in guaranteeing that this leads to the PoA expression we hope for.

Proof of Theorem 4. For each player i , let \vec{w} be a profile of types drawn independently from $F_1 \times \cdots \times F_n$. Consider the utility of deviation $s_i^*((v_i, \vec{w}_{-i}), s_i^N(w_i))$:

$$\begin{aligned} \mathbf{E}_{\vec{v}} \left[u_i^{v_i}(s^N(\vec{v})) \right] &\geq \mathbf{E}_{\vec{v}, \vec{w}} \left[u_i^{v_i}(s_i^*((v_i, \vec{w}_{-i}), s_i^N(w_i)), s_{-i}^N(\vec{v}_{-i})) \right] \\ &= \mathbf{E}_{\vec{v}, \vec{w}} \left[u_i^{w_i}(s_i^*((w_i, \vec{w}_{-i}), s_i^N(v_i)), s_{-i}^N(v_{-i})) \right] \\ &= \mathbf{E}_{\vec{v}, \vec{w}} \left[u_i^{w_i}(s_i^*(\vec{w}, s_i^N(v_i)), s_{-i}^N(\vec{v}_{-i})) \right]. \end{aligned}$$

In the first equality, we used the independence among v_i, w_i, \vec{w}_{-i} and \vec{v}_{-i} . Note that the probability that we see a pair (v_i, w_i) is the same as the probability we see (w_i, v_i) .

Summing up, we have

$$\begin{aligned}
\mathbf{E}_{\vec{v}} \left[\sum_i u_i^{v_i}(s^N(\vec{v})) \right] &\geq \mathbf{E}_{\vec{v}, \vec{w}} \left[\sum_i u_i^{w_i}(s_i^*(\vec{w}, s_i^N(v_i)), s_{-i}^N(\vec{v}_{-i})) \right] \\
&\geq \mathbf{E}_{\vec{v}, \vec{w}} \left[\lambda \text{OPT}(w) - \mu \sum_i p_i(s^N(\vec{v})) \right] \\
&= \lambda \mathbf{E}_{\vec{w}} [\text{OPT}(w)] - \mu \mathbf{E}_{\vec{v}} \left[\sum_i p_i(s^N(\vec{v})) \right].
\end{aligned}$$

The second inequality is the consequence of the smoothness of the game. Note that OPT , which is the sum of bidders' values in the optimal allocation, is a function of the actual types \vec{w} of the bidders, whereas the payments depend only on the bidders' strategies, and are therefore are functions of $s^N(\vec{v})$.

Now the same argument as in the proof of Theorem 2 gives us the BPOA bound. \square

5 Composability: Simultaneous item auctions retain smoothness

Another advantage of the smoothness framework is that it allows us to compose (λ, μ) -smooth mechanisms while maintaining smoothness, where by composing we mean running multiple mechanisms simultaneously, and we allow bidders to have combinatorial preferences over the outcomes of these mechanisms (under some restrictions on the valuations). For example, we have seen that the first price auction is $(1 - \frac{1}{e}, 1)$ -smooth. Suppose we have m items to sell, and each bidder has combinatorial valuations over these items, we may consider running the following simultaneous item auction:

Definition 2. In a simultaneous first price auction, every bidder puts a bid b_{ij} on each item simultaneously. Then each item is allocated to the highest bidder for that item and each bidder pays for her bid on each item she wins.

By the homework you should have convinced yourselves that any pure Nash of this auction maximizes the social welfare. The next theorem shows that when each bidder has *XOS* valuations, then the simultaneous first price auction inherits the smoothness of the single item first price auction, and has Bayesian PoA of $1 - 1/e$.

Definition 3. A valuation $v : 2^M \rightarrow \mathbb{R}$ is *XOS* if for every $S \subseteq M$, there exists an additive valuation a_S such that $a_S(S) = v(S)$, $a_S(T) \leq v(T)$ for all $T \subseteq S$, and $a_S(j) = 0$ for all $j \notin S$.

Exercise 2. Show that any submodular valuation is *XOS*.

Theorem 6 (Syrngkanis and Tardos, 2013). *A combinatorial auction that runs a (λ, μ) -smooth auction simultaneously for each item is (λ, μ) -smooth, if all bidders have *XOS* valuations.*⁴

⁴The result in ?? is more general, by extending the definition of *XOS* from valuations over items to valuations over outcomes.

Corollary 7. *A simultaneous first price auction in which all bidders have XOS valuations is $(1 - \frac{1}{e}, 1)$ -smooth, and hence has Bayesian PoA of $\frac{e}{e-1}$.*

Proof of Theorem 6. We prove the case for Nash equilibrium, and the proof of Theorem 4 extends the analysis to Bayesian PoA. Let $p_{ij}(\vec{s})$ denote the expected payment made by player i in the auction for item j , when players play strategies \vec{s} .

Given valuation profile (v_1, \dots, v_n) , let (S_1^*, \dots, S_n^*) be an optimal allocation that maximizes social welfare. Recall that by Definition 3, for each bidder i there is an additive valuation a_{i, S_i^*} such that $a_{i, S_i^*}(T) \leq v_i(T)T$ for all $T \subseteq M$ and $a_{i, S_i^*}(S_i^*) = v_i(S_i^*)$. For each bidder i , let s_i^* be the deviation where for each item j , the bidder bids according to the deviation $s_i^*((a_{i, S_i^*}, a_{-i, S_{-i}^*}), s_i^N(v_i))$ given by the definition of smoothness (2). For each item $j \in S_{i^*}^*$, by smoothness we have

$$\sum_i u_i^{a_{i, S_i^*}}(s_i^*, s_{-i}^N) \geq \lambda a_{i^*, S_{i^*}^*}(j) - \mu \sum_i p_{i, j}(s^N(\vec{v})).$$

Summing over the items, we have

$$\sum_i u_i^{a_{i, S_i^*}}(s_i^*, s_{-i}^N) \geq \lambda \sum_i a_{i, S_i^*}(S_i^*) - \sum_i p_i(s^N) = \lambda \text{OPT} - \sum_i p_i(s^N),$$

where we crucially used that $a_{i, S_i^*}(S_i^*) = v_i(S_i^*)$. By the equilibrium condition, we have

$$\sum_i u_i^{v_i}(s_i^N, s_{-i}^N) \geq \sum_i u_i^{v_i}(s_i^*, s_{-i}^N) \geq \sum_i u_i^{a_{i, S_i^*}}(s_i^*, s_{-i}^N) \geq \lambda \text{OPT} - \sum_i p_i(s^N),$$

where the second inequality comes from the fact that $v_i(S) \geq a_{i, S_i^*}(S)$ for any $S \subseteq M$. This shows the smoothness of the simultaneous auction. \square

6 The case of subadditive valuations: a tantalizing non-smooth analysis

As the class of valuations gets larger, computing good approximations of the optimal social welfare becomes harder, and designing tractable auctions that are incentive compatible are even more so. We have seen that for gross substitute valuations one can run the VCG auction in time polynomial in n and m , and achieves optimal social welfare. In class we have seen that with maximal-in-distributional-range mechanisms, one can design $O(1)$ approximately optimal mechanisms for some classes of valuations that are complement free. In general, however, designing $O(1)$ approximately optimal incentive compatible auctions is very challenging, even when one is allowed unbounded computational power but polynomial amount of communication. In fact, strong lower bounds are known when communication is restricted to value oracles (Dobzinski, 2011; Dughmi and Vondrák, 2015), and, even allowing general communications, the best currently known incentive compatible approximation for submodular and XOS valuations is $O(\sqrt{\log m})$ (Dobzinski, 2016). In comparison, it is remarkable that simple auctions such as the simultaneous first price auction guarantee $O(1)$ approximations to the optimal social welfare at any equilibrium.

To further appreciate the power of the simultaneous item auction, we look at the most general class of valuations that are complement free: a valuation function $v : 2^M \rightarrow \mathbb{R}$ is said to be *subadditive* if for all $S, T \subseteq M$, $v(S) + v(T) \geq v(S \cup T)$.

Exercise 3. Show that any XOS valuation is subadditive.

Remark 1. Over the semester we have seen a hierarchy: gross substitute valuations are submodular, submodular valuations are XOS, and XOS valuations are subadditive.

The following theorem confirms that the simultaneous first price auction has $O(1)$ PoA for all complement-free valuations.

Theorem 8. (Feldman et al., 2013) *The simultaneous first price auction with subadditive bidders has Bayesian PoA of at most 2, assuming bidders' types are drawn from a product distribution.*

The proof of Theorem 8 is not exactly a smoothness argument, because the deviations we consider will depend on the other bidders' Nash strategies (and not only on their types). However, the overall argument still resembles that for the smooth games.

Proof. We will prove the PoA for Nash equilibrium only. The extension argument to the Bayesian PoA is a twist on (although not a replica of) that of Theorem 4, and is left as an exercise.

Let (s_1^N, \dots, s_n^N) be a Nash equilibrium. Given valuation profiles (v_1, \dots, v_n) , let (S_1^*, \dots, S_n^*) be an allocation that maximizes the social welfare. Let $q_{-i} \in \mathbb{R}^M$ be a vector whose coordinate for any $j \notin S_i^*$ is 0, and for any $j \in S_i^*$ is equal to the random variable representing the highest bid put by a bidder other than i on j , when each bidder bids according to her Nash strategy. Let B_{-i} be the distribution of q_{-i} . For each bidder i , consider the following deviation s_i^* : draw a vector b_i^* from the distribution B_{-i} , and bid that on the items. In plain language, the deviation bids 0 on all items outside S_i^* , and imitates on S_i^* the randomized prices set by the other bidders in the equilibrium.

The key observation is that

$$u_i^{v_i}(s_i^*, s_{-i}^N) \geq \frac{1}{2}v_i(S_i^*) - \sum_{i \in N, j \in S_i^*} p_{ij}(s^N).$$

We only need to argue that the expected value for bidder i when deviating to s_i^* is at least $\frac{1}{2}v_i(S_i^*)$, as the expected payment she makes is obviously upper bounded by $\sum_{i \in N, j \in S_i^*} p_{ij}(s^N)$ — the sum of her bids altogether is $\sum_{j \in S_i^*} \mathbf{E}[q_{-i}(j)]$, which is bounded by $\sum_{i \in N, j \in S_i^*} p_{ij}(s^N) = \sum_{j \in S_i^*} \mathbf{E}_{b_k \sim s_k^N}[\max_{k \in N} b_k(j)]$. Now to see the bound on the value, notice that for any realization of q_{-i} and b_i^* , suppose S is the set of items won by bidder i when she bids b_i^* and the highest bids put by the other bidders are q_{-i} , then $S_i^* \setminus S$ is the set she wins when she bids q_{-i} and the other bidders' highest bids are b_i^* . By construction, these events occur with exactly the same probability, and therefore her expected value for the set she wins, when conditioning on any realization of the two vectors without specifying which one is which, is equal to $\frac{1}{2}v_i(S) + \frac{1}{2}v_i(S_i^* \setminus S) \geq \frac{1}{2}v_i(S_i^*)$, where the inequality comes from subadditivity. As this bound holds for any conditioning, it holds in expectation as well.

Now applying the Nash condition, we see that

$$\sum_i u_i^{v_i}(s_i^N, s_{-i}^N) \geq \sum_i u_i^{v_i}(s_i^*, s_{-i}^N) \geq \frac{1}{2} \sum_i v_i(S_i^*) - \sum_{i \in N, j \in M} p_{ij}(s^N) = \frac{1}{2} \text{OPT} - \sum_i p_i(s^N).$$

Cancelling out the payments on both sides leads to the conclusion. \square

Remark 2. As a consequence of a more general result by Roughgarden (2014), any auction that uses subexponential amount of communication must have PoA of at least 2 for subadditive valuations. In this sense, the simultaneous first price auction has the best price of anarchy among all such auctions. Roughgarden’s proof makes use of lower bounds for nondeterministic communication protocols, which are beyond the scope of this course. The interested reader should refer to the paper or the lecture notes by Tim Roughgarden.

References

- Dobzinski, S. (2011). An impossibility result for truthful combinatorial auctions with submodular valuations. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 139–148.
- Dobzinski, S. (2016). Breaking the logarithmic barrier for truthful combinatorial auctions with submodular bidders. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 940–948.
- Dughmi, S. and Vondrák, J. (2015). Limitations of randomized mechanisms for combinatorial auctions. *Games and Economic Behavior*, 92:370–400.
- Feldman, M., Fu, H., Gravin, N., and Lucier, B. (2013). Simultaneous auctions are (almost) efficient. In *Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*, pages 201–210.
- Roughgarden, T. (2014). Barriers to near-optimal equilibria. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 71–80.
- Roughgarden, T. (2015). Intrinsic robustness of the price of anarchy. *J. ACM*, 62(5):32.
- Syrngkanis, V. (2012). Bayesian games and the smoothness framework. *CoRR*, abs/1203.5155.
- Syrngkanis, V. and Tardos, É. (2013). Composable and efficient mechanisms. In *Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*, pages 211–220.