Games of Incomplete Information

Up to this point, we have assumed that players know all relevant information about each other…

including the payoffs that each receives from the various outcomes of the game

and what they know about the structure of the game

Such games are known as games of complete information.

This is a very strong assumption.

Firms in an industry necessarily know each other's costs?

Firm bargaining with a union necessarily know the disutility that union members will feel if they go out on strike for a month?

In many circumstances, players have what is known as *incomplete information*.

In theory: need to consider a player's beliefs about other players' preferences…

his beliefs about their beliefs about his preferences…

and so on, much in the spirit of rationalizability.

Alternative approach (Harsanyi 1967-68).

Each player's preferences are determined by the realization of a random variable.

Actual realization is observed only by the player, but ex ante probability distribution is assumed to be common knowledge.

Incomplete information is reinterpreted as a game of imperfect information:

Nature makes the first move, choosing realizations of the random variables that determine each player's preference type

Each player observes the realization of only his own random variable.

This describes a Bayesian game.

Example 8.E.1: Modified prisoners' dilemma.

With probability μ, prisoner 2 has the usual preferences ("type I preferences")

With probability $(1 - \mu)$, prisoner 2 hates to rat on his accomplice (this is type II).

Psychic penalty equal to 6 years in prison for confessing.

Prisoner 1 always has the usual preferences.

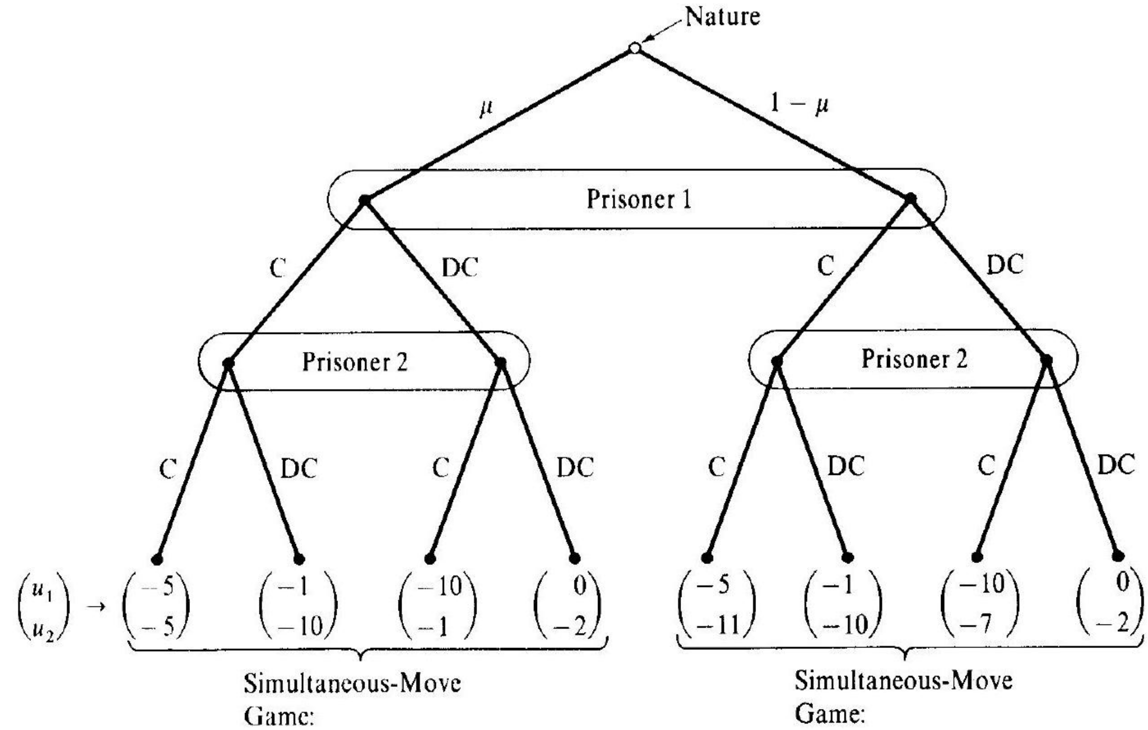

Prisoner 2

		DC	C
Prisoner 1	DC	0, −2	−10, −1
	C	−1, −10	−5, −5

Prisoner 1

	DC	C
DC	0, −2	−10, −7
C	−1, −10	−5, −11

A pure strategy (a complete contingent plan) for player 2 can be viewed as:

A function that for each possible realization of his preference type indicates what action he will take.

Hence, prisoner 2 now has four possible pure strategies:

(confess if type I, confess if type II)

(confess if type I, don't confess if type II)

(don't confess if type I, confess if type II)

(don't confess if type I, don't confess if type II).

Player 1 does not observe player 2's type, and so a pure strategy for player 1 in this game is simply a (noncontingent) choice of either "confess" or "don't confess."

Each player i has a payoff function $u_i(s_i, s_{-i}, \theta_i)$

$\theta_i \in \Theta_i$ is a random variable chosen by nature that is observed only by player i.

Joint probability distribution of the θ_i 's is $F(\theta_1, \ldots, \theta_1)$, common knowledge among the players.

Letting $\Theta = \Theta_1 \times \cdots \times \Theta_I$, a Bayesian game is summarized as:

$$[I, \{S_i\}, \{u_i(\cdot)\}, \Theta, F(\cdot)]$$

A pure strategy for player $\boldsymbol{i}$ in a Bayesian game is a function $\boldsymbol{s_i(\theta_i)}$, or decision rule, that gives the player's strategy choice for each realization of his type $\boldsymbol{\theta_i}$.

Player i 's pure strategy set $\mathcal{C}_i'$ is therefore the set of all such functions.

Player i 's expected payoff given a profile of pure strategies for the I players $(s_1(\cdot), \ldots, s_I(\cdot))$:

$$\tilde{u}_i(s_1(\cdot), \ldots, s_I(\cdot)) = E_\theta[u_i(s_1(\theta_1), \ldots, s_I(\theta_I), \theta_i)] \qquad (1)$$

Turn now to the Nash equilibrium of this game of imperfect information.

Definition 8.E.1: A (pure strategy) Bayesian Nash equilibrium for the Bayesian game $[I, \{S_i\}, \{u_i(\cdot)\}, \Theta, F(\cdot)]$ is a profile of decision rules $(s_1(\cdot), \ldots, s_I(\cdot))$ that constitutes a Nash equilibrium of game $\Gamma_N = [I, \{\mathcal{P}_i\}, \{\tilde{u}_i(\cdot)\}]$.

That is, for every $i = 1, \ldots, I$,

$$\tilde{u}_i(s_i(\cdot), s_{-i}(\cdot)) \geq \tilde{u}_i(s_i'(\cdot), s_{-i}(\cdot))$$

for all $s_i'(\cdot) \in \mathcal{C}_i'$, where $\tilde{u}_i(s_i(\cdot), s_{...}(\cdot)) = E_\theta[u_i(s_1(\theta_1), \ldots, s_I(\theta_I), \theta_i)]$, as defined before.

In a (pure strategy) Bayesian Nash equilibrium, each player must be playing a best response to the conditional distribution of his opponents' strategies *for each type that he might end up having.*

Proposition 8.E.1: A profile of decision rules $(s_1(\cdot), \ldots, s_I(\cdot))$ is a Bayesian Nash equilibrium in Bayesian game $[I, \{S_i\}, \{u_i(\cdot)\}, \Theta, F(\cdot)]$ if and only if, for all i and, for all $s_i' \in S_i$,

$$E_{\theta_{-i}}[u_i(s_i(\bar{\theta}_i), s_{-i}(\theta_{-i}), \bar{\theta}_i) \mid \bar{\theta}_i] \geq E_{\theta_{-i}}[u_i(s_i', s_{-i}(\theta_{-j}), \bar{\theta}_i) \mid \bar{\theta}_i] \qquad (2)$$

Proof:

Show first that (1) implies (2).

To do so, let's show that if (2) fails, then (1) fails.

Assume (2) fails: that is, equation (2) does not hold for some player i for some $\bar{\theta}_i \in \Theta_i$ that occurs with positive probability.

Then player i could do better by changing his strategy choice in the event he gets realization $\bar{\theta}_i$, contradicting $(s_1(\cdot), \ldots, s_l(\cdot))$ being a Bayesian Nash equilibrium.

Hence (1) fails.

Show now that (2) implies (1).

If condition (2) holds for all $\bar{\theta}_i \in \Theta_i$ occurring with positive probability, then player i cannot improve on the payoff he receives by playing strategy $s_i(\cdot)$:

Simply take the average over θ_i over both sides of the inequality in (2).

In essence, we can think of *each type of player i as being a separate player* who maximizes his payoff given his conditional probability distribution over the strategy choices of his rivals.

Example 8.E. 1 Continued:

Type I of prisoner 2 has a dominant strategy: "confess".

Type II of prisoner 2 also has a dominant strategy: "don't confess."

Given this behavior by prisoner 2 , prisoner l's best response is to play "don't confess" if:

$$\underbrace{-10\mu + 0(1-\mu)}_{U_1^e(Don't\ confess)} > \underbrace{-5\mu - 1(1-\mu)}_{U_1^e(Confess)}$$

$$\mu < \frac{1}{6}$$

And play "confess" if $\mu > \frac{1}{6}$. (He is indifferent if $\mu = \frac{1}{6}$.)

Example 8.E.2:

Two firms: 1 and 2.

Any independent invention by one of the firms is shared fully with the other.

There is a new invention that either of the two firms could potentially develop.

To develop this new product costs a firm $c \in (0,1)$.

The benefit of the invention to each firm i *is known only by that firm*.

Each firm i has a type θ_i that is independently drawn from a uniform distribution on [0,1].

Its benefit from the invention is θ_i is $(\theta_i)^2$.

Timing:

Firms each privately observe their own type.

Then they each simultaneously choose either to develop the invention or not.

Let us now solve for the Bayesian Nash equilibrium of this game.

Write $s_i(\theta_i) = 1$ if type θ_i of firm i develops the invention and $s_i(\theta_i) = 0$ if it does not.

If firm i develops it, its payoff is $(\theta_i)^2 - c$ regardless of whether firm j does so.

If firm i decides not to develop it...

it will have an expected payoff equal to $(\theta_i)^2 \cdot \text{Prob}\left(s_j(\theta_j) = 1\right)$.

That is, it won't have the cost, but will only have the benefit if the other firm makes the investment.

Hence, firm i 's best response is to invest if and only if its type θ_i is such that:

$$(\theta_i)^2 - c \geq (\theta_i)^2 \cdot \text{Prob}\left(s_j(\theta_j) = 1\right)$$

$$\theta_i \geq \left[\frac{c}{1 - \text{Prob}\left(s_j(\theta_j) = 1\right)}\right]^{1/2}$$

For any given strategy of firm j, firm i 's best response is a *cutoff rule*:

It invests for all θ_i above the value on the right-hand side above, and does not for all θ_i below it.

If firm i existed in isolation, it would be indifferent about developing the Zigger when $\theta_i = \sqrt{c}$.

But we see that this cutoff is always (weakly) above it: free-riding.

Suppose then that $\hat{\theta}_1, \hat{\theta}_2 \in (0,1)$ are the cutoff values for firms 1 and 2 respectively in a Bayesian Nash equilibrium

(it can be shown that $0 < \hat{\theta}_i < 1$ for $i = 1,2$ in any Bayesian Nash equilibrium of this game)

Then uniform distribution implies:

$$\text{Prob}\left(s_j(\theta_j) = 1\right) = 1 - \hat{\theta}_j$$

Hence:

$$\hat{\theta}_i = \left[\frac{c}{1 - (1 - \hat{\theta}_j)}\right]^{1/2} = \left[\frac{c}{\hat{\theta}_j}\right]^{1/2} \Longrightarrow \left(\hat{\theta}_i\right)^2 \hat{\theta}_j = c$$

An analogous computation holds for the other firm:

$$\left(\hat{\theta}_j\right)^2 \hat{\theta}_i = c$$

Hence:

$$\left(\hat{\theta}_i\right)^2 \hat{\theta}_j = \left(\hat{\theta}_j\right)^2 \hat{\theta}_i \Longrightarrow \hat{\theta}_i = \hat{\theta}_j$$

So:

$$\left(\hat{\theta}_i\right)^2 \hat{\theta}_j = c \Longrightarrow \left(\hat{\theta}_i\right)^2 \hat{\theta}_i = c \Longrightarrow \hat{\theta}_i^{\,3} = c \Longrightarrow \hat{\theta}_i = c^{1/3}$$

Consider again a Nash equilibrium:

> Suppose we start with a game of complete information that has a unique mixed strategy equilibrium in which players actually randomize.
>
> Change the game by introducing many different types (formally, a continuum) of each player, with the realizations of the various players' types being statistically independent of one another.
>
> All types of a player have identical preferences.

A (pure strategy) Bayesian Nash equilibrium of this Bayesian game is then precisely equivalent to a mixed strategy Nash equilibrium of the original complete information game.

The Possibility of Mistakes: Trembling-Hand Perfection

Consider the game below:

	L	R
U	1,1	$0,-3$
D	$-3,0$	0,0

(D, R) is a Nash equilibrium involving play of weakly dominated strategies.

> However, caution might preclude the use of such strategies.

Trembling-hand perfect Nash equilibrium: Nash equilibria that are robust to the possibility that, with some very small probability, players make mistakes.

For any normal form game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$, define a perturbed game:

$$\Gamma_\varepsilon = [I, \{\Delta_\varepsilon(S_i)\}, \{u_i(\cdot)\}]$$

by choosing for each player i and strategy $s_i \in S_i$ a number $\varepsilon_i(s_i) \in (0,1)$, with $\sum_{s_i} \varepsilon_i(s_i) < 1$.

Define player i 's perturbed strategy set:

$$\Delta_\varepsilon(S_i) = \left\{\sigma_i : \sigma_i(s_i) \geq \varepsilon_i(s_i) \text{ for all } s_i \in S_i \text{ and } \sum_{s_i \varepsilon S_i} \sigma_i(s_i) = 1\right\}$$

That is, perturbed game Γ_ε is derived from the original game Γ_N by requiring that:

each player i play every one of his strategies s_i with at least some minimal positive probability $\varepsilon_i(s_i)$

$\varepsilon_i(s_i)$ is interpreted as the unavoidable probability that strategy s_i gets played by mistake.

We want to consider as predictions in game Γ_N only those Nash equilibria σ that are robust to the possibility that players make mistakes.

The robustness test we employ can be stated roughly as:

To consider σ as a robust equilibrium, we want there to be at least some slight perturbations of Γ_N whose equilibria are close to σ.

Definition 8.F.1: A Nash equilibrium σ of game $\Gamma_N = [I, \{\Delta(S_j)\}, \{u_i(\cdot)\}]$ is (normal form) trembling-hand perfect if there is some sequence of perturbed games $\{\Gamma_{\varepsilon^k}\}_{k=1}^{\infty(n)}$ that converges to Γ_N [in the sense that $\lim_{k\to\infty} \varepsilon_i^k(s_i) = 0$ for all i and $s_i \in S_i$], for which there is some associated sequence of Nash equilibria $\{\sigma^k\}_{k=1}^{\infty}$ that converges to σ (i.e., such that $\lim_{k\to\infty} \sigma^k = \sigma$).

This is a mild test of robustness:

We require only that *some* perturbed games exist that have equilibria arbitrarily close to σ.

A stronger test would require that the equilibrium σ be robust to all perturbations close to the original game.

An alternative formulation makes it easier to check whether a Nash equilibrium is trembling-hand perfect:

Proposition 8.F.1: A Nash equilibrium σ of game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i(\cdot)\}]$ is (normal form) trembling-hand perfect if and only if there is some sequence of totally mixed strategies $\{\sigma^k\}_{k=1}^{xi}$ such that $\lim_{k\to\infty} \sigma^k = \sigma$ and σ_i is a best response to every element of sequence $\{\sigma_{-i}^k\}_{k=1}^{k}$ for all $i = 1, \dots, I$.

Immediate consequence:

Proposition 8.F.2: If $\sigma = (\sigma_1, \dots, \sigma_l)$ is a (normal form) trembling-hand perfect Nash equilibrium, then σ_i is not a weakly dominated strategy for any $i = 1, \dots, I$. Hence, in any (normal form) trembling-hand perfect Nash equilibrium, no weakly dominated pure strategy can be played with positive probability.

The converse, that any Nash equilibrium not involving play of a weakly dominated strategy is necessarily trembling-hand perfect, turns out to be true for two-player games but not for games with more than two players.

Thus, trembling-hand perfection can rule out more than just Nash equilibria involving weakly dominated strategies.

Auctions

Many economic transactions are conducted through auctions.

- Governments auction off treasury bonds, foreign currency, public enterprises, exploration rights...
- Artwork, antiques, cars, and houses.
- Government contracts.
- Firms acquire inputs or subcontract work.
- Battles for the acquisition of a firm between business groups or private equity funds.

Four main types of auctions:

- Ascending bid auction (open, oral, or English Type):
  - Price is raised until only one participant remains, who wins the auctioned item at the final price.
- Descending bid auction (Dutch Type):
  - Auctioneer starts at a very high price and reduces it continuously until someone accepts it at that announced price.
  - Winner pays the item at that price.
- First-price sealed-bid auction:
  - Each participant submits their bid in a sealed envelope without seeing the other players' bids
  - Item is auctioned to the participant who submitted the highest bid at the price declared in their envelope.
- Second-price sealed-bid auction (Vickrey Auction):
  - Participants submit their bids in sealed envelopes.
  - Highest bid wins the item but pays the price declared by the second highest bid.

Auctions also differ in terms of how the auctioned item's value is assessed by its participants.

Private value auction:

- Each participant's valuation is known only to themselves
- Example: auction of a work of art.

Common value auction:

Value of an item is the same for all participants

But participants have different private information about that value.

Example: oil well or a company for sale may have the same value for all potential buyers, but each one has different estimates of that value.

Our focus: sealed-bid games

Simpler to analyze

In the case of private values, strategically equivalent to the descending bid auction

Second-price sealed-bid auction is strategically equivalent to the ascending bid auction.

Setup

Let's consider independent private values.

Each bidder knows his own valuation.

The valuation of any given bidder is an independent random variable.

The distribution is common knowledge.

Definition: Auction with independent private values:

- $N = \{1, \ldots, n\}$: set of bidders

- $A_i \subset \mathbb{R}_+$: set of strategies: it's a bid. If player i chooses $a_i \in A_i$, it means his bid is a_i.

- Set of types for player i: $\Theta_i = [v_{min}, v_{max}]$, with $v_{min} \geq 0$.

  For example, player i might assign a minimum value of \$50 to the auctioned object and a maximum value of \$80 – only player i knows the exact value, but the others know this range (which can be from zero to infinity if the other players do not have relevant information).

- Each player i's belief, where i believes that the valuations of their opponents are drawn from independent draws from a distribution F that is strictly increasing and continuous on $[v_{min}, v_{max}]$.

  E.g., If the distribution is uniform, density will be $\frac{1}{v_{max}-v_{min}}$.

- Payoff function for each player, defined for all $a \in A$ and $v \in \Theta$:

$$u_i(a, v) = \begin{cases} v_i - P(a), & if\ a_j \leq a_i, for\ all\ j \neq i, \\ 0, & if\ a_j > a_i\ for\ some\ j \neq i \end{cases}$$

$P(a)$ is the price paid by the winner if bids are a.

In the first row, we have the payoff if the player wins the auction: it's what he obtains ($v_i - P(a)$) if his bid is the highest ($a_j \leq a_i, for\ all\ j \neq i$, that is, all other players make lower bids).

We are ignoring the possibility of a tie; in this case, we can assume that the good is divided equally, or that the winner is chosen randomly – it makes no difference to our results.

Second-Price Auction (Vickrey)

In this type of auction, the participant who bids the highest wins but pays a price equal to the second-highest bid.

Despite having multiple Bayesian equilibria, bidding one's own valuation of the object is a weakly dominant strategy for any player i.

To understand this, imagine you're player i and let x represent the highest bid among your competitors in the auction.

You're considering bidding:

$a_i' < v_i$: a bid less than your valuation,

exactly v_i: a bid equal to your valuation, or

$a_i'' > v_i$: a bid higher than your valuation.

 Depending on the value of x, the following table shows the payoffs you would receive for each possible action.

	$x \leq a_i''$	$a_i' < x < v_i$	$x = v_i$	$v_i < x < a_i''$	$a_i'' < x$
a_i'	Ganha/empata; paga x	Perde	Perde	Perde	Perde
v_i	Ganha; paga x	Ganha; paga x	Empata; paga v_i	Perde	Perde
a_i''	Ganha; paga x	Ganha; paga x	Ganha; paga v_i	Ganha/empata; paga x	Perde

 If you bid below v_i:

- Sometimes, you lose when you should have won (i.e., $a_i < x < v_i$).

 If you bid above v_i:

- Sometimes, you win when you should have lost (i.e., $a_i > x > v_i$).

First-Price Auction

In first-price auctions, the participant with the highest bid wins the item and pays exactly their bid amount.

Let's denote the bid of a player with valuation v_i as $\beta_i(v_i)$ and look for a symmetric equilibrium, meaning $\beta_i(v_i) = \beta(v)$ for every player $i \in N$

All players use the same function β(v)β(v) to determine their bid based on their valuation v.

First, although we won't prove it here, we assume that the strategies $\beta_i(v_i)$, and thus $\beta(v)$, are strictly increasing and continuous functions on $[v_{min}, v_{max}]$.

Intuitively, this means that the higher the value an individual assigns to the auctioned item, the higher their bid will be.

We'll construct the expected payoff for a player with valuation vv who places a bid b, assuming all other players are bidding according to β.

Without loss of generality, let's call this player "Player 1" (it doesn't matter if we label them as Player 2, 5, or i).

Note first that if this player wins the auction, their payoff is $(v - b)$: this is their "net value," meaning the value vvthey assign to the item minus the payment b they make.

If they lose the auction, their payoff is zero.

Therefore, their expected payoff is:

$$(v - b) \times Prob(jogador\ 1\ ganha\ o\ leil\text{ão}) + 0 \times Prob(jogador\ perde\ o\ leil\text{ão})$$

The second term is zero, so we can ignore it.

We now need an expression for the probability that the player wins the auction. This happens if they place the highest bid among all other players.

Let's denote the bids of the other players as b_2, b_3, ..., b_N (remember there are n participants in total in the auction).

$$Prob(jogador\ ganha\ o\ leil\text{ão}) = Prob(b\ \text{é}\ maior\ que\ todos\ os\ demais\ lances)$$

$$= Prob(b > b_2, b > b_3, \dots, b > b_N)$$

Now, let's remember that we're assuming the other players are using the strategy $b_2 = \beta(v_2)$,...,$b_n = \beta(v_N)$.

We can rewrite this last expression as:

$$= Prob\big(b > \beta(v_2), b > \beta(v_3), \dots, b > \beta(v_N)\big)$$

We're assuming that the valuations v_2, ..., v_n are independent — that is, they are independent random variables.

Remember: if X and Y are independent random variables, then $P(X\ and\ Y) = P(X)P(Y)$

We'll use this to rewrite the last expression as:

$$= Prob(b > \beta(v_2)) \times Prob(b > \beta(v_3)) \times \dots \times Prob(b > \beta(v_N))$$

The probability that $b > \beta(v_i)$ is equal to the probability that $\beta^{-1}(b) > v_i$:

we're applying the inverse function to both sides, and we keep the inequality sign because β is an increasing function — so its inverse is also increasing.

Let's rewrite the expression above:

$$= Prob(\beta^{-1}(b) > v_2) \times Prob(\beta^{-1}(b) > v_3) \times \ldots \times Prob(\beta^{-1}(b) > v_N)$$

Now let's use the fact that the random variables v_2 to v_n have the same probability distribution F.

This means that the probability that v_2 is less than $\beta^{-1}(b)$ is the same as the probability that v_3 is less than $\beta^{-1}(b)$, and so on.

Furthermore, we have an expression for this probability: generally, the "probability that a random variable is less than a given value" is simply its distribution F evaluated at that value.

The expression above becomes:

$$= F(\beta^{-1}(b)) \times F(\beta^{-1}(b)) \times \ldots \times F(\beta^{-1}(b))$$

We now have the term $F(\beta^{-1}(b))$ multiplied by itself (N-1) times (one for each of the other players, excluding Player 1, whose bid we're evaluating).

This expression can be further simplified:

$$= F(\beta^{-1}(b))^{N-1}$$

Therefore, the expected utility of Player 1 is simply:

$$(v - b) \times Prob(jogador\ 1\ ganha\ o\ leilão) = (v - b)F\left(\beta^{-1}(b)\right)^{N-1}$$

Antes de resolver o caso geral, vamos fazer uma versão simples, para facilitar o entendimento.

Vamos supor que haja apenas dois jogadores: N=2, e portanto N-1=1.

Além disso, vamos supor que a função $\beta(v)$ seja particularmente simples: $\beta(v) = \gamma v$, em que$\gamma > 0$ é simplesmente um número real positivo.

Logo, a função inversa é simplesmente $^1/_\gamma$.

Por último, vamos considerar uma distribuição uniforme no intervalo (0,1), que tem densidade constante f(x)=1 e distribuição F(x)=x

A expressão anterior se torna:

$$= (v - b)\frac{b}{\gamma}$$

Para maximizar essa função, vamos obter a condição de primeira ordem – ou seja, vamos derivar e igualar a zero em relação a b, o lance do jogador 1.

Vamos usar a derivada do produto, em que o primeiro termo é $(v - b)$ e o segundo é $\frac{b}{\gamma}$. Obtemos então:

$$-\frac{b}{\gamma} + (v - b)\frac{1}{\gamma} = 0$$

Podemos reorganizar um pouco essa expressão:

$$(v-b)=b$$ [redacted]

Logo, 2b=v, e obtemos o lance ótimo do jogador:

$$b=\frac{v}{2}$$

Ou seja, o lance ótimo é metade do valor que o jogador atribui ao bem.

Assim, o jogador equilibra dois objetivos: maximizar a chance de ganhar o leilão (o que sugere um lance alto) e não pagar muito caro caso seja vencedor (o que sugere um lance baixo).

Vamos voltar agora ao caso geral.

Cada participante quer maximizar sua utilidade esperada, que já calculamos:

$$(v-b)F\left(\beta^{-1}(b)\right)^{N-1}$$

Lembrando: o termo $(v-b)$ é o quanto o jogador obtém em casa de vitória no leilão (sua valoração v menos seu lance b).

Já o termo $F\left(\beta^{-1}(b)\right)^{N-1}$ é a probabilidade de vitória: ou seja, é a probabilidade de que seu lance b seja maior do que o lance dos demais (N-1) jogadores.

Vamos então maximizar essa utilidade esperada.

Para tanto, vamos derivar e igualar a zero para obter a condição de primeira ordem:

$$-F\left(\beta^{-1}(b)\right)^{N-1}+(v-b)(N-1)F\left(\beta^{-1}(b)\right)^{N-2}F'\left(\beta^{-1}(b)\right)\frac{1}{\beta'\left(\beta^{-1}(b)\right)}=0$$

Conseguimos obter essa condição de primeira ordem do fato de β ser diferenciável em todo seu domínio (lembre-se que assumimos que essa função é estritamente crescente e contínua) e pelo teorema da função inversa.

Para $\beta(v)$ ser um equilíbrio a condição de primeira ordem deve ser obedecida quando substituímos b por $\beta(v)$:

$$-F(v)^{N-1}+(v-\beta(v))(N-1)F(v)^{N-2}F'(v)\frac{1}{\beta'(v)}=0$$

Ou, reorganizando os termos:

$$-\beta'(v)F(v)^{N-1}+(v-\beta(v))(N-1)F(v)^{N-2}F'(v)=0$$

$$\beta'(v)F(v)^{N-1}=(v-\beta(v))(N-1)F(v)^{N-2}F'(v)$$

$$\beta'(v)F(v)^{N-1}+\beta(v)(N-1)F(v)^{N-2}F'(v)=v(N-1)F(v)^{N-2}F'(v) \qquad \text{(i)}$$

Essa expressão define uma equação diferencial em β, a função que queremos encontrar.

Vamos usar uma estratégia comum para resolver equações diferenciais. Observe com atenção o lado esquerdo: trata-se da derivada da função $\beta(v)F(v)^{N-1}$.

Vamos verificar isso.

De fato, você pode derivar $\beta(v)F(v)^{N-1}$em relação à variável v, usando a regra do produto: derivada do primeiro termo ($\beta(v)$) vezes o segundo ($F(v)^{N-1}$), mais a derivada do segundo ($F(v)^{N-1}$) vezes o primeiro ($\beta(v)$).

Ficamos então com:

$$[\beta'(v)F(v)^{N-1}] + \beta(v)[F(v)^{N-1}]' \qquad \text{(ii)}$$

Para concluir, precisamos usar a regra da cadeia para derivar o segundo termo da expressão (ii).

Afinal, temos uma composição de funções: temos a função $F(v)$, e elevamos essa função ao expoente (N-1).

Usando a regra da cadeia, a expressão (ii) pode ser reescrita como:

$$[\beta'(v)F(v)^{N-1}] + \beta(v)[(N-1)F(v)^{N-1}F'(v)]$$

Com isso, concluímos que o lado esquerdo da expressão (i) é, de fato, a derivada da função $\beta(v)F(v)^{N-1}$.

Podemos então reescrever a equação (i):

$$\frac{d\beta(v)F(v)^{N-1}}{dv} = v(N-1)F(v)^{N-2}F'(v)$$

Vamos reorganizar ligeiramente:

$$d\beta(v)F(v)^{N-1} = v(N-1)F(v)^{N-2}F'(v)dv$$

Podemos agora integrar dos dois lados dessa expressão entre v_{min} e v.

Vamos mudar o nome da variável de integração de v para x, pois v está num limite de integração (lembre-se de que um limite de integração não pode ser uma variável de integração – mas não se preocupe, isso é apenas notação).

$$\int_{v_{min}}^{v} d\beta(x)F(x)^{N-1} = \int_{v_{min}}^{v} x(N-1)F(x)^{N-2}F'(x)dx \qquad \text{(iii)}$$

Mais uma vez, vamos olhar com atenção para o lado esquerdo.

Qual é o integrando, ou seja, a função que está sendo integrada?

Note que nem sequer aparece o integrando! (Formalmente, isso significa que estamos integrando a função constante igual a 1).

É como calcular a seguinte integral:

$$\int_{a}^{b} dy = y(b) - y(a)$$

Vamos aplicar isso para reescrever a equação (iii). Apenas o lado esquerdo vai mudar:

$$\beta(v)F(v)^{N-1} - \beta(v_{min})F(v_{min})^{N-1} = \int_{v_{min}}^{v} x(N-1)F(x)^{N-2}F'(x)dv$$

Observe que $F(v_{min}) = 0$: a função distribuição sempre vale zero no início do suporte! (Formalmente, $F(v_{min})$ é a probabilidade de que a variável aleatória v seja menor que o menor valor possível v_{min}, o que é impossível, e tem portanto probabilidade zero.)

A última equação se torna então:

$\beta(v)F(v)^{N-1} = \int_{v_{min}}^{v} x(N-1)F(x)^{N-2}F'(x)dx$ (iv)

Vamos agora fazer a integral por partes para simplificar o lado direito dessa expressão.

Lembre-se da integral por partes:

$$\int_a^b u(x)v'(x)dx = u(b)v(b) - u(a)v(a) - \int_a^b u'(x)v(x)dx$$

Para aplicar a fórmula da integral por partes à expressão (iv), faremos:

$$u(x) = x$$

$$v'(x) = (N-1)F(x)^{N-2}F'(x)$$

Podemos então calcular:

$$u'(x) = 1$$

$$v(x) = F(x)^{N-1}$$

(Para obter v(x), usamos a mesma conta que fizemos para reescrever o lado esquerdo da equação (i).)

O lado direito da equação (iv) se torna (reforçando: esse é apenas o lado direito!):

$$\int_{v_{min}}^{v} x(N-1)F(x)^{N-2}F'(x)dx = vF(v)^{N-1} - v_{min}F(v_{min})^{N-1} - \int_{v_{min}}^{v} 1 \cdot F(x)^{N-1}dx$$

Vamos usar novamente $F(v_{min}) = 0$ para reescrever essa expressão:

$= vF(v)^{N-1} - \int_{v_{min}}^{v} F(x)^{N-1}dx$ (v)

A expressão (v) é portanto apenas outra forma de escrever o lado direito da equação (iv).

Vamos reescrever a equação (iv):

$$\beta(v)F(v)^{N-1} = vF(v)^{N-1} - \int_{v_{min}}^{v} F(x)^{N-1}dx$$

Isolando $\beta(v)$ obtemos:

$$\beta(v) = v - \frac{\int_{v_{min}}^{v} F(x)^{N-1}\,dx}{F(v)^{N-1}}$$

Note que $\beta(v_{min}) = v_{min}$, pois:

$$\beta(v_{min}) = v_{min} - \frac{\int_{v_{min}}^{v_{min}} F(x)^{N-1}\,dx}{F(v)^{N-1}} = v_{min} - \frac{0}{F(v)^{N-1}} = v_{min}$$

Usamos aqui o fato de que quando os limites de integração são iguais, a integral vale zero.

A interpretação é simples: o jogador que atribui o menor valor possível ao bem (v_{min}) dá um lance igual a esse valor.

Todos os jogadores dão lances menores que suas valorações individuais a respeito do objeto do leilão: $v > v_{min}$ implica que o segundo termo do lado direito da expressão para $\beta(v)$ é estritamente negativo.

Lembre-se: o jogador busca equilibrar os objetivos de maximizar a chance de ganhar o leilão, mas não pagar muito caro em caso de vitória.

Como exemplo, vamos calcular β assumindo que F é uniforme em $[0,1]$, ou seja, $F(x) = x$.

- **Exemplo (Lance ótimo quando F é uniforme):** Vamos resolver de dois jeitos. Primeiro, de maneira mais imediata, podemos aplicar a expressão de $\beta(v)$ obtida acima. Nesse caso:

$$\beta(v) = v - \frac{\int_{v_{min}}^{V} F(x)^{N-1}\,dx}{F(v)^{N-1}} = v - \frac{1}{v^{N-1}}\frac{v^N}{N} = \frac{N-1}{N}v$$

Vamos agora resolver de maneira explícita.
Queremos encontrar um equilíbrio simétrico da forma $\beta(v) = av$, onde a é uma constante positiva qualquer.
O valor esperado do payoff do jogador de tipo v que dá o lance b quando todos os outros jogadores dão lances β é dado por:

$$(v-b)Prob(maior\ lance\ \text{é}\ b) = (v-b)\big(prob(av \le b)\big)^{N-1}$$
$$= (v-b)\left(\frac{b}{a}\right)^{N-1}$$

A condição de primeira ordem que maximiza o payoff esperado é:

$$(v-b)(N-1) = b$$

Que, resolvendo para b, nos dá:

$$b = \frac{N-1}{N}v$$

Observe que quando o número N de participantes do leilão é grande, o lance *b* se aproxima da valor *v* que o jogador atribui ao bem: se há muita competição, o jogador vai ter que pagar caro se quiser ter alguma chance de vitória.
Vemos isso matematicamente: (N-1)/N se aproxima de 1 quanto N aumenta indefinidamente.

Seção: Equivalência em Leilões

Se você quiser vender um bem através de um leilão, que modelo irá usar?

Inglês ou Holandês?

É útil fazer envelope fechado?

Essas decisões são tomadas todos os dias em licitações públicas, por exemplo.

Um importante resultado da teoria dos leilões é o "Teorema da Equivalência das Receitas", que afirma que, sob certas condições, muitos tipos diferentes de leilões geram exatamente a mesma receita esperada para o leiloeiro.

Uma consequência desse teorema, que usaremos abaixo, é que, em equilíbrio, cada um desses mecanismos de leilão irá gerar o mesmo excedente para cada participante do leilão, condicional à sua valoração v.

Esse resultado é conhecido como “Equivalência de Payoff”.

Em particular, um participante com valoração v tem, em equilíbrio, a seguinte utilidade esperada:

$$E[U(v_i)] = \int_{v_{min}}^{v_i} F(t)^{n-1} d(t)$$

Essa expressão pode ser muito útil como “atalho” para determinarmos os lances em equilíbrio de um leilão, especialmente para casos mais complicados que leilões como o de segundo preço, onde o equilíbrio é simplesmente $b^*(v_i) = v_i$ (ou seja, o lance ótimo é a própria valoração no leilão de segundo preço).

Vamos utilizar também a seguinte expressão para o payoff esperado de um participante de um leilão:

$$E[U(v_i)] = \big(v_i - b(v_i)\big)F(v_i)^{n-1}$$

Vamos obter o lance ótimo $b^*(v_i)$ de um participante em um leilão de primeiro preço

Basta utilizar as duas expressões dadas no enunciado e igualá-las:

$$\big(v_i - b(v_i)\big)F(v_i)^{N-1} = \int_{v_{min}}^{v} F(t)^{n-1} d(t)$$

Isolando $b(v_i)$ para obter uma expressão para o lance ótimo, chegamos diretamente à resposta do exercício:

$$b^*(v_i) = v_i - \frac{\int_{v_{min}}^{v_i} F(t)^{n-1} d(t)}{F(v_i)^{n-1}}$$