CS364B: Frontiers in Mechanism Design Lecture #15: The Price of Anarchy of Bayes-Nash Equilibria *

Tim Roughgarden†

February 26, 2014

1 The Story So Far

We're square in the middle of Part IV of the course, where we take auction simplicity as a hard constraint and seek conditions under which an auction's equilibria are guaranteed to have near-optimal welfare. We'd like to answer this question for different auction formats, different valuation classes, and different equilibrium concepts.

1.1 Recap: Simultaneous Single-Item Auctions

Last lecture we studied simultaneous second-price auctions (S2A), where each item is sold separately using a Vickrey auction. Bidders still have combinatorial valuations (e.g., submodular valuations), and hence have up to 2^m private parameters, but the auction format only allows to report m parameters (one bid per item). The main result of last lecture was the following.

Theorem 1.1 *Suppose that:*

1. *Every bidder i has a submodular (or even XOS[1]) valuation;*
2. $\mathbf{b}$ *is a pure Nash equilibrium that satisfies weak no overbidding (WNO).*[2]

Then the welfare $\sum_{i=1}^n v_i(S_i(\mathbf{b}))$ *of* $\mathbf{b}$ *is at least 50% times the maximum possible,*

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

[1]Recall that an XOS valuation is the maximum of additive valuations — for some additive valuations $\mathbf{a}_i^1, \ldots, \mathbf{a}_i^r$, $v_i(S) = \max_{\ell=1}^r \{\sum_{j \in S} a_{ij}^\ell\}$ for every $S \subseteq U$.

[2]Recall that $\mathbf{b}$ satisfies WNO if $\sum_{j \in S_i(\mathbf{b})} b_{ij} \leq v_i(S)$ for every i, where $S_i(\mathbf{b})$ denotes the items i wins in $\mathbf{b}$.

Recall that the no overbidding condition is necessary for non-trivial welfare guarantees even in the Vickrey auction, because of "bluffing" equilibria with arbitrarily low welfare.

1.2 Recap: Extension Theorems

Theorem 1.1 is only about pure Nash equilibria, which is a full-information equilibrium concept. Auctions are generally better modeled as games of incomplete information, where the Bayes-Nash equilibrium is the appropriate equilibrium concept. The goal for this lecture is to understand the extend to which Theorem 1.1 continuous to hold for Bayes-Nash equilibria when bidders valuations are private but drawn from a commonly known prior distribution.[3]

Our plan is to derive POA bounds for Bayes-Nash equilibria via "extension theorems." Recall from CS364A (Lecture #14) that an extension theorem extends a POA bound from a stronger equilibrium concept (which is easier to analyze) to a weaker (which is better motivated). Last quarter, we considered only full-information games. We weren't happy with POA bounds for pure Nash equilibria because they might not exist or might be hard (e.g., PLS-complete) to compute; we weren't happy with POA bounds for mixed Nash equilibria because they might also be hard (e.g., $PPAD$-complete) to compute. Correlated and coarse correlated equilibria are tractable and learnable equilibrium concepts, and we used an extension theorem for "smooth games" to losslessly extend POA bounds for pure Nash equilibria to POA bounds for coarse correlated equilibria. We're in a similar position now: last lecture we prove POA bounds for the easily analyzes but poorly motivated pure Nash equilibria of full-information S2A's; can we extend these bounds to Bayes-Nash equilibria of their more general incomplete-information counterparts?

1.3 Recap: Proof of Theorem 1.1

Extension theorems apply only to a restricted class of POA bounds, and it is important we isolate where the different hypotheses are applied last lecture's POA proof. Let's recap the

[3]In fact, if we only cared about the price of anarchy of pure Nash equilibria of S2As, we could have used previous results to deduce a bound of $\frac{1}{2}$. In Lecture #7, we proved that if bidders with submodular valuations bid sincerely in the Kelso-Crawford auction, then the auction terminates with an allocation with welfare at least 50% of the maximum possible. This same guarantee applies to pure Nash equilibria of S2A's; see the Exercises for details.

main steps. For a judicious choice of hypothetical deviations $\mathbf{b}_1^*, \ldots, \mathbf{b}_n^*$, we showed that

$$\underbrace{\sum_{i=1}^{n} v_i(S_i(\mathbf{b}))}_{\text{welfare of } \mathbf{b}} \geq \sum_{i=1}^{n} u_i(\mathbf{b}) \tag{1}$$

$$\geq \sum_{i=1}^{n} u_i(\mathbf{b}_i^*, \mathbf{b}_{-i}) \tag{2}$$

$$\geq \text{OPT welfare} - \sum_{i=1}^{n} \sum_{j \in S_i(\mathbf{b})} b_{ij} \tag{3}$$

$$\geq \text{OPT welfare} - \sum_{i=1}^{n} v_i(S_i(\mathbf{b})). \tag{4}$$

Inequalities (1)–(4) immediately imply Theorem 1.1. The role of the hypotheses in this derivation are:

1. Inequality (1) depends only on the nonnegativity of prices.
2. Inequality (2) depends only on the assumption that $\mathbf{b}$ is a pure Nash equilibrium.
3. Inequality (3) depends on the choice of $\mathbf{b}_1^*, \ldots, \mathbf{b}_n^*$; given this choice, the inequality follows from algebraic manipulations (see last lecture).
4. Inequality (4) depends only on the WNO condition, that $\sum_{j \in S_i(\mathbf{b})} b_{ij} \leq v_i(S_i(\mathbf{b}))$ for every i.

1.4 Consequence: S2A Games are $(1, 1)$-Smooth

The first important point is that the assumption that $\mathbf{b}$ is a pure Nash equilibrium is used only in the second step; in particular, the manipulations in the third step are valid for every bid profile $\mathbf{b}$, not just equilibria (as the reader should verify). To explain the second important point, we need to recall how the hypothetical deviations $\mathbf{b}_1^*, \ldots, \mathbf{b}_n^*$ are defined. Let $S_1^*, \ldots, S_n^*$ be an allocation that maximizes the welfare $\sum_{i=1}^{n} v_i(S_i^*)$. For bidder i, let a_i^* be an additive valuation in the XOS representation of v_i such that $\sum_{j \in S_i^*} a_{ij}^* = v_i(S_i^*)$ (and $\sum_{j \in S} a_{ij}^* \leq v_i(S)$ for every $S \subseteq U$). Then, $\mathbf{b}_i^*$ is defined as

$$b_{ij}^* = \begin{cases} a_{ij}^* & \text{if } j \in S_i^* \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

The important observation is that each $\mathbf{b}_i^*$ is defined independently of $\mathbf{b}$.

These two observations — that the Nash equilibrium hypothesis is used only in a minimal way in (2), with the hypothetical deviations $\mathbf{b}_1^*, \ldots, \mathbf{b}^n$ chosen independently of the particular equilibrium $\mathbf{b}$ — means that (1)–(4) is a "smoothness proof" in the sense of CS364A.

Specifically, inequalities (3) and (4) imply that full-information S2A games are $(1, 1)$-smooth payoff-maximization games, provided we restrict to bid profiles $\mathbf{b}$ that satisfy the WNO condition. The extension theorem from CS364A (Lecture #14) immediately implies that every mixed Nash equilibrium, correlated equilibrium, or coarse correlated equilibrium that randomizes only over bid vectors that satisfy WNO has expected welfare at least 50% of the maximum possible.[4]

2 The POA of Bayes-Nash Equilibria of S2As

The observation in Section 1.4 is progress — certainly we're happier with bounds for many equilibrium concepts rather than just pure Nash equilibria. Still, all of these equilibrium concepts are traditionally meant for full-information games, and auctions are most naturally modeled as games of incomplete information. Can we also extend our POA bounds to Bayes-Nash equilibria?

2.1 The Bayes-Nash Price of Anarchy

The Bayes-Nash (BN) POA is defined with respect to a prior distribution $\mathbf{F}$ as follows:

$$\frac{\inf_{\text{BNE } \sigma} \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\text{welfare of } \sigma(\mathbf{v})]}{\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\text{optimal welfare for } \mathbf{v}]}. \tag{6}$$

In the denominator, the optimal welfare is a function of the valuation profile $\mathbf{v}$, and our benchmark is the expected (over $\mathbf{v} \sim \mathbf{F}$) optimal welfare. In the numerator, we use the expected welfare of the worst Bayes-Nash equilibrium. For simplicity, we've stated the BN POA (6) only for pure Bayes-Nash equilibria. As we know, for S2A we'll have to restrict our analysis to Bayes-Nash equilibria that satisfy some type of no overbidding condition. With a mixed Bayes-Nash equilibrium σ, the numerator contains a second expectation, over the random actions chosen in a strategy profile $\sigma(\mathbf{v})$. Our positive results will all hold for mixed Bayes-Nash equilibria.

The Bayes-Nash POA (6) depends on the prior distribution $\mathbf{F}$. Which prior should we study? At the present level of generality, any particular choice would be arbitrary, so what we'd really like are bounds agnostic to the prior. Most ambitiously, maybe we can prove good BN POA bounds for *every* prior $\mathbf{F}$?

2.2 Impossibility Result for Correlated Valuations

In this section we show the impossibility of good Bayes-Nash POA bounds for S2A's with an arbitrary prior distribution $\mathbf{F}$. Concretely, we construct a correlated prior distribution $\mathbf{F}$ over unit-demand valuations for n bidders for which the Bayes-Nash POA is $\Omega(n^{1/4})$. This

[4]The proof of the extension theorem can also be reworked so that the no overbidding condition only has to hold in expectation; we'll see some examples of this later this lecture.

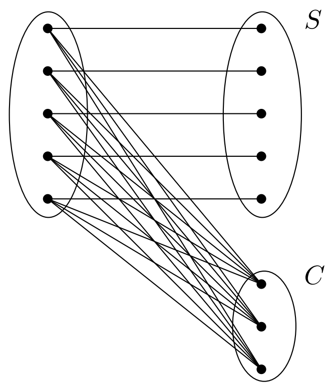

Figure 1: Correlated distribution with bad Bayes-Nash Price of Anarchy.

example will also develop your intuition for the twisted scenarios you can implement with arbitrary correlated distributions.

The example has n bidders and $n + \sqrt{n} - 1$ items. Each bidder will, with probability 1, have a 0-1 unit-demand valuation. Thus, we can think of the welfare-maximization problem as unweighted bipartite matching.

To simplify the exposition, we assume that each item has a reserve price of $1/n^{1/4}$. This assumption is not necessary, since each reserve price can be simulated with a "dummy bidder" who only wants that one item and has valuation $1/n^{1/4}$ for it.[5]

We now describe the valuation distribution $\mathbf{F}$ via a randomized algorithm for generating bidders' valuations. First, we choose a subset $C \subseteq U$ of $\sqrt{n} - 1$ "common" items uniformly at random from U. All bidders want the items of C. (Evidently, the distribution $\mathbf{F}$ is correlated.) There are n "special items" $S = U \setminus C$, and each bidders want a distinct one. Precisely, we randomly order the special items $S = \{1, 2, \ldots, n\}$, and define unit-demand valuations by

$$v_{ij} = \begin{cases} 1 & \text{if } j = i \text{ or } j \in C \\ 0 & \text{otherwise.} \end{cases}$$

The corresponding bipartite graph is the union of a perfect matching (with S) and a complete bipartite graph (with S); see Figure 1. Because of the perfect matching, the optimal welfare is n with probability 1.

We now consider the Bayes-Nash equilibria. We assume that every bidder's bid space is the set of bid vectors that satisfy the SNO condition (i.e., $\sum_{j\ inS} b_{ij} \leq v_i(S)$ for every $S \subseteq U$). A more complex variation on the example shows a polynomial lower bound on the BN POA under weaker no overbidding assumptions (see Exercises).[6]

[5]Formally, every Bayes-Nash equilibrium with reserve prices can be extended to one with dummy bidders by having all of the dummy bidders bid their valuations. The welfare of the Bayes-Nash equilibrium increases by at most the sum of the reserve prices, which in this case is $O(n^{3/4})$.

[6]We are ignoring the issue of the existence of Bayes-Nash equilibria which, perhaps surprisingly, is not well understood in settings like the present one which have continuous type spaces, continuous auctions spaces, and discontinuous payoffs. An easy solution is discretization: assuming that valuations are bounded, and both valuations and bids are restricted to be multiples of an arbitrarily small positive number ϵ, the

When a bidder learns its valuation — the $\sqrt{n}$ items that it wants — it knows that one of the desired items is special and that the rest are common. By the symmetry in the construction, however, it has no idea which of the items is special. Intuitively, this means that almost no bidders will successfully acquire their special item, instead wasting their bids on the common items.

To make this intuition precise, first note that the SNO condition prevents a bidder from bidding at least $1/n^{1/4}$ on more than $n^{1/4}$ items (its maximum value is 1). By symmetry, the probability that a bidder bids at least $n^{1/4}$ (and thus has a chance of winning) its special good is at most $n^{1/4}/\sqrt{n} = 1/n^{1/4}$. It follows that the expected number of bidders that acquire their special good is at most $n^{3/4}$. Since the common goods contribute at most $\sqrt{n}$ to the welfare, the expected welfare of every Bayes-Nash equilibrium that satisfies the SNO condition is $O(n^{3/4})$. Thus, the Bayes-Nash POA is $\Omega(n^{1/4})$ in this example.

2.3 Bayes-Nash POA Bound for Independent Valuations

The example in Section 2.2 shows that we need to restrict attention to product prior distributions.[7] Observe that a point mass distribution $\mathbf{F}$ can be represented as a product distribution (of single-dimensional point masses). Point mass distributions — where all players' valuations are known with certainty — correspond to full-information games. Thus, bounds for the Bayes-Nash POA with respect to an arbitrary product prior cannot be better than POA bounds for the full-information case. Hence, the best-case scenario is that we can prove that the Bayes-Nash POA is at least $\frac{1}{2}$ with respect to an arbitrary product prior distribution (subject to no overbidding).

Theorem 2.1 *For every product prior distribution* $\mathbf{F}$ *over XOS valuations, and every (mixed) Bayes-Nash equilibrium* σ *that satisfies a "no overbidding" condition,*

$$\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\mathit{welfare}(\sigma(\mathbf{v}))] \geq \frac{1}{2} \cdot \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\mathit{OPT\ welfare}(\mathbf{v})]\,. \tag{7}$$

This result is from [1]; the interpretation we give is from [2]. It will be clear in the proof exactly what "no overbidding" condition is needed on σ — basically an "on average" (over $\mathbf{v} \sim \mathbf{F}$) version of the WNO condition. For simplicity, we've stated inequality (7) and the proof below for pure Bayes-Nash equilibria; the general case is different only by having a second expectation on the left-hand side of (7), over the random action choices of $\sigma(\mathbf{v})$. A remarkable aspect of the result is that the bound does not depend on the prior distribution $\mathbf{F}$ at all, even though the set of Bayes-Nash equilibria depends on $\mathbf{F}$ in a complex way.

Proof of Theorem 2.1: Our starting point in the "smoothness-type" inequality (3). We proved last lecture that, for every valuation profile $\mathbf{v}$, there exist hypothetical deviations

existence of (mixed-strategy) Bayes-Nash equilibria reduces to Nash's existence theorem for finite games (see Exercises).

[7] It is an open question whether or not there are good Bayes-Nash POA bounds for interesting restricted classes of correlated prior distributions.

$\mathbf{b}_1^*, \ldots, \mathbf{b}_n^*$ such that, for every bid vector $\mathbf{b}$,

$$\sum_{i=1}^{n} u_i(\mathbf{b}_i^*, \mathbf{b}_{-i}) \geq \text{OPT welfare}(\mathbf{v}) - \sum_{i=1}^{n} \sum_{j \in S_i(\mathbf{b})} b_{ij}. \tag{8}$$

Recall the definition (5) of the deviation $\mathbf{b}_i^*$: if S_i^* is the bundle i gets in a welfare-maximizing allocation for the profile $\mathbf{v}$, then go "all-in" for S_i^*. This definition depends on $\mathbf{v}$ (and not just v_i) but is independent of $\mathbf{b}$, and we use the notation $\mathbf{b}_1^*(\mathbf{v}), \ldots, \mathbf{b}_n^*(\mathbf{v})$ henceforth.

The rest of the proof is effectively an "extension theorem" that transmutes a smoothness-type inequality of the form (8) into a bound on the Bayes-Nash POA. We give the proof only for S2A's (with no overbidding), but it will be clear that the proof applies to every class of games in which one can prove an inequality of the form (8) and relate the final term in (8) to the welfare of $\mathbf{b}$.

Let σ denote a Bayes-Nash equilibrium. As in all POA bounds, we need to figure out how to invoke the equilibrium hypothesis. In light of (2.1), the natural idea is to use the $\mathbf{b}_i^*(\mathbf{v})$ as the hypothetical deviation for bidder i. Unfortunately, *this makes no sense*. The reason is that, in a game of incomplete information, players take actions knowing only the prior $\mathbf{F}$, their own valuation v_i, and the strategies σ_{-i} used by the other players. The whole point is that bidder i does not know $\mathbf{v}_{-i}$ when it takes an action, and hence does not have sufficient information to execute the deviation $\mathbf{b}_i^*(\mathbf{v})$.

We salvage the above idea using the "doppelganger trick." While bidder i does not know $\mathbf{v}_{-i}$, it does know the prior $\mathbf{F}_{-i}$ from which $\mathbf{v}_{-i}$ is drawn. So, it does have sufficient information to sample "doppelganger" valuations $\mathbf{w} \sim \mathbf{F}$ and then deviate to the bid vector $\mathbf{b}_i^*(v_i, \mathbf{w}_{-i})$. That is, the bidder computes the hypothetical optimal bundles S_i^* it gets when the other bidders have valuations $\mathbf{w}_{-i}$, and goes all-in for S_i^*. Call this (randomized) deviation $\sigma_i^*(v_i)$. This "best approximation" of the deviation $\mathbf{b}_i^*(\mathbf{v})$ turns out to be sufficient for the proof.

Because σ is a Bayes-Nash equilibrium, bidder i's expected utility only goes down when it has the valuation v_i and switches from action $\sigma_i(v_i)$ to action $\sigma_i^*(v_i)$:

$$\begin{aligned} \mathbf{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}}[u_i(\sigma(\mathbf{v}))] &\geq \mathbf{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}}[u_i(\sigma_i^*(v_i), \sigma_{-i}(\mathbf{v}_{-i}))] \\ &= \mathbf{E}_{\mathbf{v}_{-i} \sim \mathbf{F}_{-i}, \mathbf{w} \sim \mathbf{F}}[u_i(\mathbf{b}_i^*(v_i, \mathbf{w}_{-i}), \sigma_{-i}(\mathbf{v}_{-i}))], \end{aligned}$$

where the equation follows from the definition of the deviation $\sigma_i^*(v_i)$. Next, integrate out over v_i and exchange v_i and w_i (which are i.i.d. draws from F_i) to obtain

$$\mathbf{E}_{\mathbf{v} \sim \mathbf{F}}[u_i(\sigma(\mathbf{v}))] \geq \mathbf{E}_{\mathbf{v}, \mathbf{w} \sim \mathbf{F}}[u_i(\mathbf{b}_i^*(\mathbf{w}), \sigma_{-i}(\mathbf{v}_{-i}))],$$

with the exchange justified by the assumption that $\mathbf{F}$ is a product distribution (as you should check). Summing over all bidders i, using linearity of expectation, and applying the smooth-type inequality (8) gives

$$\begin{aligned}
\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\sum_{i=1}^{n} u_i(\sigma(\mathbf{v}))\right] &\geq \mathbf{E}_{\mathbf{v},\mathbf{w}\sim\mathbf{F}}\left[\sum_{i=1}^{n} u_i(\mathbf{b}_i^*(\mathbf{w}),\sigma_{-i}(\mathbf{v}_{-i}))\right] \\
&\geq \mathbf{E}_{\mathbf{v},\mathbf{w}\sim\mathbf{F}}\left[\text{OPT welfare}(\mathbf{w}) - \sum_{i=1}^{n}\sum_{j\in S_i(\mathbf{b})}\sigma_i(v_i)_j\right] \\
&= \mathbf{E}_{\mathbf{w}\sim\mathbf{F}}[\text{OPT welfare}(\mathbf{w})] - \sum_{i=1}^{n}\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\sum_{j\in S_i(\mathbf{b})}\sigma_i(v_i)_j\right]
\end{aligned}$$

The last term looks like a jumble, but the ith summand is simply the expected sum of i's bids on the bundle it wins in the equilibrium $\mathbf{b}$. It is now clear what "no overbidding" condition we need to control this term, namely that it is bounded above by the expected welfare bidder i contributes in this equilibrium:

$$\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\sum_{j\in S_i(\mathbf{b})}\sigma_i(v_i)_j\right] \leq \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[v_i(S_i(\mathbf{b}))]\,; \tag{9}$$

Inequality (9) is the WNO condition from last lecture, averaged over valuation profiles $\mathbf{v}\sim\mathbf{F}$.

Assembling our inequalities completes the proof:

$$\begin{aligned}
\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\text{welfare}(\sigma(\mathbf{v}))] &\geq \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\sum_{i=1}^{n} u_i(\sigma(\mathbf{v}))\right] &(10)\\
&\geq \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\text{OPT welfare}(\mathbf{v})] - \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\left[\sum_{i=1}^{n}\sum_{j\in S_i(\mathbf{b})}\sigma_i(v_i)_j\right] &(11)\\
&\geq \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\text{OPT welfare}(\mathbf{v})] - \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}\Bigg[\underbrace{\sum_{i=1}^{n} v_i(S_i(\mathbf{b}))}_{\text{welfare}(\sigma(\mathbf{v}))}\Bigg], &(12)
\end{aligned}$$

where inequality (10) follows from the nonnegativity of prices, inequality (11) from the preceding derivation, and inequality (12) from the no overbidding condition 9. Rearranging completes the proof. ■

2.4 An Extension Theorem for the Bayes-Nash POA

In CS364A we gave an extension theorem that, assuming a (λ,μ)-smoothness condition — a sufficient condition for the POA of pure Nash equilibria to be at least $\frac{\lambda}{1+\mu}$ — yields a bound of $\frac{\lambda}{1+\mu}$ on the POA of coarse correlated equilibria (equivalently, no-regret sequences).

The proof of Theorem 2.1 can be viewed as an extension theorem for Bayes-Nash equilibria. Consider an arbitrary auction game that satisfies the following (λ, μ)-smoothness condition: for every valuation profile $\mathbf{v}$, there exist deviations $\mathbf{b}_1^*, \ldots, \mathbf{b}_n^*$ such that, for every bid vector $\mathbf{b}$,

$$\sum_{i=1}^{n} u_i(\mathbf{b}_i^*, \mathbf{b}_{-i}) \geq \lambda \cdot \text{OPT welfare}(\mathbf{v}) - \mu \sum_{i=1}^{n} p_i(\mathbf{b}),$$

where $p_i(\mathbf{b})$ is the total payment made by bidder i in the outcome $\mathbf{b}$. This is a sufficient condition for the POA of (full-information) pure Nash equilibria to be at least $\frac{\lambda}{1+\mu}$, provided the payment of each bidder i is at most its welfare (WNO). The doppelganger trick in the proof of Theorem 2.1 shows that, for an arbitrary product prior distribution, it is also a sufficient condition for the POA of Bayes-Nash equilibria that satisfy WNO on average to be at least $\frac{\lambda}{1+\mu}$. We will see two more extension theorems for Bayes-Nash equilibria in the next lecture.

References

[1] G. Christodoulou, A. Kovács, and M. Schapira. Bayesian combinatorial auctions. In *35th International Colloqium on Automata, Languages and Programming (ICALP)*, pages 820–832, 2008.

[2] T. Roughgarden. The price of anarchy in games of incomplete information. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, pages 862–879, 2012.