# CS364B: Frontiers in Mechanism Design
# Lecture #12: Bayesian Incentive-Compatibility *

Tim Roughgarden†

February 19, 2014

## 1 The Big Picture

The last several lectures brought us to the research frontier of the design of polynomial-time DSIC mechanisms for NP-hard welfare-maximization problems. The complexity of the state-of-the-art DSIC mechanisms in some scenarios and the outright non-existence of good DSIC mechanisms in other scenarios motivate relaxing the incentive-compatibility guarantee.

What is a meaningful relaxation of the DSIC guarantee? How do we model the behavior of players that have no dominant strategy? Full-information concepts like Nash equilibria are not well suited for auction settings, where players' generally don't know each others valuations. Last lecture presented an interesting relaxation, implementation in undominated strategies, but not many possibility results are known for this concept. This lecture introduces the more traditional notion of Bayesian incentive-compatibility. The idea is that a player acts to maximize its utility given what it knows, using a prior distribution to average over what it doesn't know.

## 2 Bayes-Nash Equilibria

We begin by defining Bayesian games and their equilibria. We first recall from Lecture #1 the basic notions of a game of incomplete information. Each player $i$ has a *type space* $T_i$. A player's type encodes all of its preference information – that is, given an outcome and a player's type, its utility is uniquely defined. For this course, where we concentrate on auction settings, we use the terms "type" and "valuation" interchangeably. For us, a canonical example of a type space would be something like the set of all monotone submodular valuations on an item set $U$.



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

Each player also has an action space $A_i$. In mechanism design, $A_i$ is the various ways in which a bidder $i$ can interact with a mechanism. In a direct revelation mechanism, $A_i = T_i$. In auctions, this means that each player submits a (possibly false) valuation. Most but not all of the auctions that we've studied have been direct-revelation. The exceptions were the ascending auctions that we studied in Part I, where the actions can be history-dependent and hence action spaces are generally much larger than the type spaces. We emphasized this point in Lecture #1, to explain why EPIC implementations are generally strictly harder to come by than DSIC implementations. In Part IV of the course, motivated by auction simplicity, we'll study some auction formats where the action space is much *smaller* than the type space.

Recall that a *strategy* in a game of incomplete information is a function $\sigma_i : T_i \to A_i$ from types to actions. That is, a strategy specifies how a player acts given what it wants. In an auction, $\sigma_i$ is a bidding strategy — how a player bids, given its valuation. Direct revelation and sincere bidding are simple examples of strategies; so is "always bid half your valuation." A strategy can also be randomized, so the range of $\sigma_i$ is more properly the set of distributions over $A_i$.

When a player has no dominant strategy, its utility-maximizing action depends on what other players are doing. In full-information games, the Nash equilibrium concept treats the equilibrium profile as common knowledge, so that all players know that they (and others) are playing a best response. But the whole point of games of incomplete information is to allow uncertainty in what players know about each other. The Bayes-Nash equilibrium concept hews as closely as possible to the Nash equilibrium concept, subject to uncertainty in players' types, in the following sense. When a player $i$ reasons about what its best response is, it makes the following assumptions:

1. Player $i$ is certain about its own type $t_i$.
2. Player $i$ is certain about a *prior distribution* $\mathbf{F}$ on $\mathbf{t}$, and that $\mathbf{t} \sim \mathbf{F}$. Thus, $\mathbf{t}_{-i}$ is drawn from the posterior of $F$ on $T_{-i}$ given $t_i$.
3. Player $i$ is certain that the strategies $\sigma_{-i}$ are employed by the other players.

The first and third assumptions are also present in the usual interpretation of Nash equilibria — players know their own payoffs and the strategies chosen by others. The second assumption reduces to the full-information case when $\mathbf{F}$ is a point mass. With a general prior $\mathbf{F}$, bidders reason about the uncertainty in other players' preferences by taken expectations with respect to the commonly known prior.

Keep in mind that a bidder $i$ cares about the unknown types of others $\mathbf{t}_{-i}$ only inasmuch as they determine their actions.[1] The known prior $\mathbf{F}$ and strategies $\sigma_{-i}$ induce a distribution over the actions $\mathbf{a}_{-i}$ of others. This action distribution is something that player $i$ can reason about directly, by comparing the expected utilities achieved by its own actions.

[1] The assumption that bidder $i$'s utility depends only on others actions and not on their types is called the *private value* model. There is plenty of well-motivated work on *interdependent values*, where a players' utility also depends directly on others' types, but this is outside the scope of this course.

**Definition 2.1** A *Bayes-Nash equiibrium (BNE)* of a game of incomplete information is a strateagy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ such that, for every player $i$ and possible type $t_i$ of $i$, $\sigma_i(t_i)$ maximizes $i$'s expected utility $\mathbf{E}_{\mathbf{t}_{-i} \sim \mathbf{F}_{-i}|t_i}[u_i(a_i, \sigma_{-i}(\mathbf{v}_{-i}))]$ over all possible actions, where the expectation is conditional on $t_i$.

In words, every player always randomizes only over actions that maximize its expected utility with respect to the distribution induced on others' actions by $F_{v_i}$ and $\sigma_{-i}$.

We next contrast BNE with two more stringent equilibrium notions that we've studied previously. Recall from Lecture #1 that a strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ is an *ex post Nash equilibrium (EPNE)* of $\sigma_i(v_i)$ is a best response to $\sigma_{-i}(\mathbf{v}_{-i})$ for *every* $\mathbf{v}_{-i}$, not just on average with respect to some distribution over $\sigma_{-i}(\mathbf{v}_{-i})$. Evidently, an EPNE is a BNE with respect to every possible prior distribution. A strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ is a *dominant-strategy equilibrium (DSE)* if $\sigma_i(v_i)$ is a best response to every *action* profile $\mathbf{a}_{-i}$, where $\mathbf{a}_{-i}$ has the form $\sigma_{-i}(\mathbf{v}_{-i})$ or not. Clearly, every DSE is a EPNE; as we saw in Lecture #1, the converse need not hold.

# 3 Example: First-Price Auctions

We illustrate the Bayes-Nash equilibrium concept in a simple example, Consider a first-price auction, where the highest bidder wins and pays its bid. Suppose there are two bidders, with valuations drawn i.i.d. from the uniform distribution on $[0, 1]$. We claim that the strategy profile in which each bidder always bids half its value — $(\sigma_1(v_1) = v_1/2, \sigma_2(v_2) = v_2/2)$ — is a Bayes-Nash equilibrium.

To verify the claim, consider the first bidder; the same argument applies to the second. Suppose the first bidder's valuation is $v_1$. The first bidder reasons about its action under the assumption that (i) $v_2$ is drawn uniformly from $[0, 1]$ and (ii) the second bidder bids half its value. Thus, the first bidder assumes that the other bid is uniformly distributed between 0 and $\frac{1}{2}$. This means that its expected utility from bidding $b_1$ is

$$\underbrace{(v_1 - b_1)}_{\downarrow \text{ in } b_1} \cdot \underbrace{\mathbf{Pr}_{v_2}[\text{bidder 1 wins}]}_{\uparrow \text{ in } b_1}.$$

We also have

$$\mathbf{Pr}_{v_2}[\text{bidder 1 wins}] = \mathbf{Pr}_{v_2}[b_1 > b_2] = \mathbf{Pr}_{v_2}[v_2 < 2b_1] = \min\{2b_1, 1\}. \tag{1}$$

The expected utility of the first bidder is therefore $\min\{2b_1, 1\} \cdot (v_1 - b_1)$, and the unique maximizer is $b_1 = \frac{v_1}{2}$.[2]

[2]As an aside, note that the expected revenue in this BNE is $\mathbf{E}_{v_1,v_2}\left[\max\{\frac{v_1}{2}, \frac{v_2}{2}\}\right] = \frac{1}{3}$. The expected revenue of the second-price auction in this setting is $\mathbf{E}_{v_1,v_2}[\min\{v_1, v_2\}] = \frac{1}{3}$ — exactly the same! This is a special case of *revenue equivalence*, which states two BNE that induce the same allocation rule — in this case, the welfare-maximizing rule — generate the same expected revenue.

As the derivation in (1) makes clear, the BNE of the first-price auction depends on the details of the setup, including the number $n$ of bidders and the prior $\mathbf{F}$ (unlike a second-price auction).[3] With $n$ bidders with valuations drawn i.i.d. from the uniform distribution on $[0, 1]$, the BNE is $\sigma_i(v_i) = \frac{n-1}{n} v_i$ (see Exercises). Thus, the amount of bid shading at equilibrium decreases as the competition increases.[4]

# 4 Bayesian Incentive-Compatible (BIC) Mechanisms

The definition of a Bayes-Nash equilibrium (Definition 2.1) demands less than the equilibrium concepts we've been focusing on, ex post Nash equilibria and dominant-strategy equilibria. There is hope that we accomplish things with BNE that are impossible with EPNE and DSE — and this is indeed the case, as we show in the next lecture.

**New, Weaker Goal:** Given a valuation distribution $\mathbf{F}$, design a simple and/or polynomial-time mechanism $M$ such that there is a Bayes-Nash equilibrium $(\sigma_1, \ldots, \sigma_n)$ that has near-optimal expected welfare:

$$\mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\text{welfare of } M(\sigma(\mathbf{v}))] \geq \alpha \cdot \mathbf{E}_{\mathbf{v}\sim\mathbf{F}}[\text{OPT welfare}(\mathbf{v})],$$

with $\alpha$ as close to 1 as possible.

In Lecture #4 of CS364A we discussed the *Revelation Principle* for mechanisms with dominant strategies, which used a simulation argument to reduce the problem of designing a good mechanism with dominant strategies to that of designing a good (direct-revelation) DSIC mechanism. The same simulation argument applies to mechanisms with a good Bayes-Nash equilibrium. In more detail, consider a valuation distribution $\mathbf{F}$, a mechanism $M$, and a Bayes-Nash equilibrium $\sigma = (\sigma_1, \ldots, \sigma_n)$. Define a mechanism $M'$ as follows (Figure **??**): $M'$ takes as input reported valuations $\mathbf{v}$, feeds the input $\sigma(\mathbf{v})$ into $M$, and returns the output of $M$. Because $\sigma$ is a BNE of $M$, truthful reporting is a BNE of $M/$. The outcome distribution in the latter BNE is obviously the same as in the former. For example, if we apply the Revelation Principle to the BNE in the two-bidder first-price auction that we identified in Section 3, then we obtain the auction that awards the item to the highest bidder at a price equal to half its bid. When the two bidders' valuations are drawn i.i.d. from the uniform distribution on $[0, 1]$, truthful bidding is a BNE of this auction.

The Revelation Principle implies that, for our new mechanism design goal, we can restrict our search to direct-revelation mechanisms in which truthful reporting is a BNE. Such mechanisms are called *Bayesian incentive compatible (BIC).*

Full disclosure: the Revelation Principle for BIC mechanism design is more problematic than for DSIC mechanism design, in that the simulating mechanism may lose the simplicity of

[3]By saying *the* BNE, we are suggesting that it is unique. This is indeed the case, though it is not easy to prove; see [**?**] for details.

[4]Even more generally, at the BNE of a first price auction with i.i.d. bidder valuations, each bidder $i$ conditions on the event that its bid is the highest and bids the conditional expectation of the second-highest bid (see e.g. [**?**]).

the original mechanism. For example, in Part IV of the course, we analyze simple mechanisms in which the action space is much smaller than the valuation space. Direct revelation is, by design, forbidden in these simple mechanisms. Applying the Revelation Principle to such a mechanism result in a direct-revelation mechanism, defeating the point of the original one. For a second example, recall from Lecture #6 of CS364A the notion of a "prior-indepndent mechanism" — a mechanism whose description does not reference a prior distribution. Most of the mechanisms studied in this course are prior-independent in this sense. The BNE of a prior-independent mechanism can depend on the prior distribution, however — recall our discussion of the BNE of the (prior-independent) first-price single-item auction in Section 3. Apply the Revelation Principle to a prior-dependent BNE of a prior-independent mechanism results in a prior-dependent mechanism — the mechanism has to know how to bid on behalf of the bidders, which require knowing their BNE strategies, which generally depends on the prior. Thus the Revelation Principle transfers the informational burder from the bidders to the seller, which is not always desirable.

# 5 Single-Parameter Setting

Last quarter we only studied DSIC mechanisms. We discussed and applied a satisfyingly complete theory for *single-parameter* problems. In this section we note that the single-parameter theory carries over to BIC mechanism design as well.

Recall that a *single-parameter enviroment* has the following ingredients:

- $n$ bidders, each with a single private valuation $v_i$.
- A feasible set $X$, where each member $\mathbf{x} \in X$ is an $n$-vector $(x_1, \ldots, x_n)$.

We think of $x_i$ as the "amount of stuff" that bidder $i$ receives, and $v_i$ as its value "per unit of stuff." In a direct-revelation mechanism with allocation rule $\mathbf{x}$ and payment rule $\mathbf{p}$, each bidder $i$ strives to maximize its quasi-linear utility $v_i \cdot x_i(\mathbf{b}) - p_i(\mathbf{b})$.

Last quarter, we made good use of *Myerson's Lemma*. The first part of Myerson's Lemma states that an allocation rule $\mathbf{x}$ is (DSIC-)implementable, meaning that there is a payment rule $\mathbf{p}$ such that $(\mathbf{x}, \mathbf{p})$ is a DSIC mechanism, if and only if $\mathbf{x}$ is monotone, meaning that $x_i(z, \mathbf{v}_{-i})$ is nondecreasing in $z$ for every $i$ and $\mathbf{v}_{-i}$. For monotone allocation rules, Myerson's Lemma also gives us a formula for the (unique, up to a constant) payment rule resulting in a DSIC mechanism.

The analog of Myerson's Lemma for BIC mechanism concerns *interim allocation and payment rules*, which are fundamental concepts. For an allocation rule $\mathbf{x}$ and product prior distrubtion $\mathbf{F}$, the interim allocation rule is

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

While an allocation rule takes as input a full valuation profile $\mathbf{v}$, an interim allocation rule takes as input only $i$'s valuation $v_i$ and then averages over $\mathbf{v}_{-i}$ with respect to the prior $\mathbf{F}_{-i}$.[5]

[5] If $\mathbf{F}$ is a correlated valuation distribution, the interim rule averages over the conditional distribution on $\mathbf{v}_{-i}$.

For example, if $x_i(\mathbf{v})$ is always either 0 or 1, then $x_i(v_i)$ is just the probability (over $\mathbf{v}_{-i}$) of winning with a reported valuation of $v_i$. The interim payment rule is defned analogously:

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

The general version of Myerson's Lemma is the following.

**Theorem 5.1 ([?])** *XXX*

The key steps of the proof are the same as in the DSIC case covered last quarter (Lecture #3), and we leave the details as an exercise.

As one would expect, the monotonicity requirement in Theorem 5.1 is only weaker than that for DSIC mechanisms — monotone allocation rules can only induce monotone interim allocation rules. On the hand, the requirement in Theorem 5.1 is non-trivial. For example, the "second-highest bidder wins" allocation rule for a single-item auction does not typically meet the monotonicity condition in Theorem 5.1.