GAME THEORY

Jean-Christophe Spiliotis

Michaelmas Term 2017

Part II

Applied Game Theory:Who Gets What—and Why?

3 Auctions and Mechanisms

Definition. An auction is mathematically pinned down as

$$\mathcal{A} = \langle \mathcal{B}, \pi, \mu \rangle ,$$

where

- $\mathcal{B}$ is the *bidding set*, i.e. the set of feasible bidding profiles: with bidders 1, 2 and 3, an element of $\mathcal{B}$ would be $\mathbf{b} = (b_1, b_2, b_3)$. Bidding set and bidding profile are analogous to strategy set and strategy profile, respectively, in the general definition of a game.
- π is the *allocation rule*: for each bidder i, $\pi_i(\mathbf{b})$ is the probability that i wins the auction given the bidding profile $\mathbf{b}$.
- μ is the *payment rule*: $\mu_i(\mathbf{b})$ is i's expected payment in the auction.

First principles. Let $U_i(v_i)$ the expected payoff from an individual with valuation v_i. It is given by the expression $U_i(v_i) = P_i(v_i)v_i - Q_i(v_i)$, where $P_i(v_i)$ is the probability that an individual wins the auction as a function of their valuation v_i, and $Q_i(v_i)$ is the expected payment by an individual with valuation v_i. Notice the use of i-subscripts, since bidders' choices of bid may depend on individual preferences (e.g. risk aversion).

In what follows, we make the assumption of *independent private valuations* (IPV): only i observes v_i (i.e. valuations are private) and v_i depends only on i, not on anything related to other bidders (i.e. they are independent). Other possible assumptions include common or affiliated valuations, and will not be considered in this course.

Incentive-compatible direct auctions. Central to mechanism design is the notion of incentive compatibility. (Remember that an auction *is* a mechanism, more specifically a selling mechanism, and refer to section **3.3** for more general details.)

Suppose that you are a bidder in an auction; for ease of exposition, and *without loss of generality*, consider a sealed-bid first-price auction. However, this auction operates under somewhat unusual rules. There is a *proxy bidder* who announces a function $\beta_i(\cdot)$ for all bidders $i \in N$ which will determine the value of each i's bid:

1. The bidding functions $\{\beta_i(\cdot)\}_{i\in N}$ are revealed to all bidders
2. Instead of writing your bid in the sealed envelope, you report a valuation w_i
3. The proxy bidder observes $\{w_i\}_{i\in N}$ and computes the associated bidding profile $\widetilde{\mathbf{b}} = \big(\beta_1(w_1), \dots, \beta_n(w_n)\big)$
4. The auction takes place as usual with bidding profile $\widetilde{\mathbf{b}}$. $\pi(\widetilde{\mathbf{b}})$ and $\mu(\widetilde{\mathbf{b}})$ are computed—in the case of a sealed-bid first-price auction, the highest bid wins and the buyer pays her bid.

Such an auction is said to be *direct*, insofar as bidders report directly a valuation instead of their bid. However, it is not necessarily *truthful*: nothing stops you from misreporting your type—indeed, it is often lucrative to con a mechanism by pretending to be someone you aren't! Fortunately, Myerson's lemma provides necessary and sufficient conditions for the case when it is **not** profitable to misreport your type:

Lemma (Myerson, 1981)**.** *A direct mechanism is incentive compatible if and only if, for each $i \in N$,*

1. *$P_i(\cdot)$ is nondecreasing*
2. $U_i(v_i) = U_i(\underline{v}) + \displaystyle\int_{\underline{v}}^{v_i} P_i(x)dx$

Proof. This lemma is twofold, therefore we have to show that *(i)* incentive compatibility implies points 1. and 2., and *(ii)* points 1. and 2. imply incentive compatibility.

In order to prove *(i)*, let's begin by stating what it means for a direct auction to be incentive compatible. Suppose individual i has valuation v_i; she can misreport her valuation by stating $w_i > v_i$. (This is *without loss of generality* as the proof with $w_i < v_i$ is the mirror image.) Incentive compatibility requires that she cannot make herself better-off by misreporting. Recall that when i's valuation is v, her payoff from reporting x is $U_i(x, v) = P_i(x)v - Q_i(x)$. The incentive-compatibility constraint on an individual with valuation v_i is therefore:

$U_i(v_i) \geq P_i(w_i)v_i - Q_i(w_i)$; adding and substracting $P_i(w_i)w_i$, this becomes

$U_i(v_i) \geq P_i(w_i)w_i + P_i(w_i)(v_i - w_i) - Q_i(w_i)$ where we recognise $U_i(w_i) = P_i(w_i)w_i - Q_i(w_i)$ which would be i's payoff from bidding w_i if her true valuation was w_i. Therefore,

$$U_i(v_i) \geq U_i(w_i) + P(w_i)(v_i - w_i) \tag{1}$$

Similarly, starting from the incentive-compatibility constraint on an individual with valuation w_i, $U_i(w_i) \geq P_i(v_i)w_i - Q_i(v_i)$, we can obtain the symmetrical result

$$U_i(w_i) \geq U_i(v_i) + P(v_i)(w_i - v_i) \tag{2}$$

Combining (1) and (2), we can rearrange to

$$P_i(w_i) \geq \frac{U_i(w_i) - U_i(v_i)}{w_i - v_i} \geq P_i(v_i)$$

It is immediate that $P_i(w_i) \geq P_i(v_i)\ \forall w_i > v_i$, which proves 1. Now, let's see what happens to this inequality as $w_i \longrightarrow v_i$. First, the term between the two inequalities is of the form

$$\frac{f(x+h) - f(x)}{h},$$

where $h := w_i - v_i$ and $x := v_i$. As $h \longrightarrow 0$, this converges to $f'(x)$. (This is the definition of a derivative.) Therefore, the middle term converges to $U_i'(v_i)$ as $w_i \longrightarrow v_i$. Second, notice that $P_i(w_i) \longrightarrow P_i(v_i)$, so the middle term is "squeezed" between an upper bound and a lower bound which converge to one another. This means that the inequalities will hold with equality, and yield

$$U_i'(v_i) = P_i(v_i)$$

Finally, in order to recover an expression for $U_i(v_i)$, we use the

Theorem (Fundamental Theorem of Calculus)**.** *If the function F is continuously differentiable over $[a-\varepsilon, b+\varepsilon]$ for any arbitrarily small, positive ε, and if F differentiates to f, then $\int_a^b f(x)dx = F(b) - F(a)$.*

To apply this theorem here, we need to assume that $U_i(v_i)$ is continuously differentiable over $[\underline{v}, \bar{v}]$, which is sensible. However, we also require that $v_i \in (\underline{v}, \bar{v})$—that is, our result will not hold to individuals with either the highest or the lowest possible valuation. (This is not an issue as we have other ways to compute the payoff for such individuals. For the sake of this proof, however, we gloss over this technicality.) Under the above conditions, the FTC applies and

$$U_i(v_i) - U_i(\underline{v}) = \int_{\underline{v}}^{v_i} U_i'(x)dx$$

$$\Leftrightarrow U_i(v_i) = \int_{\underline{v}}^{v_i} P_i(x)dx + U_i(\underline{v}),$$

which proves 2. □

Solving the four standard auctions (and more!)

First, notice that the sealed-bid first-price and the descending auctions are **strategically equivalent**. This means that you can think of them as the exact same game, where the same strategies will be available to each player and therefore the outcome will also be the same. Here, the winner pays her bid and no additional information is made available throughout the course of the auction.

The sealed-bid second-price and the ascending auctions are also equivalent, but to a lesser extent and *only* with independent private valuations. To clarify the difference, note that unlike in any of the other auction formats, the ascending auction forces bidders to release additional information as the game unfolds; for instance, a strategy available to a bidder in an ascending auction could be "drop out of the auction when half the other bidders dropped out, even if my valuation hasn't been reached yet." Such a strategy is not available to a bidder in the sealed-bid second-price auction; therefore, these auctions are not strategically equivalent. However, in effect the highest bidder pays the second-highest bidder's bid, so these two auctions will be **outcome equivalent**, a weaker equivalence concept.

Now, a solution concept for an auction means the exact same thing that for any other game. We seek to predict the outcome of the auction in terms of players' strategies (here, bids); these strategies are themselves contingent on each player's type (here, valuation). Here, bidder i's equilibrium bid is denoted by the expression $\beta_i(v_i)$.

In fact, auctions with independent private valuations yield a symmetric equilibrium—that is, the function $\beta_i(\cdot)$ is the same for each i. This means that once the auction format is chosen, we can find a function $\beta(v_i)$ that will determine bidder i's equilibrium bid.

Sealed-bid second-price and ascending auctions

For each player, bidding their own valuation is a dominant strategy. There is therefore a dominant-strategy equilibrium in which $\beta(v_i) = v_i$ for all i.

Sealed-bid first-price and descending auctions

Here, there is no dominant-strategy equilibrium. Therefore, the equilibrium concept that applies here is a Nash equilibrium, more specifically Bayes-Nash as an auction is a Bayesian games (since players have types/valuations). The equilibrium combines what I call the "first-principles" expression (FP) and Myerson's lemma (ML), which holds in equilibrium of any incentive-compatible mechanism by the Revelation Principle (see the Krishna textbook from the reading list). For these specific auctions, FP can be rewritten as

$$U_i(v_i) = P_i(v_i)v_i - Q_i(v_i) = P_i(v_i)v_i - P_i(v_i)\beta(v_i) = [F(v_i)]^n v_i - [F(v_i)]^n \beta(v_i)$$

where $F(k)$ is the distribution of valuations among the bidders. Similarly, ML becomes

$$U_i(v_i) = U_i(\underline{v}) + \int_{\underline{v}}^{v_i} P_i(x)dx = 0 + \int_{\underline{v}}^{v_i} [F(x)]^n dx$$

where $U_i(\underline{v}) = 0$: a bidder with the lowest possible valuation can earn no utility from the auction. Combining FP and ML and rearranging yields an expression for the BNE bid $\beta(v_i)$:

$$[F(v_i)]^n v_i - [F(v_i)]^n \beta(v_i) = \int_{\underline{v}}^{v_i} [F(x)]^n dx$$

$$\Leftrightarrow \beta(v_i) = v_i - \frac{\int_{\underline{v}}^{v_i} [F(x)]^n dx}{[F(v_i)]^n}$$