Bayesian Games and Auctions

Mihai Manea

MIT

Games of Incomplete Information

- Incomplete information: players are uncertain about the payoffs or types of others
- Often a player's type defined by his payoff function.
- More generally, types embody any private information relevant to players' decision making. . . may include a player's beliefs about other players' payoffs, his beliefs about what other players believe his beliefs are, and so on.
- Modeling incomplete information about higher order beliefs is intractable. Assume that each player's uncertainty is solely about payoffs.

Bayesian Game

A Bayesian game is a list $\mathcal{B} = (N, S, \Theta, u, p)$ where

- $N = \{1, 2, \ldots, n\}$: finite set of players
- S_i: set of pure strategies of player i; $S = S_1 \times \ldots \times S_n$
- Θ_i: set of types of player i; $\Theta = \Theta_1 \times \ldots \times \Theta_n$
- $u_i : \Theta \times S \to \mathbb{R}$ is the payoff function of player i; $u = (u_1, \ldots, u_n)$
- $p \in \Delta(\Theta)$: common prior

Often assume Θ is finite and marginal $p(\theta_i)$ is positive for each type θ_i.

Strategies of player i in $\mathcal{B}$ are mappings $s_i : \Theta_i \to S_i$ (measurable when Θ_i is uncountable).

First Price Auction

- One object is up for sale.
- Value θ_i of player $i \in N$ for the object is uniformly distributed in $\Theta_i = [0, 1]$, independently across players, i.e.,

$$p(\theta_i \leq \tilde{\theta}_i, \forall i \in N) = \prod_{i \in N} \tilde{\theta}_i, \forall \theta_i \in [0, 1], i \in N.$$

- Each player i submits a bid $s_i \in S_i = [0, \infty)$.
- The player with the highest bid wins the object (ties broken randomly) and pays his bid. Payoffs:

$$u_i(\theta, s) = \begin{cases} \frac{\theta_i - s_i}{|\{j \in N | s_i = s_j\}|} & \text{if } s_i \geq s_j, \forall j \in N \\ 0 & \text{otherwise.} \end{cases}$$

An Exchange Game

- Player $i = 1, 2$ receives a ticket on which there is a number from a finite set $\Theta_i \subset [0, 1]\ldots$ prize player i may receive.
- The two prizes are independently distributed, with the value on i's ticket distributed according to F_i.
- Each player is asked independently and simultaneously whether he wants to exchange his prize for the other player's prize: $S_i = \{agree, \ disagree\}$.
- If both players agree then the prizes are exchanged; otherwise each player receives his own prize. Payoffs:

$$u_i(\theta, s) = \begin{cases} \theta_{3-i} & \text{if } s_1 = s_2 = agree \\ \theta_i & \text{otherwise.} \end{cases}$$

Ex-Ante Representation

In the ex ante representation $G(\mathcal{B})$ of the Bayesian game $\mathcal{B}$ player i has strategies $(s_i(\theta_i))_{\theta_i \in \Theta_i} \in S_i^{\Theta_i}$—his strategies are functions from types to strategies in $\mathcal{B}$—and utility function U_i given by

$$U_i\Big(\big(\big(s_i(\theta_i)\big)_{\theta_i \in \Theta_i}\big)_{i \in N}\Big) = \mathbb{E}_p(u_i(\theta, s_1(\theta_1), \ldots, s_n(\theta_n))).$$

Interim Representation

The interim representation $IG(\mathcal{B})$ of the Bayesian game $\mathcal{B}$ has player set $\cup_i \Theta_i$. The strategy space of player θ_i is S_i. A strategy profile $(s_{\theta_i})_{i \in N, \theta_i \in \Theta_i}$ yields utility

$$U_{\theta_i}((s_{\theta_i})_{i \in N, \theta_i \in \Theta_i}) = \mathbb{E}_p(u_i(\theta, s_{\theta_1}, \ldots, s_{\theta_n})|\theta_i)$$

for player θ_i. Need $p(\theta_i) > 0$…

Bayesian Nash Equilibrium

Definition 1

In a Bayesian game $\mathcal{B} = (N, S, \Theta, u, p)$, a strategy profile $s : \Theta \to S$ is a *Bayesian Nash equilibrium* (BNE) if it corresponds to a Nash equilibrium of $IG(\mathcal{B})$, i.e., for every $i \in N, \theta_i \in \Theta_i$

$$E_{p(\cdot|\theta_i)}\left[u_i\left(\theta, s_i\left(\theta_i\right), s_{-i}\left(\theta_{-i}\right)\right)\right] \geq E_{p(\cdot|\theta_i)}\left[u_i\left(\theta, s_i', s_{-i}\left(\theta_{-i}\right)\right)\right], \forall s_i' \in S_i.$$

Interim rather than ex ante definition preferred since in models with a continuum of types the ex ante game has many spurious equilibria that differ on probability zero sets of types.

Connections to the Complete Information Games

When i plays a best-response type by type, he also optimizes ex-ante payoffs (for any probability distribution over Θ_i). Therefore, a BNE of $\mathcal{B}$ is also a Nash equilibrium of the ex-ante game $G(\mathcal{B})$.

$BNE(\mathcal{B})$: Bayesian Nash equilibria of bayesian game $\mathcal{B}$
$NE(G)$: Nash equilibria of normal-form game G

Proposition 1

For any Bayesian game $\mathcal{B}$ with a common prior p,

$$BNE(\mathcal{B}) \subseteq NE(G(\mathcal{B})).$$

If $p(\theta_i) > 0$ for all $\theta_i \in \Theta_i$ and $i \in N$, then

$$BNE(\mathcal{B}) = NE(G(\mathcal{B})).$$

Business Partnership

Two business partners work on a joint project.

- Each businessman $i = 1, 2$ can either exert effort ($e_i = 1$) or shirk ($e_i = 0$).
- Each face the same fixed (commonly known) cost for effort $c < 1$.
- Project succeeds if at least one partner puts in effort, fails otherwise.
- Players differ in how much they care about the fate of the project: i has a private, independently distributed type $\theta_i \sim U[0, 1]$ and receives payoff θ_i^2 from success.

Hence player i gets $\theta_i^2 - c$ from working, θ_i^2 from shirking if opponent j works, and 0 if both shirk.

	$e_2 = 1$	$e_2 = 0$
$e_1 = 1$	$\theta_1^2 - c, \theta_2^2 - c$	$\theta_1^2 - c, \theta_2^2$
$e_1 = 0$	$\theta_1^2, \theta_2^2 - c$	$0, 0$

Equilibrium

	$e_2 = 1$	$e_2 = 0$
$e_1 = 1$	$\theta_1^2 - c, \theta_2^2 - c$	$\theta_1^2 - c, \theta_2^2$
$e_1 = 0$	$\theta_1^2, \theta_2^2 - c$	$0, 0$

p_j: probability that j works—sufficient statistic for strategic situation faced by player i

Working is rational for i if $\theta_i^2 - c \geq p_j\theta_i^2 \iff (1 - p_j)\theta_i^2 \geq c$. Thus i must play a threshold strategy: work for

$$\theta_i \geq \theta_i^* := \sqrt{\frac{c}{1 - p_j}}.$$

Since $p_j = Prob(\theta_j \geq \theta_j^*) = 1 - \theta_j^*$, we get

$$\theta_i^* = \sqrt{\frac{c}{\theta_j^*}} = \sqrt{\frac{c}{\sqrt{\frac{c}{\theta_i^*}}}} = \sqrt[4]{c\theta_i^*},$$

so $\theta_i^* = \sqrt[3]{c}$. In equilibrium, $i = 1, 2$ works if $\theta_i \geq \sqrt[3]{c}$ and shirks otherwise.

Auctions

- single good up for sale
- n buyers bidding for the good
- buyer i has value X_i, i.i.d. with distribution F and continuous density $f = F'$; $supp(F) = [0, \omega]$
- i knows only the realization x_i of X_i

Auction Formats

- First-price sealed-bid auction: each buyer submits a single bid (in a sealed envelope) and the highest bidder obtains the good and pays his bid. Equivalent to descending-price (Dutch) auctions.
- Second-price sealed-bid auction: each buyer submits a bid and the highest bidder obtains the good and pays the second highest bid. Equivalent to open ascending-price (English) auctions.

Bidding strategies: $\beta_i : [0, \omega] \to [0, \infty)$

- What are the BNEs in the two auctions?
- Which auction generates higher revenue?

Second-Price Auction

Each bidder i submits a bid b_i, payoffs given by

$$u_i = \begin{cases} x_i - \max_{j \neq i} b_j & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{if } b_i < \max_{j \neq i} b_j \end{cases}$$

Ties broken randomly.

Proposition 2

In a second-price auction, it is a weakly dominant strategy for every player i to bid according to $\beta_i^{II}(x_i) = x_i$.

Second-Price Auction Expected Payments

$Y_1 = \max_{i \neq 1} X_i$: highest value of player 1's opponents, distributed according to G with $G(y) = F(y)^{n-1}$

Expected payment by a bidder with value x is

$$\begin{aligned} m^{II}(x) &= \text{Prob[Win]} \times E[\text{2nd highest bid} \mid x \text{ is the highest bid}] \\ &= \text{Prob[Win]} \times E[\text{2nd highest value} \mid x \text{ is the highest value}] \\ &= G(x) \times E[Y_1 | Y_1 < x] \end{aligned}$$

First-Price Auction

Each bidder i submits a bid b_i, payoffs given by

$$u_i = \begin{cases} x_i - b_i & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{if } b_i < \max_{j \neq i} b_j \end{cases}$$

Ties broken randomly.

Clearly, not optimal/equilibrium to bid own value. Trade-off: higher bids increase the probability of winning but decrease the gains.

Symmetric equilibrium: $\beta_i = \beta$ for all buyers i. Assume β strictly increasing, differentiable.

Optimal Bidding

Suppose bidder 1 has value $X_1 = x$ and considers bidding b. Clearly, $b \leq \beta(\omega)$ and $\beta(0) = 0$.

Bidder 1 wins the auction if $\max_{i \neq 1} \beta(X_i) < b$. Since β is s. increasing, $\max_{i \neq 1} \beta(X_i) = \beta(\max_{i \neq 1} X_i) = \beta(Y_1)$, so 1 wins if $Y_1 < \beta^{-1}(b)$. His expected payoff is

$$G(\beta^{-1}(b)) \times (x - b).$$

$$FOC: \quad \frac{G'(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))}(x - b) - G(\beta^{-1}(b)) = 0$$

$$b = \beta(x) \Rightarrow G(x)\beta'(x) + G'(x)\beta(x) = xg(x) \iff (G(x)\beta(x))' = xg(x)$$

$$\beta(0) = 0 \Rightarrow \beta(x) = \frac{1}{G(x)} \int_0^x yg(y)dy$$
$$= E[Y_1 | Y_1 < x].$$

Equilibrium

Proposition 3

The strategies

$$\beta^I(x) = E[Y_1 | Y_1 < x]$$

constitute a symmetric BNE in the first-price auction.

Proof

We only checked necessary conditions for equilibrium... Check that if all bidders follow strategy β^I then it is optimal for bidder 1 to follow it. Since β^I is increasing and continuous, it cannot be optimal to bid higher than $\beta^I(\omega)$. Suppose bidder 1 with value x bids $b \in [0, \beta^I(\omega)]$. $\exists z, \beta^I(z) = b$. Since bidder 1 wins if $Y_1 < z$, his payoffs are

$$\begin{aligned}
\Pi(b,x) &= G(z)[x - \beta^I(z)] \\
&= G(z)x - G(z)E[Y_1 | Y_1 < z] \\
&= G(z)x - \int_0^z yg(y)dy \\
&= G(z)x - G(z)z + \int_0^z G(y)dy \\
&= G(z)(x - z) + \int_0^z G(y)dy.
\end{aligned}$$

Then

$$\Pi(\beta^I(x), x) - \Pi(\beta^I(z), x) = G(z)(z - x) - \int_x^z G(y)dy \geq 0.$$

Shading

$$\begin{aligned}\beta^I(x) &= \frac{1}{G(x)}\int_0^x yg(y)dy \\ &= x - \int_0^x \frac{G(y)}{G(x)}dy \\ &= x - \int_0^x \left[\frac{F(y)}{F(x)}\right]^{n-1} dy\end{aligned}$$

Shading, the amount by which the bid is lower than the value, is

$$\int_0^x \left[\frac{F(y)}{F(x)}\right]^{n-1} dy.$$

Depends on n, converges to 0 as $n \to \infty$ (competition).

Example with Uniformly Distributed Values

If $F(x) = x$ for $x \in [0, 1]$, then $G(x) = x^{n-1}$ and

$$\beta^I(x) = \frac{n-1}{n}x.$$

Example with Exponentially Distributed Values

If $n = 2$ and $F(x) = 1 - \exp(-\lambda x)$ for $x \in [0, \infty)$ ($\lambda > 0$) then

$$\beta^I(x) = x - \int_0^x \frac{F(y)}{F(x)} dy$$
$$= \frac{1}{\lambda} - \frac{x \exp(-\lambda x)}{1 - \exp(-\lambda x)}$$

Note that $E[X] = 1/\lambda$.

Take $\lambda = 1$. A bidder with value $\$10^6$ will not bid more than $1. Why? Such a bidder has a lot to lose by not bidding higher but the probability of losing is small, $\exp(-10^6)$.

More generally, for $n = 2$,

$$\beta^I(x) = E[Y_1 | Y_1 < x] \leq E[Y_1] = E[X_2].$$

Revenue Comparison

Expected payment in the first-price auction by a bidder with value x is

$$m^{I}(x) = \text{Prob[Win]} \times \text{ Amount bid} = G(x) \times E[Y_1 | Y_1 < x]$$

Recall that

$$m^{II}(x) = G(x) \times E[Y_1 | Y_1 < x],$$

so both auctions yield the same revenue. Special case of the revenue equivalence theorem.

Mechanism Design

An auction is one of many mechanisms a seller can use to sell the good. The price is determined by the competition among buyers according to the rules set out by the seller—the auction format.

The seller could use other methods

- post different prices for each bidder, choose a winner at random
- ask various subsets of bidders to pay their own or others' bids

Options virtually unlimited. . .

Myerson (1981): What is the *optimal mechanism*?

Framework

- single good up for sale, worth 0 to the seller
- buyers: $1, 2, \ldots, n$
- buyers have private values, independently distributed
- buyer i's value X_i distributed according to F_i
- $supp(F_i) = [0, \omega_i] = \mathcal{X}_i$, density $f_i = F_i'$
- i knows only the realization x_i of X_i
- $\mathcal{X} = \prod_{i=1}^{n} \mathcal{X}_i$
- $f(x) = \prod_{i=1}^{n} f_i(x_i)$
- $f_{-i}(x_{-i}) = \prod_{j \neq i} f_j(x_j)$

Mechanisms

A selling mechanism (B, π, μ)

- B_i: set of messages (bids) for buyer i
- allocation rule $\pi : B \to \Delta(N)$
- payment rule $\mu : B \to \mathbb{R}^n$

The allocation rule determines, as a function of all n messages b, the probability $\pi_i(b)$ that i gets the object. Similarly the payment rule specifies a payment $\mu_i(b)$ for each buyer i.

Describe first- and second-price auctions as mechanisms...

Every mechanism induces a game of incomplete information with strategies $\beta_i : \mathcal{X}_i \to B_i$.

Direct Mechanisms

Mechanisms can be complicated, no assumptions on the messages B_i.

Direct mechanism (Q, M)

- $B_i = \mathcal{X}_i$, every buyer is asked to directly report a value
- $Q : \mathcal{X} \to \Delta(N)$ and $M : \mathcal{X} \to \mathbb{R}^n$
  - $Q_i(x)$: probability that i gets the object
  - $M_i(x)$: payment by i

If $\beta_i : \mathcal{X}_i \to \mathcal{X}_i$ with $\beta_i(x_i) = x_i$ constitutes a BNE of the induced game then we say that the direct mechanism has a *truthful* equilibrium or is *incentive compatible*.

The Revelation Principle

Proposition 4

Given a mechanism and an equilibrium for that mechanism, there exists a direct mechanism in which (i) it is an equilibrium for each buyer to report his or her value truthfully and (ii) the outcomes are the same as in the given equilibrium of the original mechanism for every type realization x.

Consider a mechanism (B, π, μ) with an equilibrium β. Define $Q : \mathcal{X} \to \Delta(N)$ and $M : \mathcal{X} \to \mathbb{R}^N$ as follows: $Q(x) = \pi(\beta(x))$ and $M(x) = \mu(\beta(x))$. The direct mechanism (Q, M) asks players to report types and does the "equilibrium computation" for them.

(ii) holds by construction.

To verify (i): if buyer i finds it profitable to report z_i instead of his true value x_i in the direct mechanism (Q, M), then i prefers the message $\beta_i(z_i)$ instead of $\beta_i(x_i)$ in the original mechanism.

Incentive Compatibility

For a direct mechanism (Q, M), define

$$
\begin{aligned}
q_i(z_i) &= \int_{X_{-i}} Q_i(z_i, x_{-i}) f_{-i}(x_{-i}) dx_{-i} \\
m_i(z_i) &= \int_{X_{-i}} M_i(z_i, x_{-i}) f_{-i}(x_{-i}) dx_{-i}
\end{aligned}
$$

Expected payoff of buyer i with value x_i who reports z_i if other buyers report truthfully

$$q_i(z_i)x_i - m_i(z_i)$$

(Q, M) is *incentive compatible* (IC) if

$$U_i(x_i) \equiv q_i(x_i)x_i - m_i(x_i) \geq q_i(z_i)x_i - m_i(z_i), \forall i, x_i, z_i$$

U_i is convex because

$$U_i(x_i) = \max\{q_i(z_i)x_i - m_i(z_i) \mid z_i \in X_i\}.$$

Payoff Formula

Since

$$q_i(x_i)z_i - m_i(x_i) = q_i(x_i)x_i - m_i(x_i) + q_i(x_i)(z_i - x_i)$$
$$= U_i(x_i) + q_i(x_i)(z_i - x_i),$$

IC requires that

$$U_i(z_i) \geq U_i(x_i) + q_i(x_i)(z_i - x_i).$$

Hence

$$q_i(z_i)(z_i - x_i) \geq U_i(z_i) - U_i(x_i) \geq q_i(x_i)(z_i - x_i).$$

For $z_i > x_i$,

$$q_i(z_i) \geq \frac{U_i(z_i) - U_i(x_i)}{z_i - x_i} \geq q_i(x_i),$$

so q_i is increasing. Since U_i is convex, it is differentiable almost everywhere,

$$U_i'(x_i) = q_i(x_i)$$
$$U_i(x_i) = U_i(0) + \int_0^{x_i} q_i(t_i)dt_i$$

Monotonicity Condition

IC implies monotonicity of q_i.

Conversely, a mechanism where U_i satisfies

$$U_i(x_i) = U_i(0) + \int_0^{x_i} q_i(t_i)dt_i$$

with q_i increasing must be incentive compatible. IC condition

$$U_i(z_i) - U_i(x_i) \geq q_i(x_i)(z_i - x_i)$$

boils down to

$$\int_{x_i}^{z_i} q_i(t_i)dt_i \geq q_i(x_i)(z_i - x_i).$$

Revenue Equivalence

$$U_i(x_i) = U_i(0) + \int_0^{x_i} q_i(t_i)dt_i$$

Theorem 1

If the direct mechanism (Q,M) is incentive compatible, then for all i and x_i,

$$m_i(x_i) = m_i(0) + q_i(x_i)x_i - \int_0^{x_i} q_i(t_i)dt_i.$$

The expected payments in any two incentive compatible mechanisms with the same allocation rule are equivalent up to a constant.

$$U_i(x_i) = q_i(x_i)x_i - m_i(x_i) = U_i(0) + \int_0^{x_i} q_i(t_i)dt_i$$

First-Price Auction Revisited

n symmetric buyers

Assuming a symmetric monotone equilibrium β in first-price auction, the highest value buyer obtains the good. Same allocation Q as in the equilibrium of the second-price auction. Buyers with value 0 bid 0, so $U_i(0) = 0$ in both auctions. By Theorem 1,

$$m^I(x) = m^{II}(x).$$

Since

$$\begin{aligned} m^I(x) &= G(x) \times \beta(x) \\ m^{II}(x) &= G(x) \times E[Y_1 | Y_1 < x], \end{aligned}$$

we obtain $\beta(x) = E[Y_1 | Y_1 < x]$.

All-Pay Auction

n symmetric buyers. Highest bidder receives the good, but all buyers have to pay their bid (as in lobbying).

Assuming a symmetric monotone equilibrium β in the all-pay auction, the highest value buyer obtains the good. Same allocation Q as in the equilibrium of the second-price auction. Buyers with value 0 bid 0, so $U_i(0) = 0$ in both auctions. By Theorem 1,

$$m^{all-pay}(x) = m^{II}(x) = G(x) \times E[Y_1 | Y_1 < x].$$

Since $m^{all-pay}(x) = \beta(x)$,

$$\beta(x) = G(x) \times E[Y_1 | Y_1 < x].$$

Underbidding compared to first- and second-price auctions.

Can use revenue equivalence with the second-price auction to derive equilibrium in any auction where we expect efficient allocation.

Individual Rationality

A mechanism is individually rational (IR) if $U_i(x_i) \geq 0$ for all $x_i \in \mathcal{X}_i$.

$$U_i(x_i) = U_i(0) + \int_0^{x_i} q_i(t_i)dt_i \geq 0, \forall x_i \in \mathcal{X}_i \iff U_i(0) = -m_i(0) \geq 0$$

Expected Revenue

In a direct mechanism (Q, M), the expected revenue of the seller is

$$E[R] = \sum_{i=1}^{n} E[m_i(X_i)].$$

Substitute the formula for m_i,

$$\begin{aligned} E[m_i(X_i)] &= \int_0^{\omega_i} m_i(x_i) f_i(x_i) dx_i \\ &= m_i(0) + \int_0^{\omega_i} q_i(x_i) x_i f_i(x_i) dx_i - \int_0^{\omega_i} \int_0^{x_i} q_i(t_i) f_i(x_i) dt_i d x_i. \end{aligned}$$

Interchanging the order of integration,

$$\int_0^{\omega_i} \int_0^{x_i} q_i(t_i) f_i(x_i) dt_i dx_i = \int_0^{\omega_i} \int_{t_i}^{\omega_i} q_i(t_i) f_i(x_i) dx_i d t_i = \int_0^{\omega_i} (1 - F_i(t_i)) q_i(t_i) dt_i$$

$$\begin{aligned} E[m_i(X_i)] &= m_i(0) + \int_0^{\omega_i} \left(x_i - \frac{1 - F_i(x_i)}{f_i(x_i)} \right) q_i(x_i) f_i(x_i) dx_i \\ &= m_i(0) + \int_X \left(x_i - \frac{1 - F_i(x_i)}{f_i(x_i)} \right) Q_i(x) f(x) dx \end{aligned}$$

Optimal Mechanism

The seller's objective is to maximize revenue,

$$\sum_{i=1}^{n} m_i(0) + \int_X \sum_{i=1}^{n} \left(x_i - \frac{1 - F_i(x_i)}{f_i(x_i)} \right) Q_i(x) f(x) dx$$

subject to IC and IR. IC is equivalent to q_i being increasing for every i and IR to $m_i(0) \leq 0$. Clearly, need to set $m_i(0) = 0$.

Maximize

$$\int_X \sum_{i=1}^{n} \psi_i(x_i) Q_i(x) f(x) dx \quad \text{where} \quad \psi_i(x_i) := x_i - \frac{1 - F_i(x_i)}{f_i(x_i)}$$

subject to q_i being increasing for every i.

$\psi_i(x_i)$: *virtual value* of player i with type x_i

Regularity condition: assume ψ_i is s. increasing for every i

Optimal Solution

Ignoring the q_i monotonicity condition, maximize for every x

$$\sum_{i=1}^{n} \psi_i(x_i) Q_i(x).$$

Set

$$Q_i(x) > 0 \iff \psi_i(x_i) = \max_{j \in N} \psi_j(x_j) \geq 0.$$

To obtain $m_i(x_i) = m_i(0) + q_i(x_i)x_i - \int_0^{x_i} q_i(t_i)dt_i$, define

$$M_i(x) = Q_i(x)x_i - \int_0^{x_i} Q_i(z_i, x_{-i})dz_i.$$

(Q, M) is an optimal mechanism. Only need to check implied q_i is increasing. If $z_i < x_i$, regularity implies $\psi_i(z_i) < \psi_i(x_i)$, which means that $Q_i(z_i, x_{-i}) \leq Q_i(x_i, x_{-i})$.

$$E[R] = E[\max(\psi_1(X_1), \psi_2(X_2), ..., \psi_n(X_n), 0)]$$

Optimal Auction

Smallest value needed for i to win against opponent types x_{-i}:

$$y_i(x_{-i}) = \inf\{z_i : \psi_i(z_i) \geq 0 \text{ and } \psi_i(z_i) \geq \psi_j(x_j), \forall j \neq i\}$$

In the optimal mechanism,

$$Q_i(z_i, x_{-i}) = \begin{cases} 1 & \text{if } z_i > y_i(x_{-i}) \\ 0 & \text{if } z_i < y_i(x_{-i}) \end{cases}.$$

Then

$$\begin{aligned} M_i(x) = Q_i(x)x_i - \int_0^{x_i} Q_i(z_i, x_{-i})dz_i &= \int_0^{x_i} (Q_i(x_i, x_{-i}) - Q_i(z_i, x_{-i}))dz_i \\ &= \begin{cases} y_i(x_{-i}) & \text{if } Q_i(x) = 1 \\ 0 & \text{if } Q_i(x) = 0 \end{cases}. \end{aligned}$$

The player with the highest positive virtual value wins. Only the winning player has to pay and he pays the smallest amount needed to win.

Symmetric Case

Suppose $F_i = F$, so $\psi_i = \psi$ and

$$y_i(x_{-i}) = \max\left(\psi^{-1}(0), \max_{j \neq i} x_j\right)$$

In the optimal mechanism,

$$Q_i(z_i, x_{-i}) = \begin{cases} 1 & \text{if } z_i > y_i(x_{-i}) \\ 0 & \text{if } z_i < y_i(x_{-i}) \end{cases}$$

and

$$M_i(x) = \begin{cases} y_i(x_{-i}) & \text{if } Q_i(x) = 1 \\ 0 & \text{if } Q_i(x) = 0 \end{cases}.$$

Proposition 5

Suppose the design problem is regular and symmetric. Then the second-price auction with a reserve price $r^* = \psi^{-1}(0)$ is an optimal mechanism.

Intuition for Virtual Values

Why is it optimal to allocate the object based on virtual values?

Consider a single buyer whose value is distributed according to F. The seller sets price p to maximize $p(1 - F(p))$,

$$FOC : 1 - F(p) - pf(p) = 0 \iff \psi(p) = 0.$$

Alternatively, setting the probability (or quantity) of purchase $q = 1 - F(p)$, seller obtains price $p(q) = F^{-1}(1 - q)$. The revenue function is

$$R(q) = q \times p(q) = qF^{-1}(1 - q)$$

with

$$R'(q) = F^{-1}(1 - q) - \frac{q}{F'(F^{-1}(1 - q))}.$$

Substituting $p = F^{-1}(1 - q)$,

$$R'(q) = p - \frac{1 - F(p)}{f(p)} = \psi(p).$$

The seller sets the monopoly price p where marginal revenue $\psi(p)$ is 0, i.e., $p = \psi^{-1}(0)$.

Optimal Auction and Virtual Values

When facing multiple buyers, the optimal mechanism calls for the seller to set a *discriminatory* reserve price $r_i^* = \psi_i^{-1}(0)$ for each buyer i. If $x_i < r_i^*$ for every buyer i, the seller keeps the object. Otherwise, the object is allocated to the buyer generating the highest *marginal revenue*. The winning buyer pays $p_i = y_i(x_{-i})$, the smallest value needed to win.

Optimal auction is *inefficient*: with positive probability, the object is not allocated to a buyer even though it is worth 0 to the seller and has positive values for buyers.

Optimal Auction Favors Weak Buyers

Two buyers with regular cdf's F_1 and F_2 s.t. *supp* $F_1 =$ *supp* $F_2 = [0, \omega]$ and

$$\frac{f_1(x)}{1 - F_1(x)} < \frac{f_2(x)}{1 - F_2(x)}, \forall x \in [0, \omega].$$

Buyer 2 is relatively disadvantaged because his value is likely to be lower. F_1 first-oder stochastically dominates F_2, i.e., $F_2(x) \geq F_1(x)$ for $x \in [0, \omega]$. Check this by integrating the inequality above for $x \in [0, z]$ to obtain $-\log(1 - F_1(z)) \leq -\log(1 - F_2(z))$.

Reserve prices r_1^* and r_2^* satisfy

$$\psi_2(r_2^*) = 0 = \psi_1(r_1^*) = r_1^* - \frac{1 - F_1(r_1^*)}{f_1(r_1^*)} < r_1^* - \frac{1 - F_2(r_1^*)}{f_2(r_1^*)} = \psi_2(r_1^*).$$

Then $\psi_2(r_2^*) < \psi_2(r_1^*)$ implies $r_2^* < r_1^*$.

More on Discrimination and Inefficiency

When both buyers have the same value $x > r_1^*$, buyer 2 obtains the good in the optimal mechanism because

$$0 < \psi_1(x) = x - \frac{1 - F_1(x)}{f_1(x)} < x - \frac{1 - F_2(x)}{f_2(x)} = \psi_2(x).$$

For small $\varepsilon > 0$, $\psi_1(x) < \psi_2(x - \varepsilon)$ so buyer 2 gets the good even if $x_2 = x - \varepsilon$ when $x_1 = x$.

Second type of *inefficiency*: object not allocated to the highest value buyer

Application to Bilateral Trade

Myerson and Satterthwaite (1983)

- seller with privately known cost C; cdf F_s, density $f_s > 0$, supp $[\underline{c}, \overline{c}]$
- buyer with privately known value V; cdf F_b, density $f_b > 0$, supp $[\underline{v}, \overline{v}]$
- $\underline{c} < \underline{v} < \overline{c} < \overline{v}$

A direct mechanism (Q, M) specifies the probability of trade $Q(c, v)$ and the transfer $M(c, v)$ from the buyer to the seller for every reported profile (c, v).

Is there any efficient mechanism ($Q(c, v) = 1$ if $c < v$ and $Q(c, v) = 0$ if $c > v$) that is individually rational and incentive compatible?

More General Mechanisms

Useful to allow for mechanisms (Q, M_s, M_b) where M_s denotes the transfer *to* the seller and M_b the transfer *from* the buyer. (Q, M) special case with $M_b = M_s$.

Alternative question: is there an efficient mechanism (Q, M_s, M_b) that is individually rational and incentive compatible, which does not run a budget deficit, i.e.,

$$\int_{\underline{c}}^{\overline{c}} \int_{\underline{v}}^{\overline{v}} (M_b(c, v) - M_s(c, v)) f_b(v) f_s(c) dvdc \geq 0?$$

If the answer is negative, then the answer to the initial question is negative.

Revenue Equivalence

$$q_b(v) = \int_{\underline{c}}^{\overline{c}} Q(c,v) f_s(c) dc \quad \& \quad m_b(v) = \int_{\underline{c}}^{\overline{c}} M_b(c,v) f_s(c) dc$$

Incentive compatibility for the buyer requires

$$U_b(v) \equiv q_b(v)v - m_b(v) \geq q_b(v')v - m_b(v'), \forall v' \in [\underline{v}, \overline{v}].$$

As before,

$$U_b(v) = U_b(\underline{v}) + \int_{\underline{v}}^{v} q_b(v') dv',$$

which implies that $m_b(v) = -U_b(\underline{v}) + f(v, Q)$.

Similarly, $U_s(c) = m_s(c) - q_s(c)c = U_s(\overline{c}) + \int_c^{\overline{c}} q_s(c') dc'$ and $m_s(c) = U_s(\overline{c}) + g(c, Q)$.

For every incentive compatible mechanism (Q, M_s, M_b),

$$\int_{\underline{c}}^{\overline{c}} \int_{\underline{v}}^{\overline{v}} (M_b(c,v) - M_s(c,v)) f_b(v) f_s(c) dv dc = -U_b(\underline{v}) - U_s(\overline{c}) + h(Q).$$

The Vickrey-Clarke-Groves (VCG) Mechanism

Fix efficient allocation Q and consider the following payments. If $v > c$ then set $M_b(c, v) = \max(c, \underline{v})$ and $M_s(c, v) = \min(v, \overline{c})$. Otherwise, $M_b(c, v) = M_s(c, v) = 0$.

(Q, M_s, M_b) is incentive compatible (similar argument to the second-price auction with reserve).

Since $U_b(\underline{v}) = U_s(\overline{c}) = 0$,

$$\begin{aligned}
h(Q) &= \int_{\underline{c}}^{\overline{c}} \int_{\underline{v}}^{\overline{v}} (M_b(c, v) - M_s(c, v)) f_b(v) f_s(c)\, dv dc \\
&= \int_{\underline{c}}^{\overline{c}} \int_{c}^{\overline{v}} (\max(c, \underline{v}) - \min(v, \overline{c})) f_b(v) f_s(c)\, dv dc \\
&= \int_{\underline{c}}^{\overline{c}} \int_{c}^{\overline{v}} (\max(c, \underline{v}) + \max(-v, -\overline{c})) f_b(v) f_s(c)\, dv dc \\
&= \int_{\underline{c}}^{\overline{c}} \int_{c}^{\overline{v}} (\max(c - v, \underline{v} - v, c - \overline{c}, \underline{v} - \overline{c})) f_b(v) f_s(c)\, dv dc < 0.
\end{aligned}$$

Negative Result

Suppose (Q, M_s, M_b) is an efficient mechanism that is individually rational and incentive compatible.

In light of the VCG mechanism, efficiency and incentive compatibility imply $h(Q) < 0$. Individual rationality requires $U_b(\underline{v}), U_s(\overline{c}) \geq 0$. Then

$$\int_{\underline{c}}^{\overline{c}} \int_{\underline{v}}^{\overline{v}} (M_b(c,v) - M_s(c,v)) f_b(v) f_s(c) dv dc = -U_b(\underline{v}) - U_s(\overline{c}) + h(Q) < 0.$$

Every efficient, individually rational, and incentive compatible mechanism must run a budget deficit.

Theorem 2

If $\underline{c} < \underline{v} < \overline{c} < \overline{v}$, there exists no efficient bilateral trade mechanism (Q, M_b, M_s) with $M_b = M_s$ that is individually rational and incentive compatible.