MS&E 246: Game Theory with Engineering Applications

Lecture 1

Ramesh Johari

Outline

- Administrative stuff
- Course introduction
- A game

Administrative details

- My e-mail: ramesh.johari@stanford.edu
- Course assistant:

  Christina Aperjis, caperjis@stanford.edu
- Website:

  eeclass.stanford.edu/msande246

  All students must sign up there,
  and keep up with announcements

Administrative details

- 6-7 problem sets

  Assigned Thursday, due following Thursday in box outside Terman 319

  No late assignments accepted
- Midterm to be held February 8 (in class)

Big picture

Economics and engineering are tied together more than ever

Game theory provides a set of tools we can use to study problems at this interface

Motivating examples

- Electronic marketplaces
  - eBay auctions:
    Fixed termination time
  - Amazon auctions
    Terminate after 10 minutes of inactivity

Which yields higher revenue?

Motivating examples

- Internet resource allocation
  - TCP: regulates flow of packets through the Internet
  - Malicious users can grab much more than “fair” share
  - How do we design “fair”, “efficient” allocation protocols that are robust to gaming?

Motivating examples

- Electricity markets
  - Electricity can't be stored, and must be reliable
  - Market failure is disastrous (e.g., California in 2000)
  - How do we design efficient, sustainable markets?

Internet provider competition

- Internet = 1000s of ASes (autonomous systems)
- Bilateral contracts between ASes:

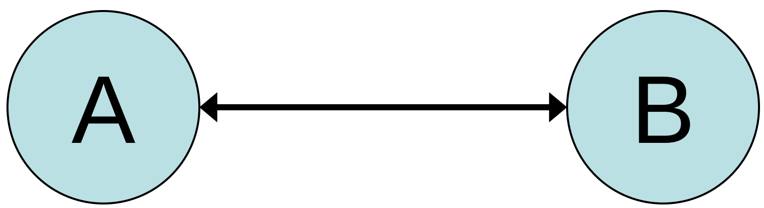

- Transit vs. peer contracts

ISP contracts

- Transit vs. peer contracts
  - Transit:
    If A pays B, then
    A agrees to carry all traffic to/from B
  - Peer:
    A and B are of similar size,
    and agree to exchange traffic
    terminating in each other's network

Problems in the ISP industry

- In 2002, seven dominant players:
  - Sprint
  - AT&T
  - MCI/UUnet
  - Qwest
  - C&W
  - Level3
  - Genuity

Problems in the ISP industry

- In 2002, seven dominant players:
  - Sprint (subsidized by wireless)
  - AT&T ACQUIRED (SBC)
  - MCI/UUnet ACQUIRED (Verizon)
  - Qwest $18B debt
  - C&W R.I.P. (in U.S.)
  - Level3 (merged w/Genuity)
  - Genuity ACQUIRED (Level3)

Econ 101, pt. 1: war of attrition

- Pricing below marginal cost

$\Rightarrow$ War of attrition (repeated game):

Lose money now in hopes of being last firm standing

Econ 101, pt. 2: Bertrand

- Example:

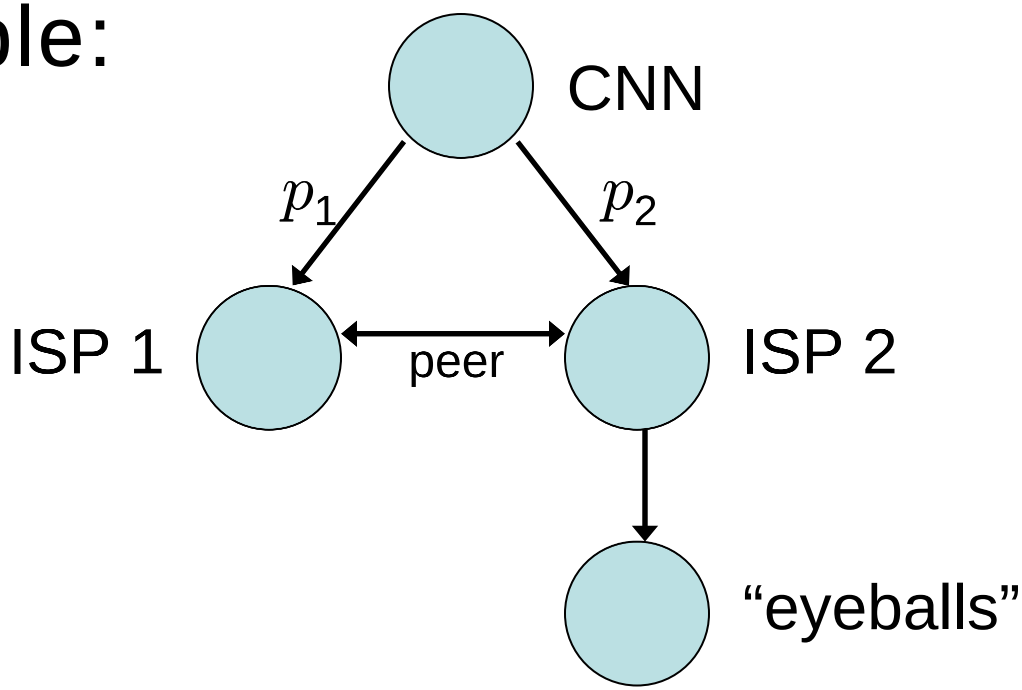

- If $p_1 < p_2$, then ISP 2’s profit = zero

The future

- Econ 101 captures the essence:
  - cutthroat pricing
  - massive financial losses
  - “last firm standing” mentality

- Question:

  Is a regulated monopoly the only endgame?

Engineering

What is the problem in Bertrand example?

ISP 2 receives no credit for the value generated.

Current protocols don't expedite transmission of value information.

$\Rightarrow$ How do we build economically robust, informative protocols?

This course

We will develop the basics of noncooperative game theory...

...but with an eye towards connection with engineering applications.

Our first game

Two players each have a budget of \$4.00.

I have \$8.00.

Each player i puts \$$w_i$ in an envelope.

I give player i a fraction $w_i/(w_1 + w_2)$ of the \$8.00 that I have.

Whatever they did not put in the envelope, they keep for themselves.

Reasoning about the game

- What is the "best" a player can do?
- What is the best they can do together?
- Should they ever bid zero?
- Is there any bid a player should *never* make?
- What is the minimum a player can guarantee himself or herself?
- *What will happen when the game is played?*

Reasoning about the game

Player 1's payoff = 8 x $w_1 / (w_1 + w_2) + 4 - w_1$

Player 1's bid \ Player 2's bid	***$0***	***$1***	***$2***	***$3***	***$4***
$0	$4.00	$4.00	$4.00	$4.00	$4.00
$1	$11.00	$7.00	$5.67	$5.00	$4.60
$2	$10.00	$7.33	$6.00	$5.20	$4.67
$3	$9.00	$7.00	$5.80	$5.00	$4.43
$4	$8.00	$6.40	$5.33	$4.57	$4.00

Reasoning about the game

Note that bidding $2 is always better than bidding $0, $3, or $4:

Player 2's bid

Player 1's bid	***$0***	***$1***	***$2***	***$3***	***$4***
~~***$0***~~	~~$4.00~~	~~$4.00~~	~~$4.00~~	~~$4.00~~	~~$4.00~~
$1	$11.00	$7.00	$5.67	$5.00	$4.60
$2	$10.00	$7.33	$6.00	$5.20	$4.67
~~***$3***~~	~~$9.00~~	~~$7.00~~	~~$5.80~~	~~$5.00~~	~~$4.43~~
~~***$4***~~	~~$8.00~~	~~$6.40~~	~~$5.33~~	~~$4.57~~	~~$4.00~~

Reasoning about the game

If we anticipate player 2 will not bid $0, $3, or $4...

Player 2's bid

Player 1's bid	***$0***	***$1***	***$2***	***$3***	***$4***
$0	$4.00	$4.00	$4.00	$4.00	$4.00
$1	$11.00	$7.00	$5.67	$5.00	$4.60
$2	$10.00	$7.33	$6.00	$5.20	$4.67
$3	$9.00	$7.00	$5.80	$5.00	$4.43
$4	$8.00	$6.40	$5.33	$4.57	$4.00

Reasoning about the game

...then we should always bid $2...

Player 1's bid \ Player 2's bid	*$0*	*$1*	*$2*	*$3*	*$4*
$0	$4.00	$4.00	$4.00	$4.00	$4.00
$1	$11.00	$7.00	$5.67	$5.00	$4.60
$2	$10.00	$7.33	$6.00	$5.20	$4.67
$3	$9.00	$7.00	$5.80	$5.00	$4.43
$4	$8.00	$6.40	$5.33	$4.57	$4.00

Reasoning about the game

...and so should player 2.

Player 2's bid

Player 1's bid

	$0	$1	$2	$3	$4
$0	$4.00	$4.00	$4.00	$4.00	$4.00
$1	$11.00	$7.00	$5.67	$5.00	$4.00
$2	$10.00	$7.33	$6.00	$5.20	$4.67
$3	$9.00	$7.00	$5.80	$5.00	$4.43
$4	$8.00	$6.40	$5.33	$4.57	$4.00

Reasoning about the game

- Does this way of reasoning about the game make sense?

- *Thought experiments:*

  What if the budgets are different?

  What if the size of the common pool is different?

MS&E 246: Lecture 2
The basics

Ramesh Johari

January 16, 2007

Course overview

(Mainly) noncooperative game theory.

Noncooperative:

Focus on individual players' incentives (note these might lead to cooperation!)

Game theory:

Analyzing the behavior of rational, self interested players

What's in a game?

1. *Players:* Who?
2. *Strategies:* What actions are available?
3. *Rules:* How? When? What do they know?
4. *Outcomes:* What results?
5. *Payoffs:*

   How do players evaluate outcomes of the game?

Example: Chess

1. *Players:* Chess masters
2. *Strategies:* Moving a piece
3. *Rules:* How pieces are moved/removed
4. *Outcomes:* Victory or defeat
5. *Payoffs:*
   Thrill of victory,
   agony of defeat

Rationality

Players are *rational* and *self-interested:*

They will *always* choose actions that maximize their payoffs, given everything they know.

Static games

We first focus on *static games.*

(one-shot games, simultaneous-move games)

For any such game, the rules say:

All players must simultaneously pick a strategy.

This immediately determines an outcome, and hence their payoff.

Knowledge

- All players know
  the *structure of the game*:

  players, strategies, rules,
  outcomes, payoffs

Common knowledge

- All players know
  the structure of the game
  - All players know all players know
    the structure
    - All players know all players know all players know
      the structure

and so on… $\Rightarrow$

We say: the structure is *common knowledge*.

This is called *complete information*.

PART I: Static games of complete information

Representation

- N : # of players
- S_n : strategies available to player n
- Outcomes:
  Composite strategy vectors
- $\Pi_n(s_1, \ldots, s_N)$:
  payoff to player n when
  player i plays strategy s_i, $i = 1, \ldots, N$

Example: A routing game

MCI and AT&T:

A Chicago customer of MCI wants to send 1 MB to an SF customer of AT&T.

A LA customer of AT&T wants to send 1 MB to an NY customer of MCI.

Providers minimize their own cost.

Key: MCI and AT&T only exchange traffic (“peer”) in NY and SF.

Example: A routing game

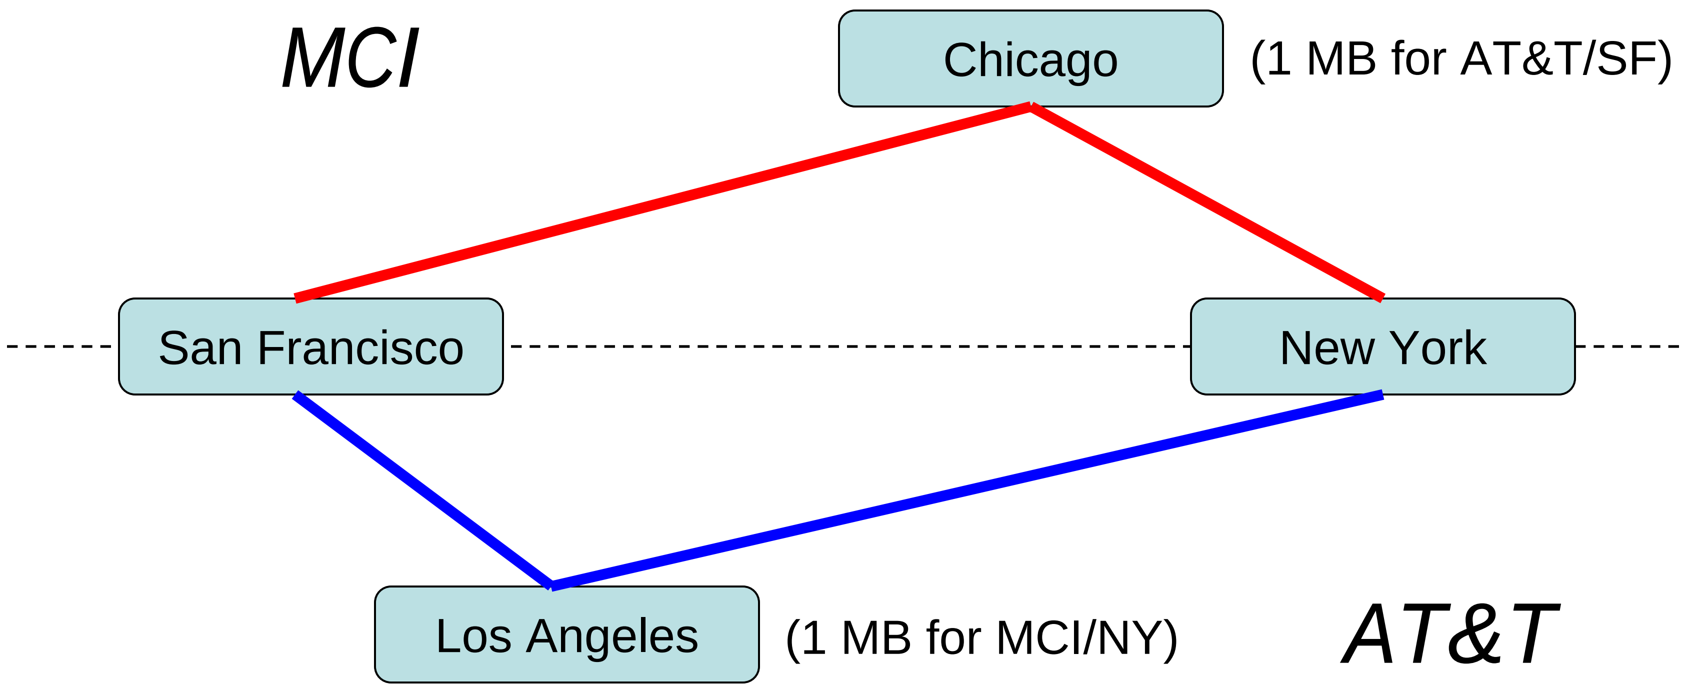

Costs (per MB): Long links = 2; Short links = 1

Example: A routing game

Players: MCI and AT&T (N = 2)

Strategies: Choice of traffic exit

$S_1 = S_2$ = { nearest exit, furthest exit }

Payoffs:

Both choose furthest exit: $\Pi_{MCI} = \Pi_{AT\&T} = -2$

Both choose nearest exit: $\Pi_{MCI} = \Pi_{AT\&T} = -4$

MCI chooses near, AT&T chooses far:

$\Pi_{MCI} = -1$, $\Pi_{AT\&T} = -5$

Example: A routing game

		AT&T	
		near	far
MCI	near	(-4,-4)	(-1,-5)
	far	(-5,-1)	(-2,-2)

Games with $N = 2$, S_n finite for each n are called *bimatrix games*.

Example: Matching pennies

Player 2

Player 1		H	T
	H	(1,-1)	(-1,1)
	T	(-1,1)	(1,-1)

This is a *zero-sum matrix game.*

Dominance

$s_n \in S_n$ is a *(weakly) dominated strategy* if

there exists $s_n^* \in S_n$ such that

$\Pi_n(s_n^*, \mathbf{s}_{-n}) \geq \Pi_n(s_n, \mathbf{s}_{-n})$,

for any choice of $\mathbf{s}_{-n}$, with
strict ineq. for at least one choice of $\mathbf{s}_{-n}$

If the ineq. is always strict, then s_n is a *strictly dominated strategy*.

Dominance

s_n* $\in S_n$ is a *weak dominant strategy* if

s_n* weakly dominates all other $s_n \in S_n$.

s_n* $\in S_n$ is a *strict dominant strategy* if

s_n* strictly dominates all other $s_n \in S_n$.

(Note: dominant strategies are unique!)

Dominant strategy equilibrium

$\mathbf{s} \in S_1 \times \cdots \times S_N$ is a

strict (or weak) dominant strategy equilibrium

if s_n is a strict (or weak) dominant strategy for each n.

Back to the routing game

		AT&T	
		near	far
MCI	near	(-4,-4)	(-1,-5)
	far	(-5,-1)	(-2,-2)

Nearest exit is strict dominant strategy for MCI.

Back to the routing game

		AT&T	
		near	far
MCI	near	(-4,-4)	(-1,-5)
	far	(-5,-1)	(-2,-2)

Nearest exit is strict dominant strategy for AT&T.

Back to the routing game

AT&T

MCI		near	far
	near	(-4,-4)	(-1,-5)
	far	(-5,-1)	(-2,-2)

Both choosing nearest exit is a strict dominant strategy equilibrium.

Example: Second price auction

- N bidders
- *Strategies:* S_n = [0, ∞); s_n = "bid"
- *Rules & outcomes:*

  High bidder wins, pays *second highest* bid
- *Payoffs:*
  - Zero if a player loses
  - If player n wins and pays t_n , then

    $\Pi_n = v_n - t_n$
  - v_n : *valuation* of player n

Example: Second price auction

- *Claim: Truthful bidding* ($s_n = v_n$) is a weak dominant strategy for player n.

- *Proof:*

  If player n considers a bid > v_n :

  Payoff may be lower when n wins, and the same (zero) when n loses

Example: Second price auction

- *Claim: Truthful bidding* ($s_n = v_n$) is a weak dominant strategy for player n.

- *Proof:*

  If player n considers a bid < v_n :

    Payoff will be same when n wins,
    but may be worse when n loses

Example: Second price auction

- We conclude:

  Truthful bidding is a (weak) dominant strategy equilibrium for the second price auction.

Example: Matching pennies

		Player 2	
		H	T
Player 1	H	(1,-1)	(-1,1)
	T	(-1,1)	(1,-1)

No dominant/dominated strategy exists!

Moral: *Dominant strategy eq. may not exist.*

Iterated strict dominance

Given a game:

- Construct a new game by *removing* a strictly dominated strategy from one of the strategy spaces S_n .
- Repeat this procedure until no strictly dominated strategies remain.

If this results in a unique strategy profile, the game is called *dominance solvable.*

Iterated strict dominance

- Note that the bidding game in Lecture 1 was dominance solvable.

- There the unique resulting strategy profile was (6,6).

Example

Player 2

Player 1		Left	Middle	Right
	Up	(1,0)	(1,2)	(0,1)
	Down	(0,3)	(0,1)	(2,0)

Example

		Player 2		
		Left	Middle	Right
Player 1	Up	(1,0)	(1,2)	(0,1)
	Down	(0,3)	(0,1)	(2,0)

Example

		Player 2		
		Left	Middle	Right
Player 1	Up	(1,0)	(1,2)	(0,1)
	Down	(0,3)	(0,1)	(2,0)

Example

Player 2

		Left	Middle	Right
Player 1	Up	(1,0)	(1,2)	(0,1)
	Down	(0,3)	(0,1)	(2,0)

Thus the game is dominance solvable.

Example: Cournot duopoly

- Two firms ($N = 2$)
- *Cournot competition*:
  each firm chooses a quantity $s_n \geq 0$
- Cost of producing s_n : $c\, s_n$
- *Demand curve*:
  Price = $P(s_1 + s_2) = a - b\,(s_1 + s_2)$
- Payoffs:
  Profit = $\Pi_n(s_1,\ s_2) = P(s_1 + s_2)\ s_n - c\ s_n$

Example: Cournot duopoly

- *Claim:*

  The Cournot duopoly is dominance solvable.

- *Proof technique*:

  First construct the *best response* for each player.

Best response

Best response set for player n to $\mathbf{s}_{-n}$:

$$R_n(\mathbf{s}_{-n}) = \text{arg max}_{s_n \in S_n} \Pi_n(s_n, \mathbf{s}_{-n})$$

[Note: arg max$_{x \in X}$ $f(x)$ is the *set* of x that maximize $f(x)$]

Example: Cournot duopoly

Calculating the best response given s_{-n}:

$$\max_{s_n \geq 0} \left[(a - bs_n - bs_{-n})s_n - cs_n\right] \implies$$

Differentiate and solve:

$$a - c - bs_{-n} - 2bs_n = 0$$

So:

$$R_n(s_{-n}) = \left[\frac{a-c}{2b} - \frac{s_{-n}}{2}\right]^+$$

Example: Cournot duopoly

For simplicity, let $t = (a - c)/b$

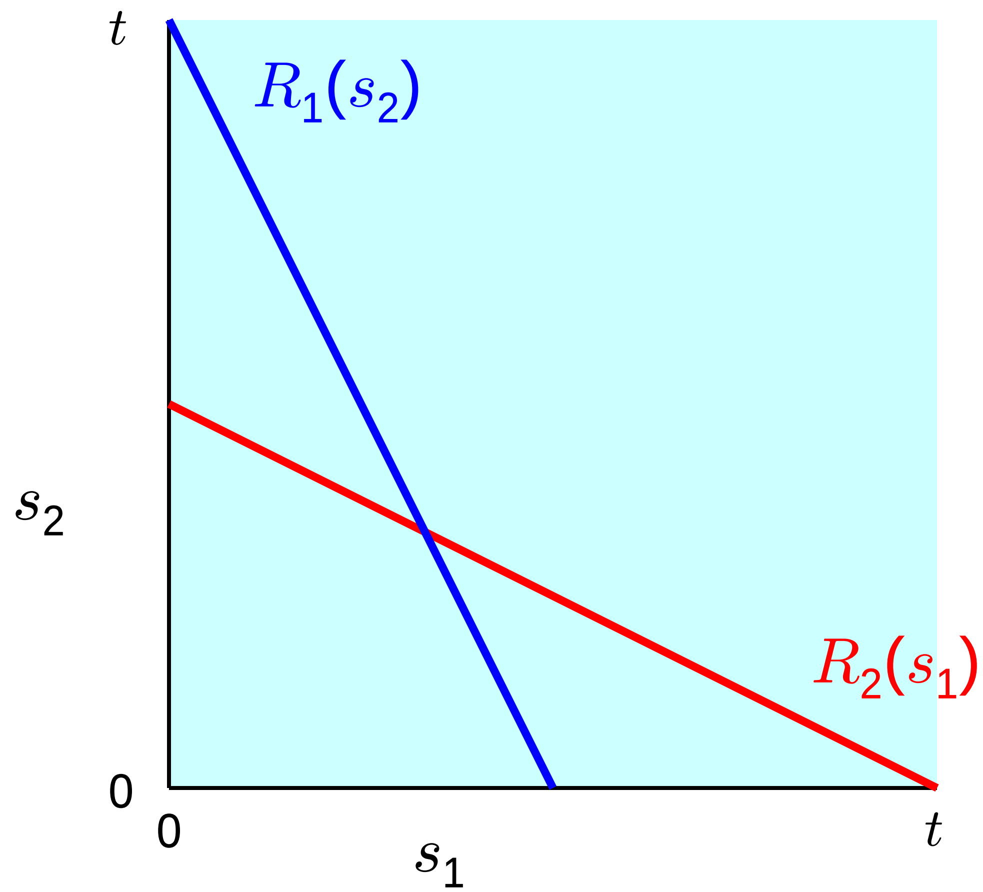

Example: Cournot duopoly

Step 1: Remove strictly dominated s_1.

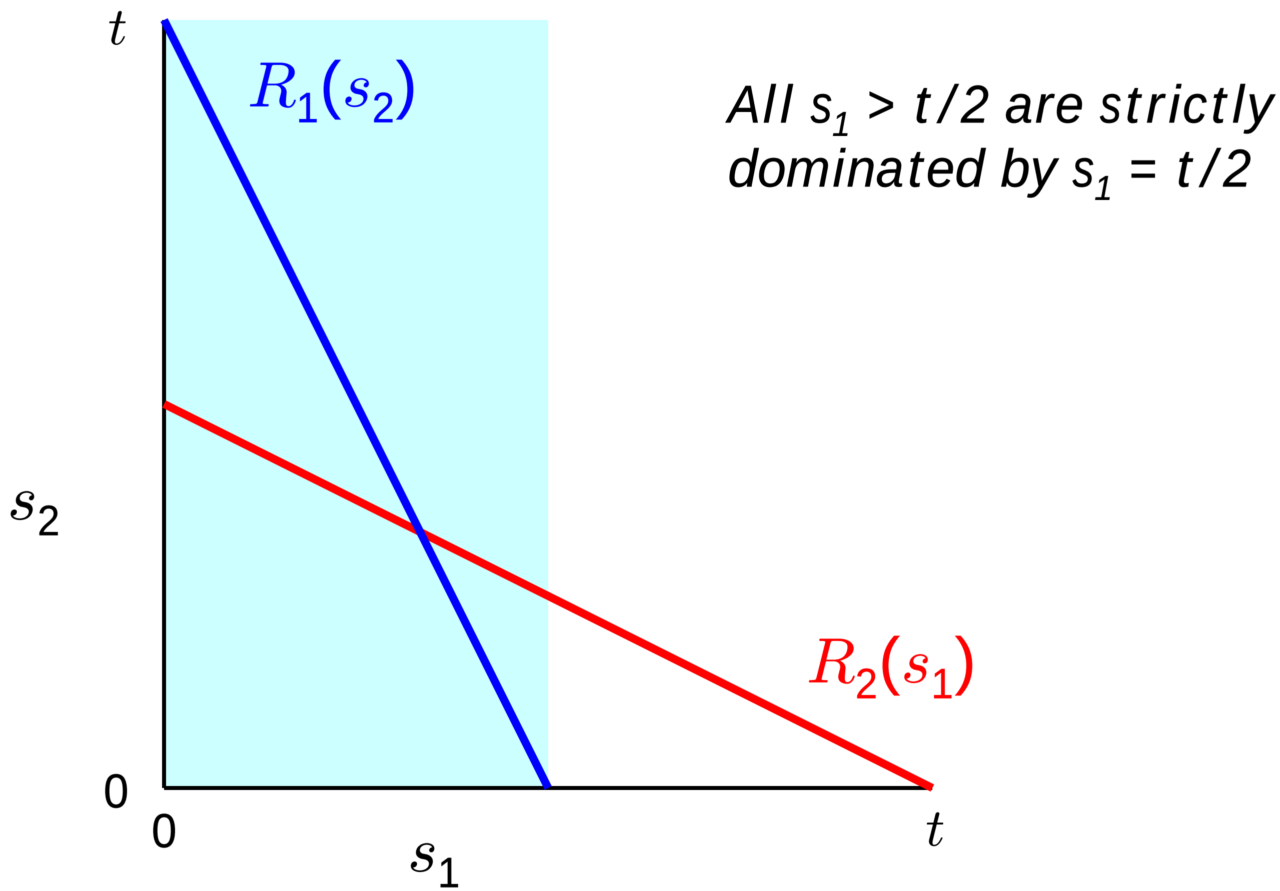

Example: Cournot duopoly

Step 2: Remove strictly dominated s_2.

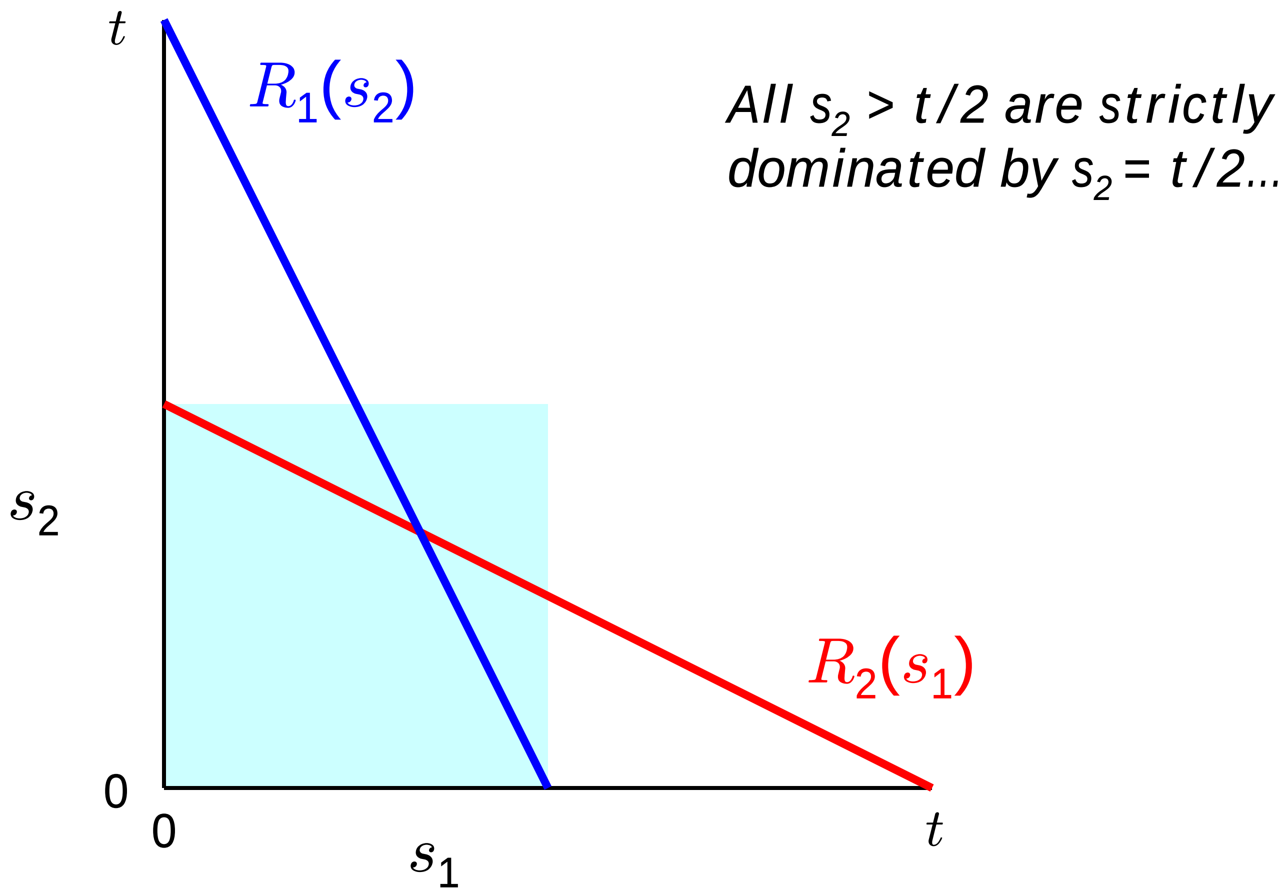

Example: Cournot duopoly

Step 2: Remove strictly dominated s_2.

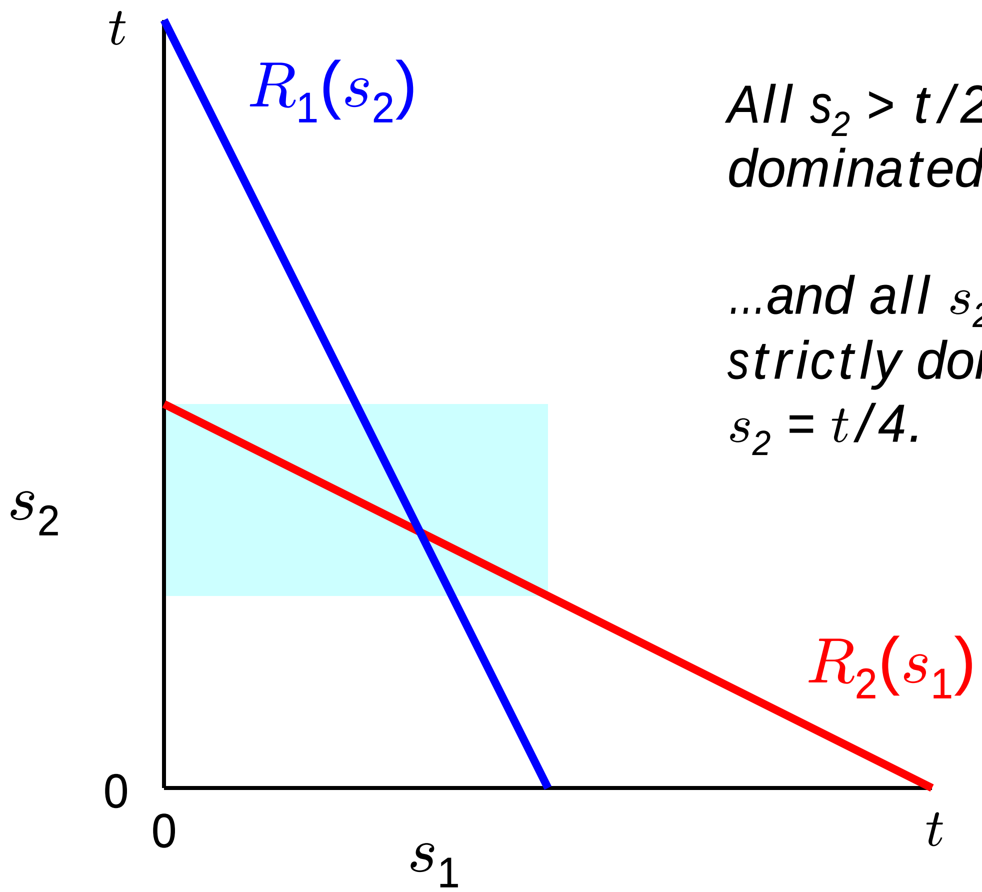

All $s_2 > t/2$ are strictly dominated by $s_2 = t/2$...

...and all $s_2 < t/4$ are strictly dominated by $s_2 = t/4$.

Example: Cournot duopoly

Step 3: Remove strictly dominated s_1.

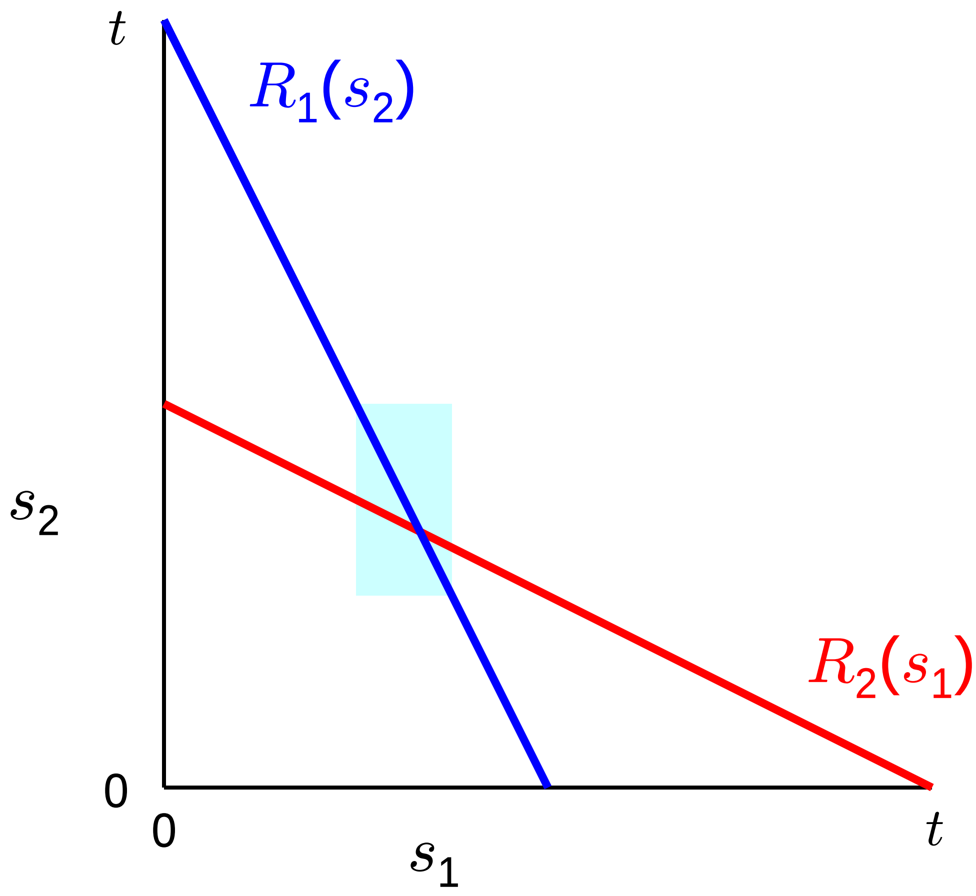

Example: Cournot duopoly

Step 4: Remove strictly dominated s_2.

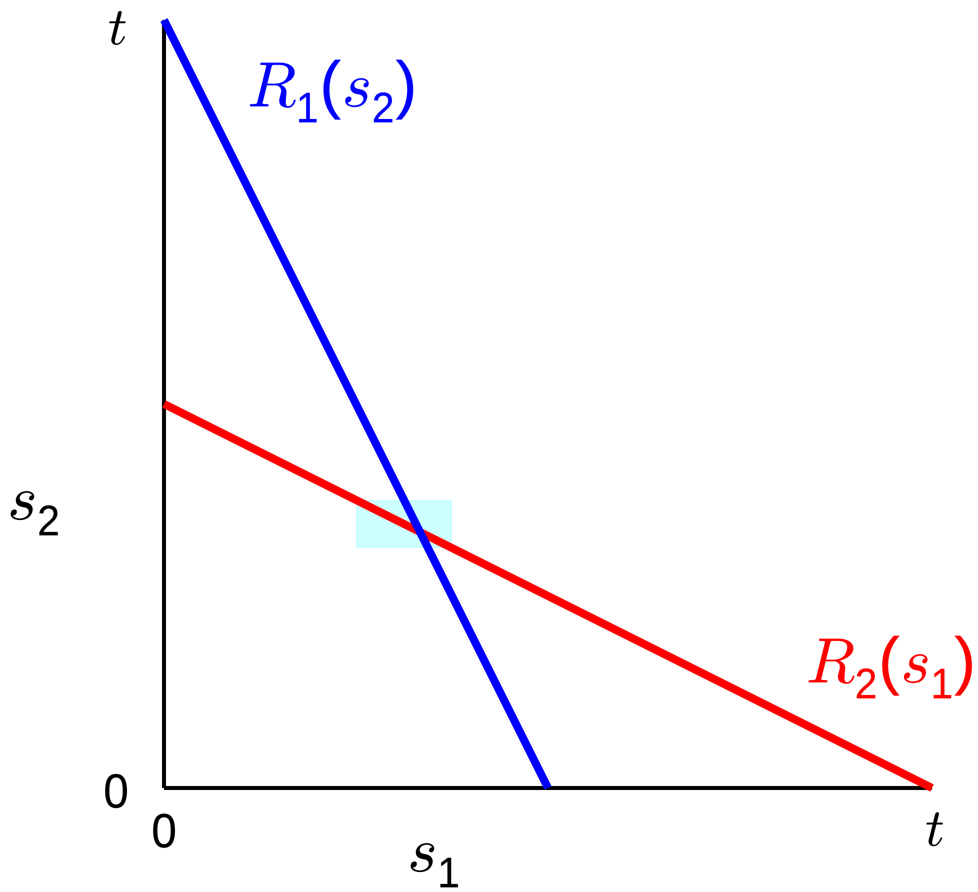

Example: Cournot duopoly

The process converges to the intersection point: $s_1 = t/3$, $s_2 = t/3$

Step #	Undominated s_1
1	$[0, t/2]$
3	$[t/4, 3t/8]$
5	$[5t/16, 11t/32]$
7	$[21t/64, 43t/128]$

Example: Cournot duopoly

Lower bound =

$$t \sum_{k=1}^{\infty} (1/2)^{2k} = t/3.$$

Upper bound =

$$t \left(1 - \sum_{k=1}^{\infty} (1/2)^{2k-1} \right) = t/3.$$

Dominance solvability: comments

- Order of elimination doesn’t matter
- Just as most games don’t have DSE, most games are not dominance solvable

Rationalizable strategies

Given a game:

- For each player n, remove strategies from each S_n that are not best responses for *any* choice of other players' strategies.
- Repeat this procedure.

Strategies that survive this process are called *rationalizable strategies.*

Rationalizable strategies

In a two player game, a strategy s_1 is rationalizable for player 1 if there exists a *chain of justification*

$$s_1 \to s_2 \to s_1' \to s_2' \to \ldots \to s_1$$

where each is a best response to the one before.

Rationalizable strategies

- If s_n is rationalizable, it also survives iterated strict dominance. (Why?)

$\Rightarrow$ For a dominance solvable game, there is a unique rationalizable strategy, and it is the one given by iterated strict dominance.

Rationalizability: example

Note that M is not *rationalizable*,
but it survives iterated strict dominance.

		Player 2		
		L	M	R
Player 1	T	(1,0)	(1,1)	(1,5)
	B	(1,5)	(1,1)	(1,0)

Rationalizable strategies

Note for later:

When “mixed” strategies are allowed,

rationalizability = iterated strict dominance

for two player games.

MS&E 246: Lecture 3
Pure strategy Nash equilibrium

Ramesh Johari

January 16, 2007

Outline

- Best response and pure strategy Nash equilibrium
- Relation to other equilibrium notions
- Examples
- Bertrand competition

Best response set

Best response set for player n to $\mathbf{s}_{-n}$:

$$R_n(\mathbf{s}_{-n}) = \text{arg max}_{s_n \in S_n} \Pi_n(s_n, \mathbf{s}_{-n})$$

[Note: arg max$_{x \in X}$ $f(x)$ is the *set* of x that maximize $f(x)$]

Nash equilibrium

Given: N-player game

A vector $\mathbf{s}$ = (s_1, ..., s_N) is a *(pure strategy) Nash equilibrium* if:

$$s_i \in R_i(\mathbf{s}_{-i})$$

for all players i.

Each *individual* plays a best response to the others.

Nash equilibrium

Pure strategy Nash equilibrium is robust to *unilateral deviations*

One of the hardest questions in game theory:

How do players know to play a Nash equilibrium?

Example: Prisoner's dilemma

Recall the routing game:

		AT&T	
		near	far
MCI	near	(-4,-4)	(-1,-5)
	far	(-5,-1)	(-2,-2)

Example: Prisoner's dilemma

Here (near,near) is the unique (pure strategy) NE:

		AT&T	
		near	far
MCI	near	(-4,-4)	(-1,-5)
	far	(-5,-1)	(-2,-2)

Summary of relationships

Given a game:

- Any DSE also survives ISD, and is a NE.

(DSE = dominant strategy equilibrium; ISD = iterated strict dominance)

Example: bidding game

Recall the bidding game from lecture 1:

			Player 2's bid		
Player 1's bid	***$0***	***$1***	***$2***	***$3***	***$4***
$0	$4.00	$4.00	$4.00	$4.00	$4.00
$1	$11.00	$7.00	$5.67	$5.00	$4.60
$2	$10.00	$7.33	$6.00	$5.20	$4.67
$3	$9.00	$7.00	$5.80	$5.00	$4.43
$4	$8.00	$6.40	$5.33	$4.57	$4.00

Example: bidding game

Here (2,2) is the unique (pure strategy) NE:

Player 2's bid

Player 1's bid

	$0	**$1**	**$2**	**$3**	**$4**
$0	$4.00	$4.00	$4.00	$4.00	$4.00
$1	$11.00	$7.00	$5.67	$5.00	$4.60
$2	$10.00	$7.33	$6.00	$5.20	$4.67
$3	$9.00	$7.00	$5.80	$5.00	$4.43
$4	$8.00	$6.40	$5.33	$4.57	$4.00

Summary of relationships

Given a game:

- Any DSE also survives ISD, and is a NE.
- If a game is dominance solvable, the resulting strategy vector is a NE

  Another example of this: the Cournot game.
- Any NE survives ISD (and is also rationalizable).

(DSE = dominant strategy equilibrium; ISD = iterated strict dominance)

Example: Cournot duopoly

Unique NE: $(t/3\ ,\ t/3)$

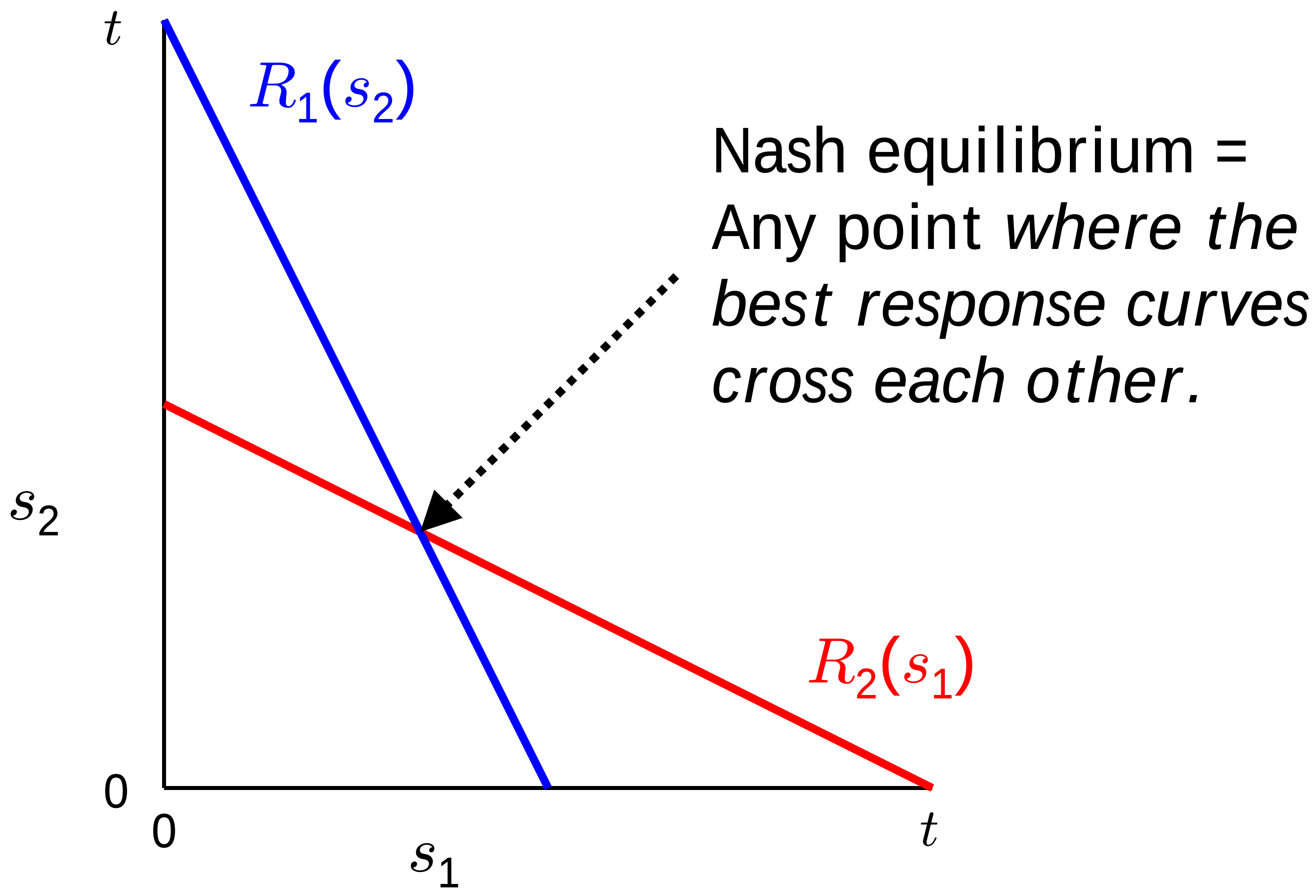

Example: coordination game

Two players trying to *coordinate* their actions:

		Player 2	
		L	*R*
Player 1	*l*	(2,1)	(0,0)
	r	(0,0)	(1,2)

Example: coordination game

Best response of player 1:

$R_1(L) = \{ l \}$, $R_1(R) = \{ r \}$

Player 2

		L	*R*
Player 1	*l*	(<u>2</u>,1)	(0,0)
	r	(0,0)	(<u>1</u>,2)

Example: coordination game

Best response of player 2:

$R_2(l) = \{ L \}$, $R_2(r) = \{ R \}$

Player 2

		L	*R*
Player 1	*l*	(2,<u>1</u>)	(0,0)
	r	(0,0)	(1,<u>2</u>)

Example: coordination game

Two Nash equilibria: (*l*, *L*) and (*r*, *R*).

Moral: NE is not a *unique predictor of play!*

		Player 2	
		L	*R*
Player 1	*l*	(2,1)	(0,0)
	r	(0,0)	(1,2)

Example: matching pennies

No pure strategy NE for this game

Moral: Pure strategy NE may not exist.

		Player 2	
		H	T
Player 1	H	(1,-1)	(-1,1)
	T	(-1,1)	(1,-1)

Example: Bertrand competition

- In *Cournot* competition, firms choose the *quantity* they will produce.

- In *Bertrand* competition, firms choose the *prices* they will charge.

Bertrand competition: model

- Two firms
- Each firm i chooses a price $p_i \geq 0$
- Each unit produced incurs a cost $c \geq 0$
- Consumers only buy from the producer offering the *lowest price*
- Demand is $D > 0$

Bertrand competition: model

- Two firms
- Each firm i chooses a price p_i
- Profit of firm i:

$$\Pi_i(p_1, p_2) = (p_i - c)D_i(p_1, p_2)$$

where

$$D_i(p_1, p_2) = \begin{cases} 0, & \text{if } p_i > p_{-i} \\ D, & \text{if } p_i < p_{-i} \\ \frac{1}{2} D, & \text{if } p_i = p_{-i} \end{cases}$$

Bertrand competition: analysis

Suppose firm 2 sets a price = $p_2 < c$.

What is the *best response set* of firm 1?

Firm 1 wants to price higher than p_2.

$$R_1(p_2) = (p_2, \infty)$$

Bertrand competition: analysis

Suppose firm 2 sets a price = $p_2 > c$.

What is the *best response set* of firm 1?

Firm 1 wants to price slightly lower than p_2

... but there is **no** *best* **response***!*

$$R_1(p_2) = \emptyset$$

Bertrand competition: analysis

Suppose firm 2 sets a price = p_2 = c.

What is the *best response set* of firm 1?

Firm 1 wants to price at or higher than c.

$$R_1(p_2) = [c, \infty)$$

Bertrand competition: analysis

Best response of firm 1:

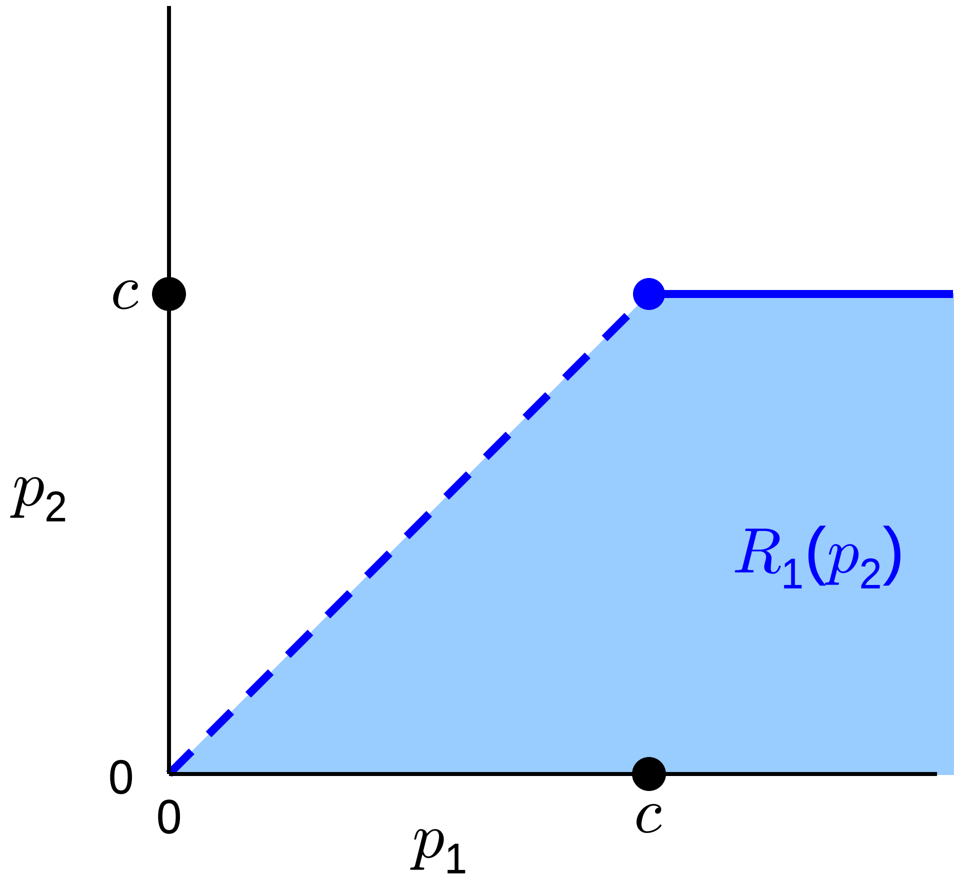

Bertrand competition: analysis

Best response of firm 2:

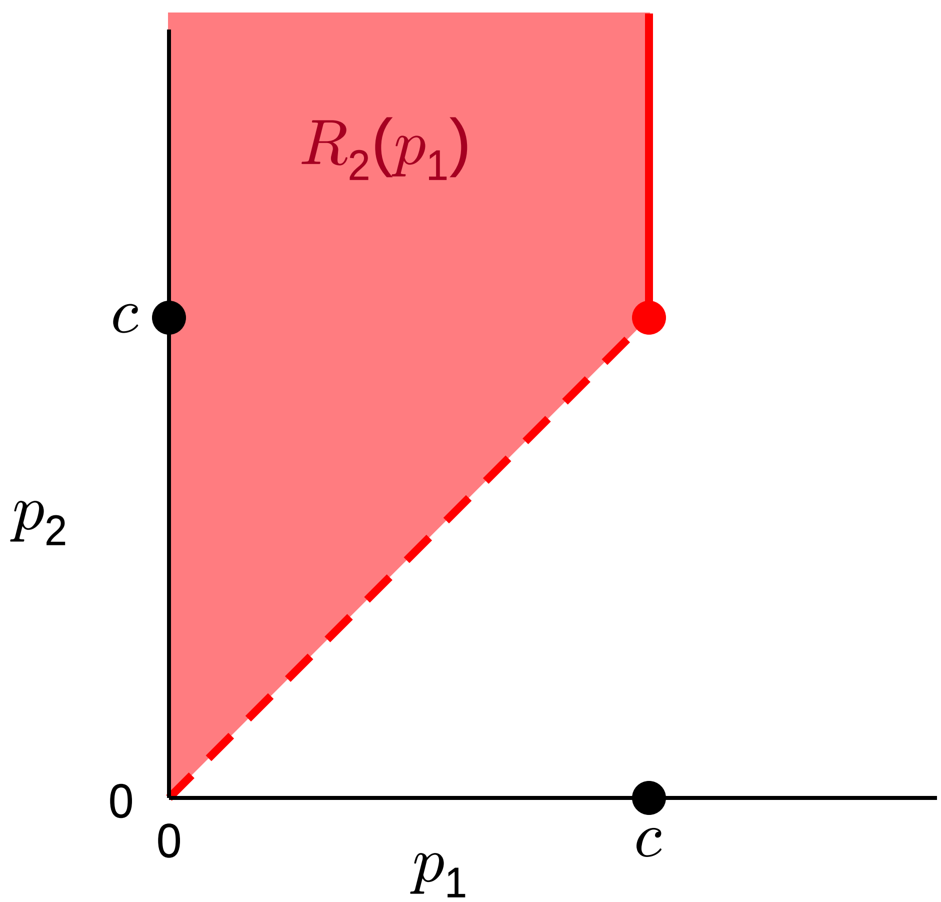

Bertrand competition: analysis

Where do they “cross”?

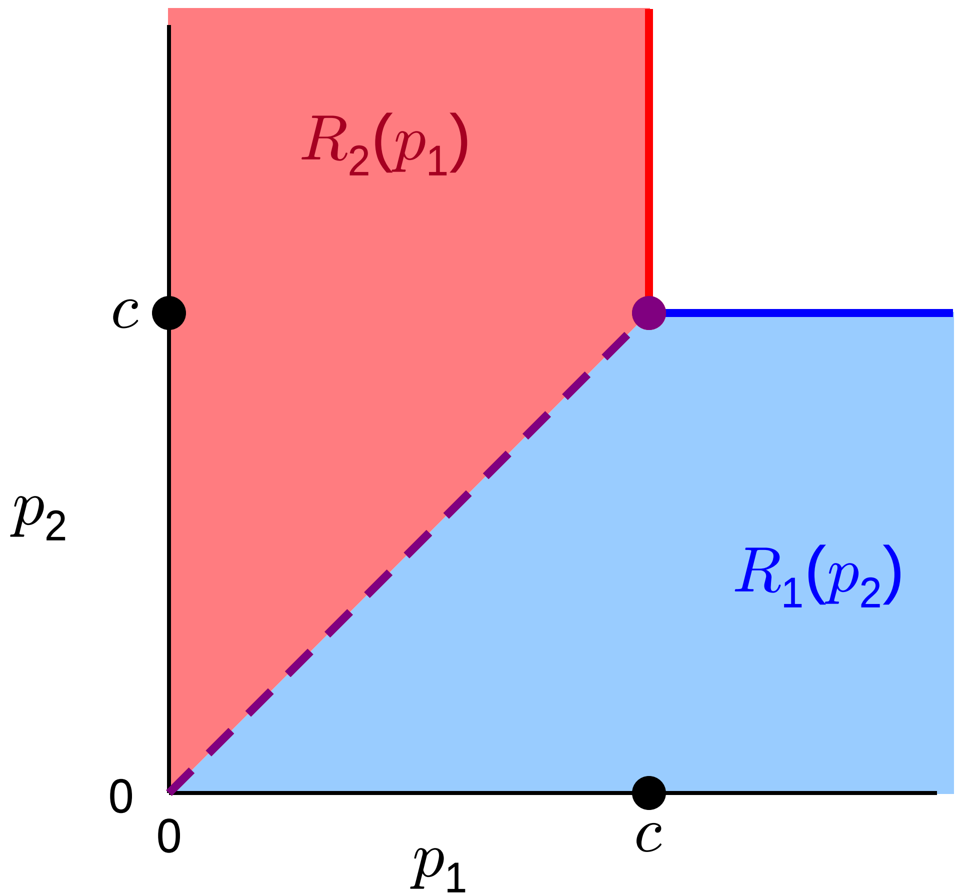

Bertrand competition: analysis

Thus the unique NE is where $p_1 = c$, $p_2 = c$.

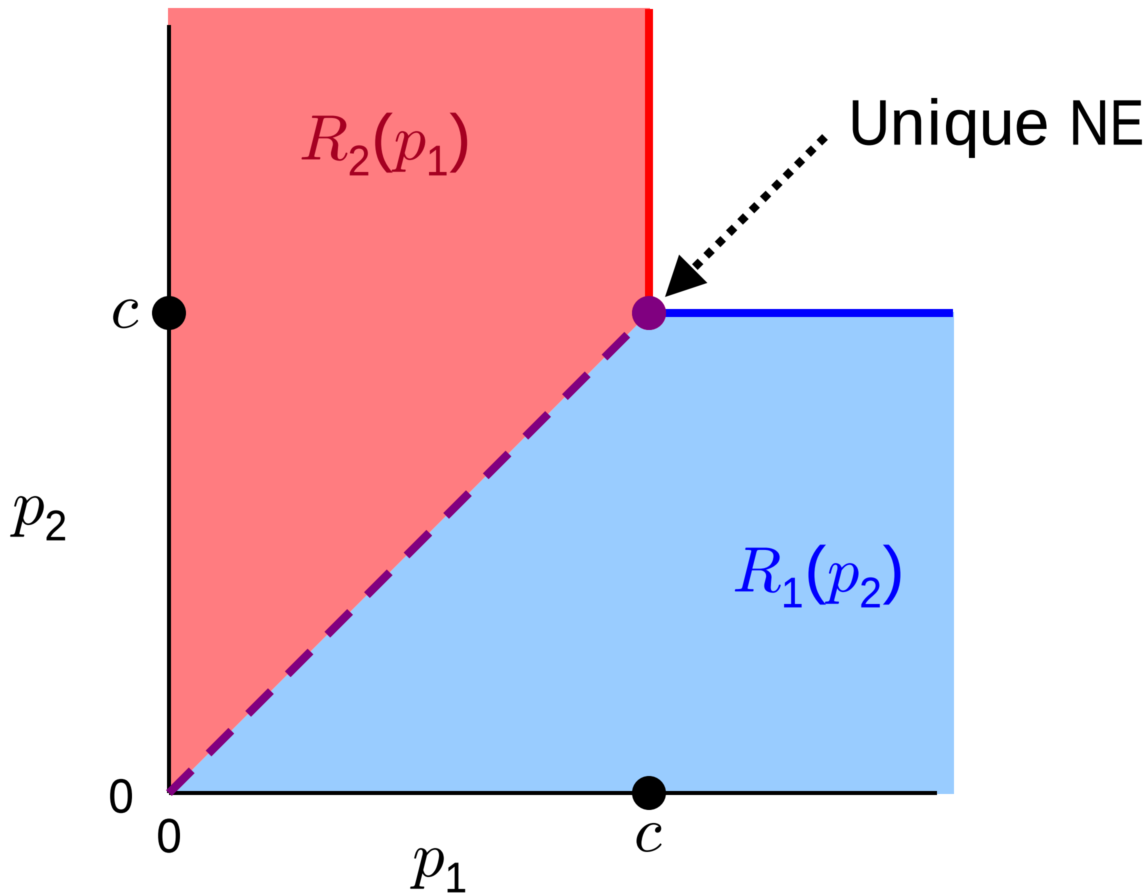

Bertrand competition

Straightforward to show:

The same result holds if demand depends on price, i.e., if the demand at price p is $D(p) > 0$.

Proof technique:

(1) Show $p_i < c$ *is never played in a NE.*

(2) Show if $c < p_1 < p_2$, *then firm 2 prefers to lower* p_2.

(3) Show if $c < p_1 = p_2$, *then firm 2 prefers to lower* p_2

Bertrand competition

What happens if $c_1 < c_2$?

No pure NE exists; however, an ε-NE exists:

Each player is happy as long as they are within ε of their optimal payoff.

ε-NE : $p_2 = c_2$, $p_1 = c_2 - \delta$

(where δ is infinitesimal)

Bertrand vs. Cournot

Assume demand is D(p) = a - p.

Interpretation: D(p) denotes the total number of consumers willing to pay *at least* p for the good.

Then the *inverse demand* is

$$P(Q) = a - Q.$$

This is the market-clearing price at which Q total units of supply would be sold.

Bertrand vs. Cournot

Assume demand is D(p) = a - p.

Then the *inverse demand* is

$$P(Q) = a - Q.$$

Assume $c < a$.

Bertrand eq.: $p_1 = p_2 = c$

Cournot eq: $q_1 = q_2 = (a - c)/3$

$\Rightarrow$ Cournot price = $a/3 + 2c/3 > c$

Bertrand vs. Cournot

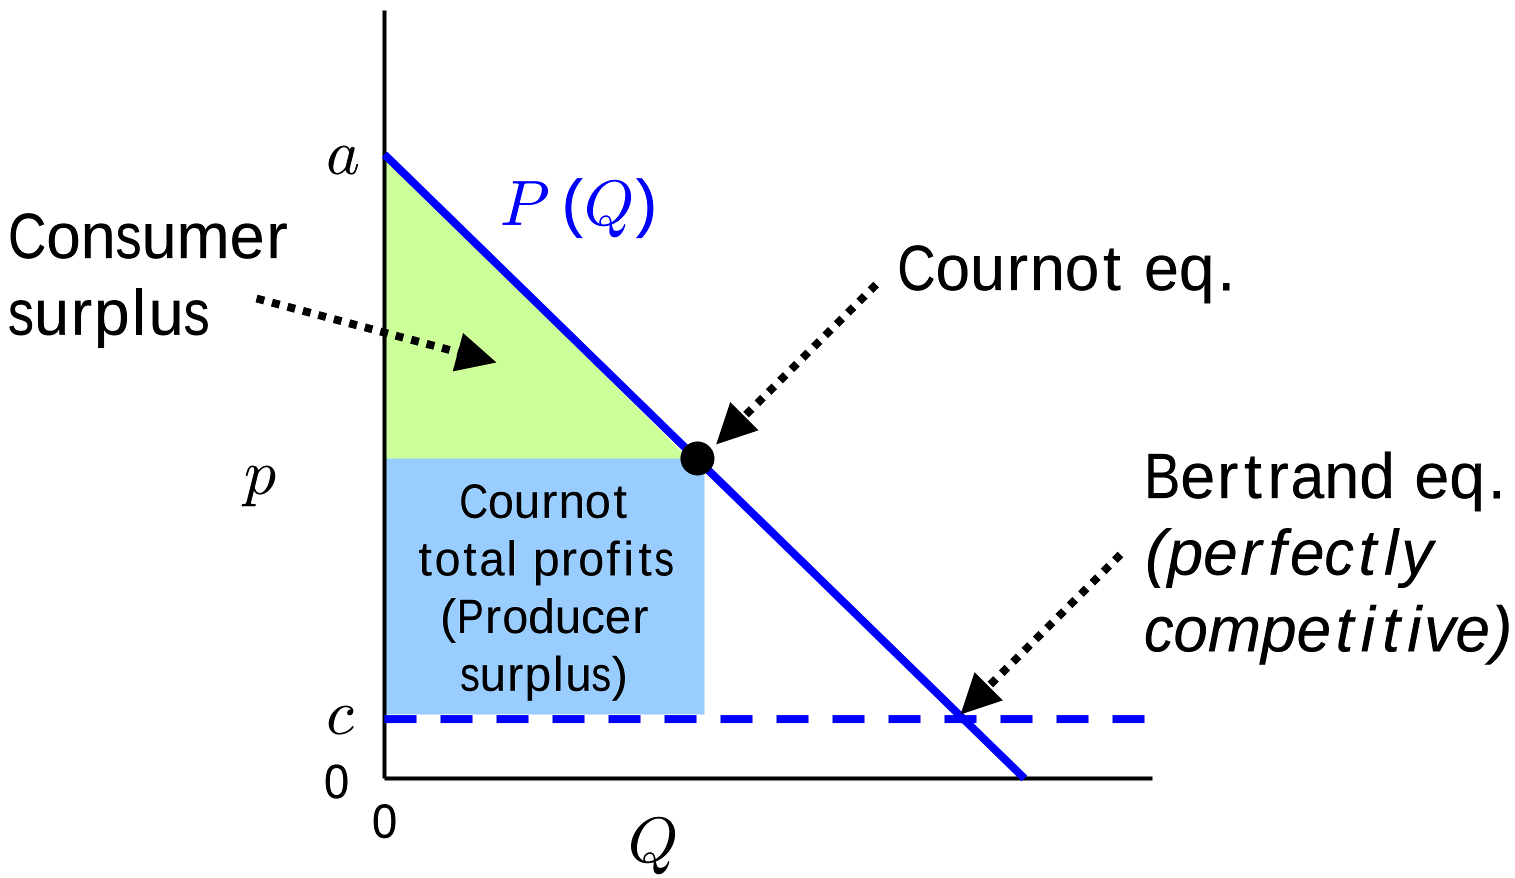

Bertrand vs. Cournot

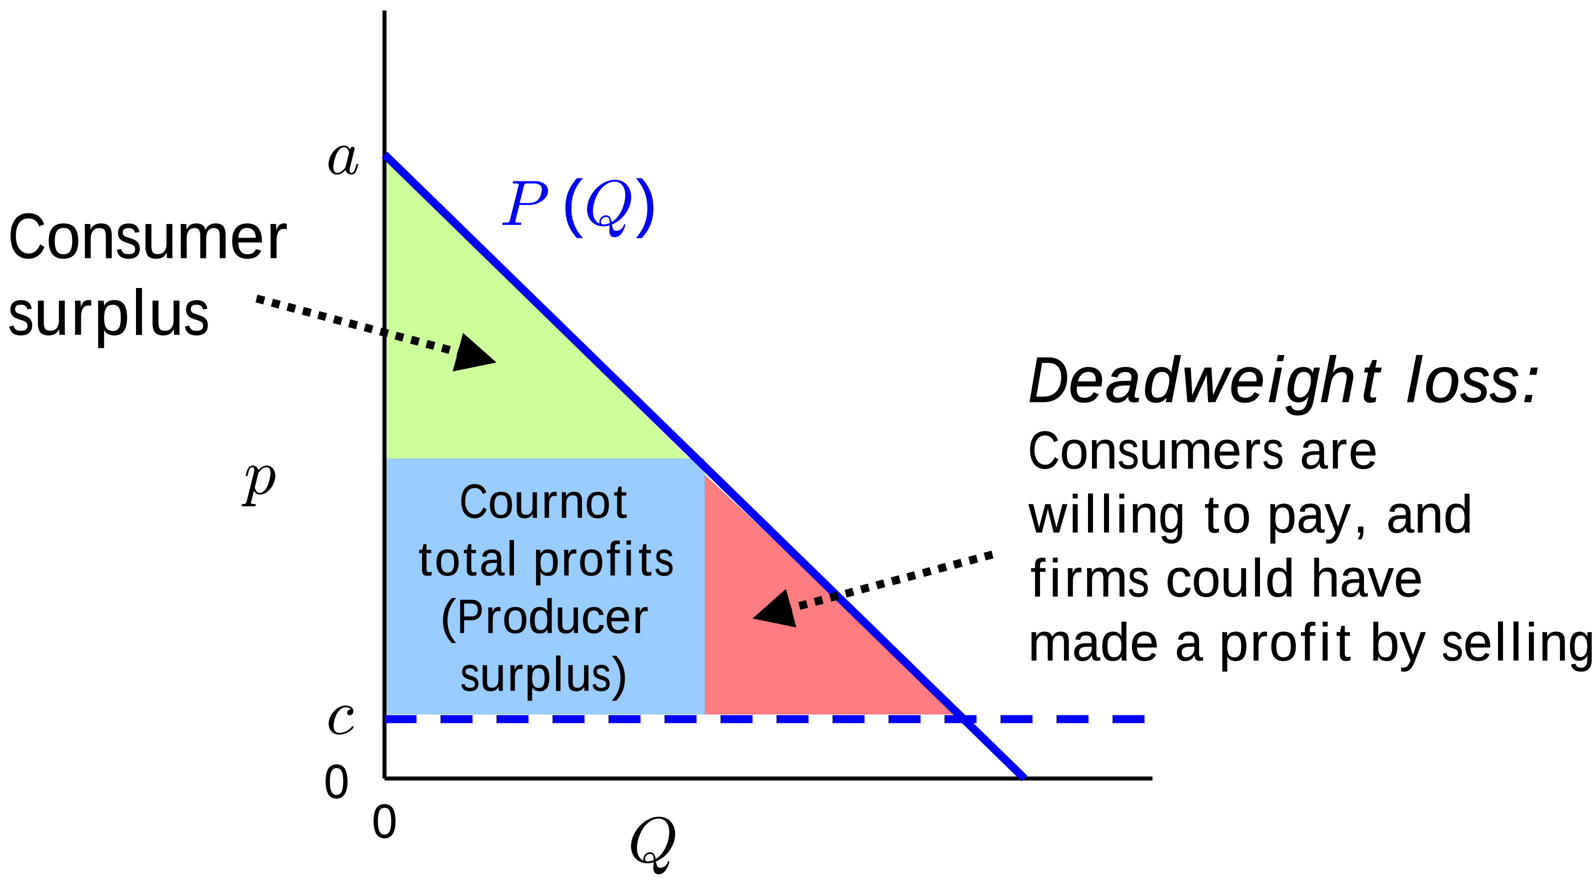

Bertrand vs. Cournot

- Cournot eq. price > Bertrand eq. price
- Bertrand price =

  marginal cost of production
- In Cournot eq., there is positive deadweight loss.

  This is because firms have *market power: they anticipate their effect on prices.*

Questions to think about

- Can a *weakly dominated* strategy be played in a Nash equilibrium?
- Can a *strictly dominated* strategy be played in a Nash equilibrium?
- Why is any NE rationalizable?
- What are real-world examples of Bertrand competition?
  Cournot competition?

Summary: Finding NE

Finding NE is typically a matter of checking the definition.

Two basic approaches...

Finding NE: Approach 1

First approach to finding NE:

(1) Compute the complete best response mapping for each player.

(2) Find where they intersect each other (graphically or otherwise).

Finding NE: Approach 2

Second approach to finding NE:

Fix a strategy vector (s_1, ..., s_N).

Check if any player has a *profitable deviation.*

If so, it cannot be a NE.

If not, it is an NE.

MS&E 246: Lecture 4
Mixed strategies

Ramesh Johari

January 18, 2007

Outline

- Mixed strategies
- Mixed strategy Nash equilibrium
- Existence of Nash equilibrium
- Examples
- Discussion of Nash equilibrium

Mixed strategies

Notation:

Given a set X, we let $\Delta(X)$ denote the set of all *probability distributions* on X.

Given a strategy space S_i for player i, the *mixed strategies* for player i are $\Delta(S_i)$.

Idea: a player can randomize over *pure strategies*.

Mixed strategies

How do we interpret mixed strategies?

Note that players only play *once*; so mixed strategies reflect *uncertainty* about what the other player might play.

Payoffs

Suppose for each player i, $\mathbf{p}_i$ is a mixed strategy for player i; i.e., it is a distribution on S_i.

We extend Π_i by taking the *expectation*:

$$\Pi_i(\mathbf{p}_1, \ldots, \mathbf{p}_N) = \sum_{s_1 \in S_1} \cdots \sum_{s_N \in S_N} p_1(s_1) \cdots p_N(s_N) \Pi_i(s_1, \ldots, s_N)$$

Mixed strategy Nash equilibrium

Given a game (N, S_1, ..., S_N, Π_1, ..., Π_N):

Create a new game with N players,
strategy spaces $\Delta(S_1)$, ..., $\Delta(S_N)$,
and expected payoffs Π_1, ..., Π_N.

A *mixed strategy Nash equilibrium*
is a Nash equilibrium of this new game.

Mixed strategy Nash equilibrium

Informally:

All players can randomize over available strategies.

In a mixed NE, player i's mixed strategy must maximize his *expected payoff*, given all other player's mixed strategies.

Mixed strategy Nash equilibrium

Key observations:

(1) All our definitions -- dominated strategies, iterated strict dominance, rationalizability -- extend to mixed strategies.

Note: any *dominant* strategy must be a *pure strategy.*

Mixed strategy Nash equilibrium

(2) We can extend the definition of *best response set* identically:

$R_i(\mathbf{p}_{-i})$ is the set of mixed strategies for player i that maximize the expected payoff $\Pi_i(\mathbf{p}_i, \mathbf{p}_{-i})$.

Mixed strategy Nash equilibrium

(2) Suppose $\mathbf{p}_i \in R_i(\mathbf{p}_{-i})$, and $p_i(s_i) > 0$.
Then $s_i \in R_i(\mathbf{p}_{-i})$.

(If not, player i could improve his payoff by not placing any weight on s_i at all.)

Mixed strategy Nash equilibrium

(3) It follows that $R_i(\mathbf{p}_{-i})$ can be constructed as follows:

(a) First find all *pure strategy* best responses to $\mathbf{p}_{-i}$; call this set $T_i(\mathbf{p}_{-i}) \subset S_i$.

(b) Then $R_i(\mathbf{p}_{-i})$ is the set of all probability distributions over T_i, i.e.:

$$R_i(\mathbf{p}_{-i}) = \Delta(T_i(\mathbf{p}_{-i}))$$

Mixed strategy Nash equilibrium

Moral:

A mixed strategy $\mathbf{p}_i$ *is*
a best response to $\mathbf{p}_{\text{-}i}$

if and only if

every s_i *with* $p_i(s_i) > 0$ *is*
a best response to $\mathbf{p}_{\text{-}i}$

Example: coordination game

We'll now apply this insight to the coordination game.

		Player 2	
		L	*R*
Player 1	*l*	(2,1)	(0,0)
	r	(0,0)	(1,2)

Example: coordination game

Suppose player 1 puts probability p_1 on *l* and probability 1 - p_1 on *r*.

Suppose player 2 puts probability p_2 on *L* and probability 1 - p_2 on *R*.

We want to find *all* Nash equilibria (pure and mixed).

Example: coordination game

- Step 1: Find best response mapping of player 1.

  Given p_2:

  $\Pi_1(l, \mathbf{p}_2) = 2\, p_2$

  $\Pi_1(r, \mathbf{p}_2) = 1 - p_2$

Example: coordination game

- Step 1: Find best response mapping of player 1.

If p_2 is:	Then best response is:
< 1/3	*r* (p_1 = 0)
> 1/3	*l* (p_1 = 1)
= 1/3	anything ($0 \leq p_1 \leq 1$)

Example: coordination game

Best response of player 1:

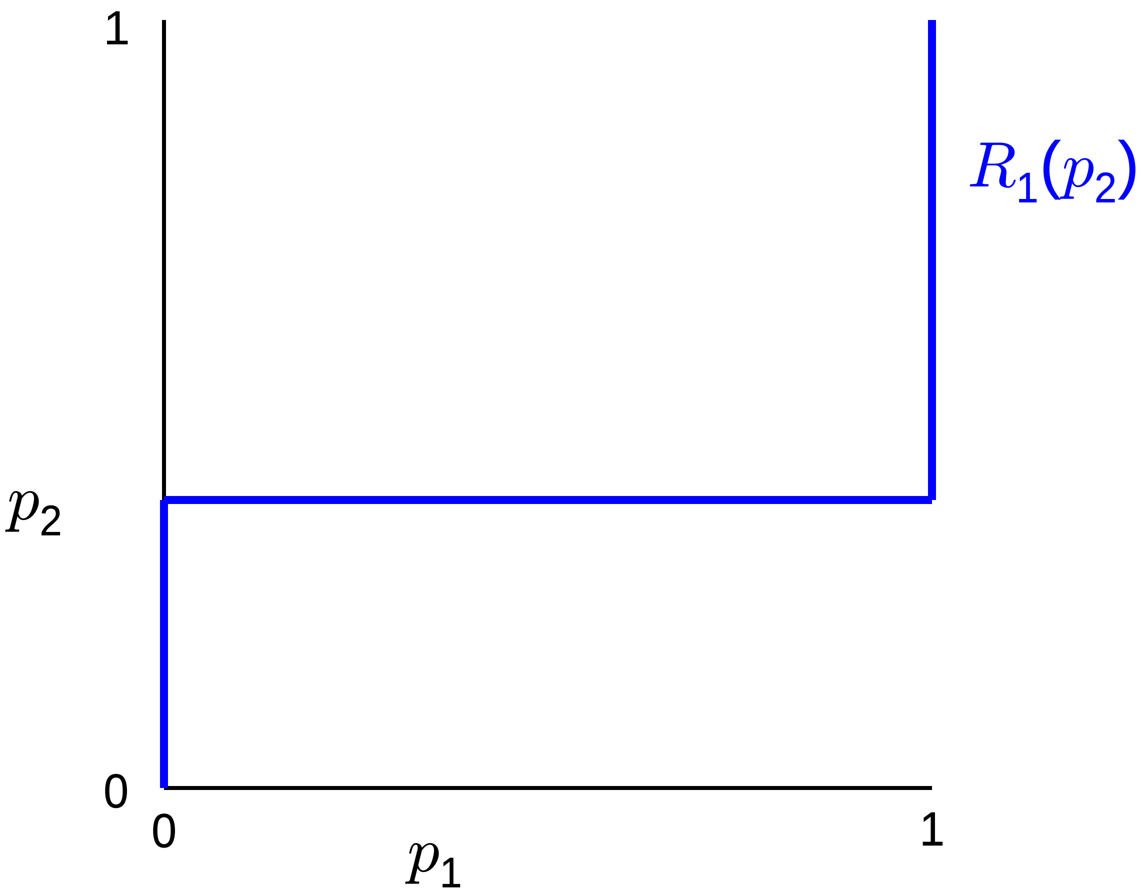

Example: coordination game

- Step 2: Find best response mapping of player 2.

If p_1 is:	Then best response is:
< 2/3	*R* (p_2 = 0)
> 2/3	*L* (p_2 = 1)
= 2/3	anything ($0 \leq p_1 \leq 1$)

Example: coordination game

Best response of player 2:

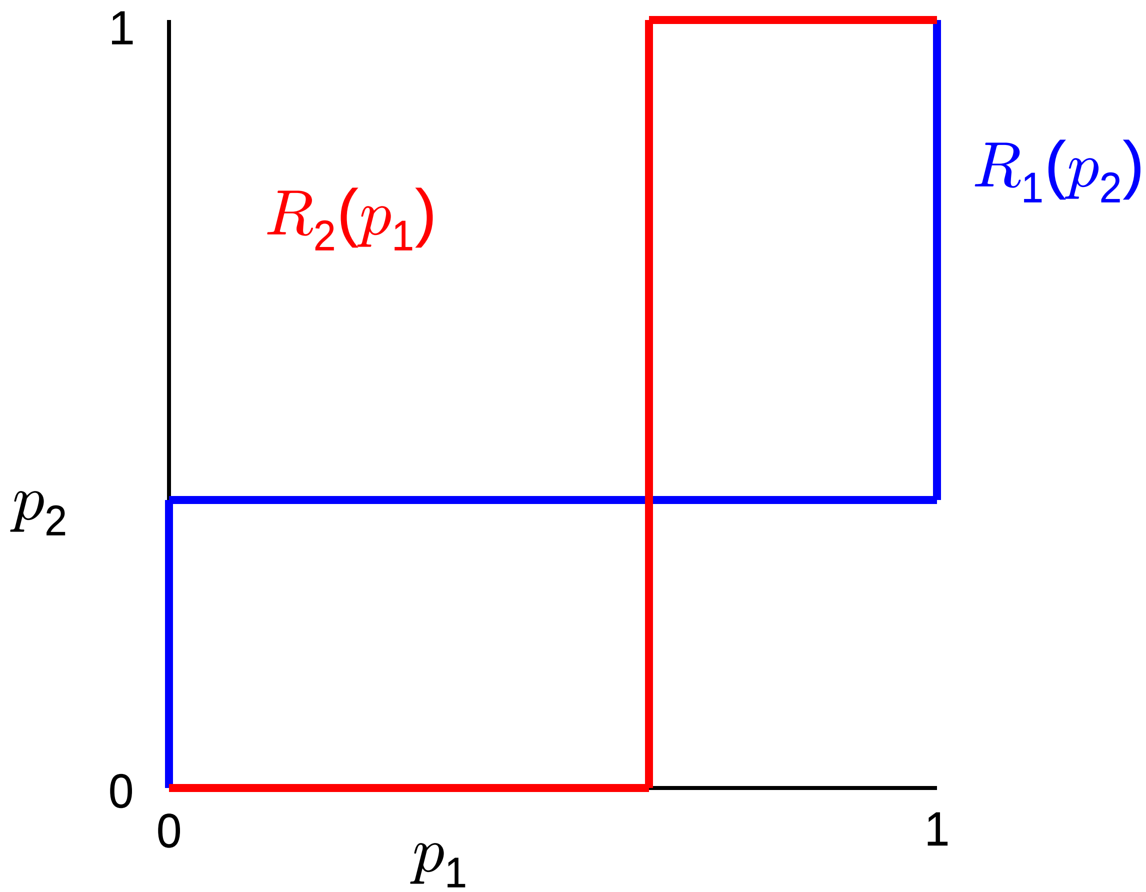

Example: coordination game

- Step 3: Find Nash equilibria.

  As before, NE occur wherever the best response mappings cross.

Example: coordination game

Nash equilibria:

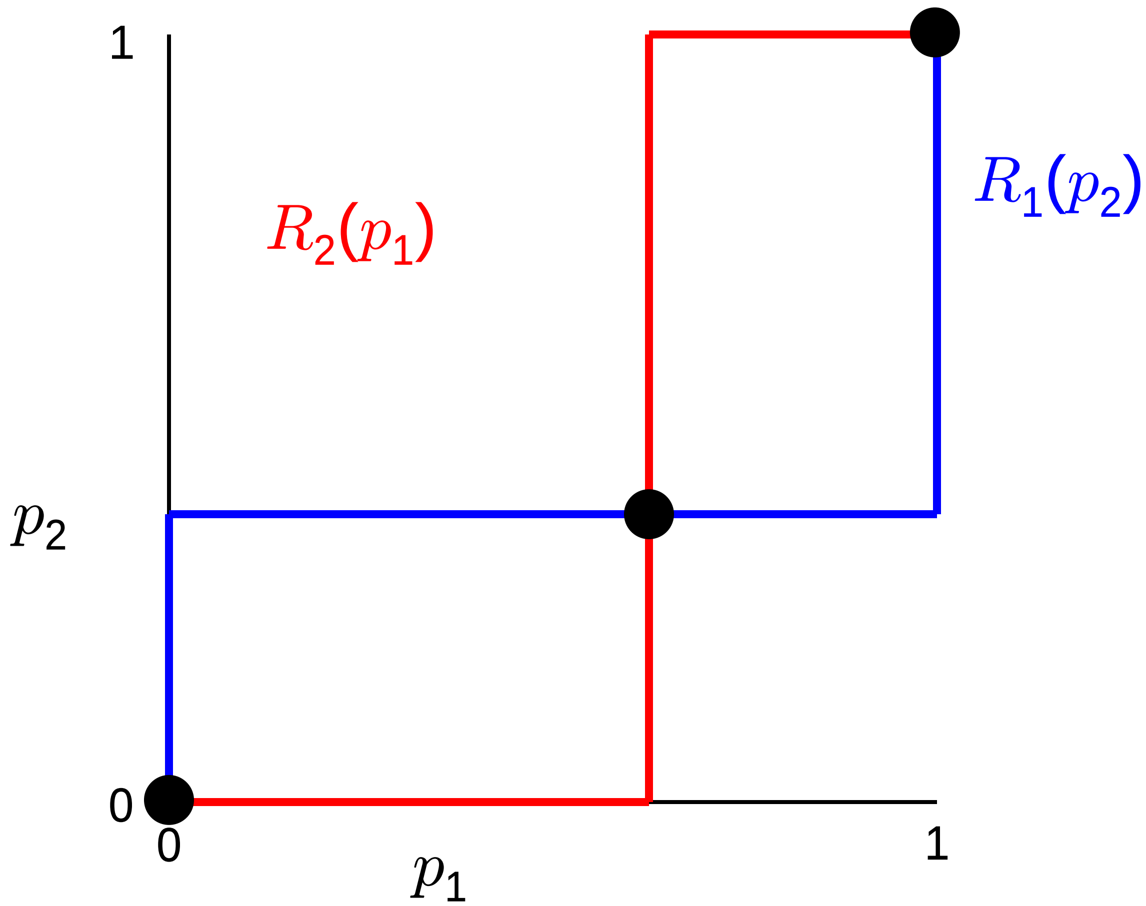

Example: coordination game

Nash equilibria:

There are 3 NE:

$p_1 = 0$, $p_2 = 0 \Rightarrow$ (*r*, *R*)

$p_1 = 1$, $p_2 = 1 \Rightarrow$ (*l*, *L*)

$p_1 = 2/3$, $p_2 = 1/3$

Note: In last NE, both players get expected payoff:

2/3 x 1/3 x 2 + 1/3 x 2/3 x 1 = 2/3.

The existence theorem

Theorem:

Any N-player game where all strategy spaces are *finite* has at least one Nash equilibrium.

Notes:

-The equilibrium may be mixed.

-There is a generalization if strategy spaces are not finite.

The existence theorem: proof

Let $X = \Delta(S_1) \times \cdots \times \Delta(S_N)$ be the product of all mixed strategy spaces.

Define BR : $X \to X$ by:

$\text{BR}_i(\mathbf{p}_1, \ldots, \mathbf{p}_N) = R_i(\mathbf{p}_{-i})$

The existence theorem: proof

Key observations:

- $\Delta(S_i)$ is a closed and bounded subset of $\mathbb{R}^{|S_i|}$
- Thus X is a closed and bounded subset of Euclidean space
- Also, X is *convex*:

  If $\mathbf{p}$, $\mathbf{p}$' are in X, then so is any point on the line segment between them.

The existence theorem: proof

Key observations (continued):

-BR is "continuous"
(*i.e.*, best responses don't change suddenly as we move through X)

(*Formal statement:*

BR has a closed graph, with
convex and nonempty images)

The existence theorem: proof

By *Kakutani's fixed point theorem,* there exists $(\mathbf{p}_1, \ldots, \mathbf{p}_N)$ such that:

$(\mathbf{p}_1, \ldots, \mathbf{p}_N) \in \mathrm{BR}(\mathbf{p}_1, \ldots, \mathbf{p}_N)$

From definition of BR, this implies:

$\mathbf{p}_i \in R_i(\mathbf{p}_{-i})$ for all i

Thus $(\mathbf{p}_1, \ldots, \mathbf{p}_N)$ is a NE.

The existence theorem

Notice that the existence theorem is not *constructive*:

It tells you *nothing* about how players reach a Nash equilibrium, or an easy process to find one.

Finding Nash equilibria in general can be computationally difficult.

Discussion of Nash equilibrium

Nash equilibrium works best when *it is unique:*

In this case, it is the only stable prediction of how rational players would play,

assuming common knowledge of rationality and the structure of the game.

Discussion of Nash equilibrium

How do we make predictions about play when there are multiple Nash equilibria?

1) Unilateral stability

Any Nash equilibrium is *unilaterally stable*:

If a regulator told players to play a given Nash equilibrium,
they have no reason to deviate.

2) Focal equilibria

In some settings, players may have prior preferences that “focus” attention on one equilibrium.

Schelling’s example (see MWG text):

Coordination game to decide where to meet in New York City.

3) Focusing by prior agreement

If players agree ahead of time on a given equilibrium, they have no reason to deviate in practice.

This is a common justification, but can break down easily in practice:
when a game is played only once,
true enforcement is not possible.

4) Long run learning

Another common defense is that if players play the game many (independent) times, they will naturally "converge" to some Nash equilibrium as a stable convention.

Again, this is dangerous reasoning: it ignores a rationality model for dynamic play.

Problems with NE

Nash equilibrium makes very strong assumptions:

-complete information

-rationality

-common knowledge of rationality

-“focusing” (if multiple NE exist)

Example

Find *all* NE (pure and mixed) of the following game:

Player 1 \ Player 2	a	b	c	d
A	(1,2)	(4,0)	(0,3)	(1,1)
B	(0,1)	(2,2)	(1,2)	(0,3)
C	(1,2)	(0,3)	(3,0)	(0,1)
D	(0.5,1)	(0,0)	(0,0)	(2,0)

MS&E 246: Lecture 5
Efficiency and fairness

Ramesh Johari

A digression

In this lecture:

We will use some of the insights of static game analysis to understand *efficiency and fairness.*

Basic setup

- N players
- S_n : strategy space of player n
- Z : space of outcomes
- $z(s_1, ..., s_N)$:
  outcome realized when $(s_1, ..., s_N)$ is played
- $\Pi_n(z)$:
  payoff to player n when outcome is z

(Pareto) Efficiency

An outcome z' *Pareto dominates* z if:

$\Pi_n(z') \geq \Pi_n(z)$ for all n,

and the inequality is strict for at least one n.

An outcome z is *Pareto efficient* if it is not Pareto dominated by any other $z' \in X$.

$\Rightarrow$ Can't make one player better off without making another worse off.

Are equilibria efficient?

Recall the Prisoner's dilemma:

		Player 1	
		defect	cooperate
Player 2	defect	(-4,-4)	(-1,-5)
	cooperate	(-5,-1)	(-2,-2)

Are equilibria efficient?

Recall the Prisoner's dilemma:

		Player 1	
		defect	cooperate
Player 2	defect	(-4,-4)	(-1,-5)
	cooperate	(-5,-1)	(-2,-2)

Unique dominant strategy eq.: (D, D).

Are equilibria efficient?

		Player 1	
		defect	cooperate
Player 2	defect	(-4,-4)	(-1,-5)
	cooperate	(-5,-1)	(-2,-2)

But (C, C) *Pareto dominates* (D, D).

Are equilibria efficient?

- Moral:
  Even when every player has a strict dominant strategy,
  the resulting equilibrium may be *inefficient*.

Resource sharing

- N users want to send data across a shared communication medium
- x_n : sending rate of user n (pkts/sec)
- $p(y)$: probability a packet is lost when total sending rate is y
- $\Pi_n(\mathbf{x})$ = *net throughput* of user n

$$= x_n \left(1 - p\left(\textstyle\sum_i x_i\right)\right)$$

Resource sharing

- Suppose: $p(y)$ = min(y/C, 1)

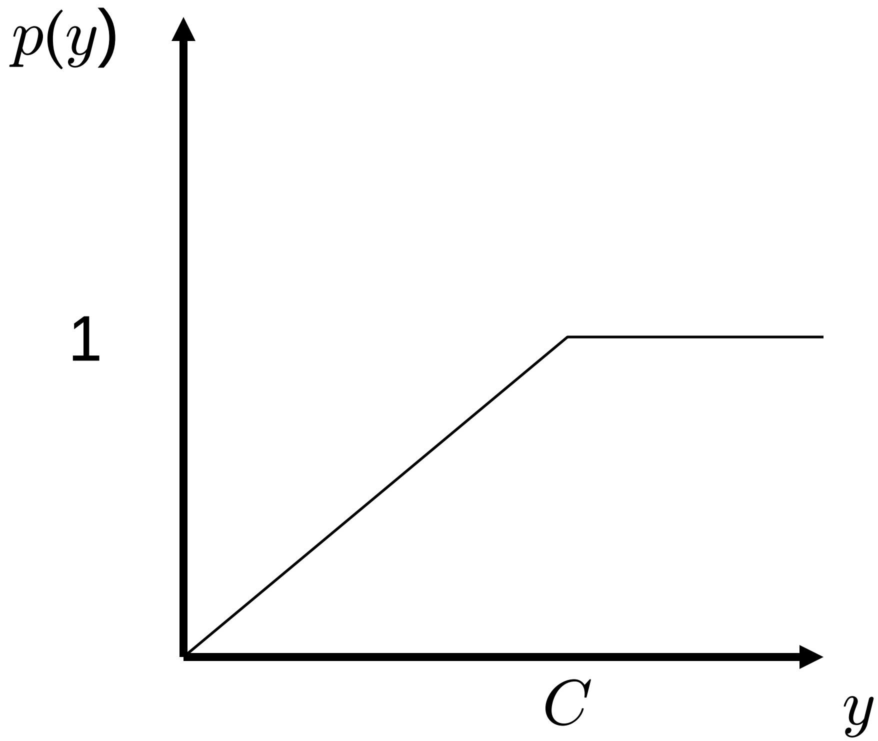

Resource sharing

- Suppose: $p(y)$ = min(y/C, 1)
- Given $\mathbf{x}$:
  Define $Y = \sum_i x_i$ and $Y_{-n} = \sum_{i \neq n} x_i$
- Thus, given $\mathbf{x}_{-n}$,

$$\Pi_n(x_n, \mathbf{x}_{-n}) = x_n \left(1 - \frac{x_n + Y_{-n}}{C}\right)$$

if $x_n + Y_{-n} \leq C$, and zero otherwise

Pure strategy Nash equilibrium

- We only search for NE s.t. $\sum_i x_i \leq C$ (Why?)
- In this region, first order conditions are:

$$1 - Y_{-n}/C - 2x_n/C = 0, \text{ for all } n$$

Pure strategy Nash equilibrium

- We only search for NE s.t. $\sum_i x_i \leq C$ (Why?)
- In this region, first order conditions are:

$$1 - Y/C = x_n / C, \text{ for all } n$$

- If we sum over n and solve for Y, we find:

$$Y^{NE} = N\,C / (N + 1)$$

- So: $x_n^{NE} = C / (N + 1)$, and

$$\Pi_n(\mathbf{x}^{NE}) = C / (N + 1)^2$$

Maximum throughput

- Note that total throughput $= \sum_n \Pi_n(\mathbf{x}) = Y(1 - p(Y)) = Y(1 - Y/C)$
- This is maximized at $Y^{\mathrm{MAX}} = C / 2$
- Define $x_n^{\mathrm{MAX}} = Y^{\mathrm{MAX}} / N = C / 2N$
- Then (if $N > 1$):

$$\Pi_n(\mathbf{x}^{\mathrm{MAX}}) = C / 4N > C / (N+1)^2 = \Pi_n(\mathbf{x}^{\mathrm{NE}})$$

So: $\mathbf{x}^{\mathrm{NE}}$ *is not efficient.*

Resource sharing: summary

- At NE, users' rates are *too high.* Why?
- When user n maximizes Π_n, he ignores reduction in throughput he causes for other players (the *negative externality*)
- AKA: Tragedy of the Commons
- If externality is *positive*,
  then NE strategies are *too low*

An interference model

- N = 2 wireless devices want to send data
- Strategy = transmit power
  $S_1 = S_2 = \{ 0, P \}$
- Each device sees the other's transmission as *interference*

An interference model

Payoff matrix ($0 < \varepsilon << R_2 < R_1$):

		Device 2	
		0	P
Device 1	0	(0, 0)	(0, R_2)
	P	(R_1, 0)	(ε, ε)

An interference model

- (P, P) is unique strict dominant strategy equilibrium (and hence unique NE)
- Note that (P, P) is not Pareto dominated by any pure strategy pair
- But…the mixed strategy pair $(\mathbf{p}_1, \mathbf{p}_2)$ with $p_1(0) = p_2(0) = p_1(P) = p_2(P) = 1/2$ Pareto dominates (P, P) if $R_n >> \varepsilon$

  (Payoffs: $\Pi_n(\mathbf{p}_1, \mathbf{p}_2) = R_n/4 + \varepsilon/4$)

An interference model

- How can coordination improve throughput?
- *Idea:*
  Suppose both devices agree to a protocol that decides when each device is allowed to transmit.

An interference model

- *Cooperative timesharing:*

  Device 1 is allowed to transmit
  a fraction q of the time.

  Device 2 is allowed to transmit
  a fraction 1 - q of the time.

  Devices can use any mixed strategy
  when they control the channel.

An interference model

Achievable payoffs via timesharing:

i.e.: over all q and all strategy pairs (σ_1, σ_2)

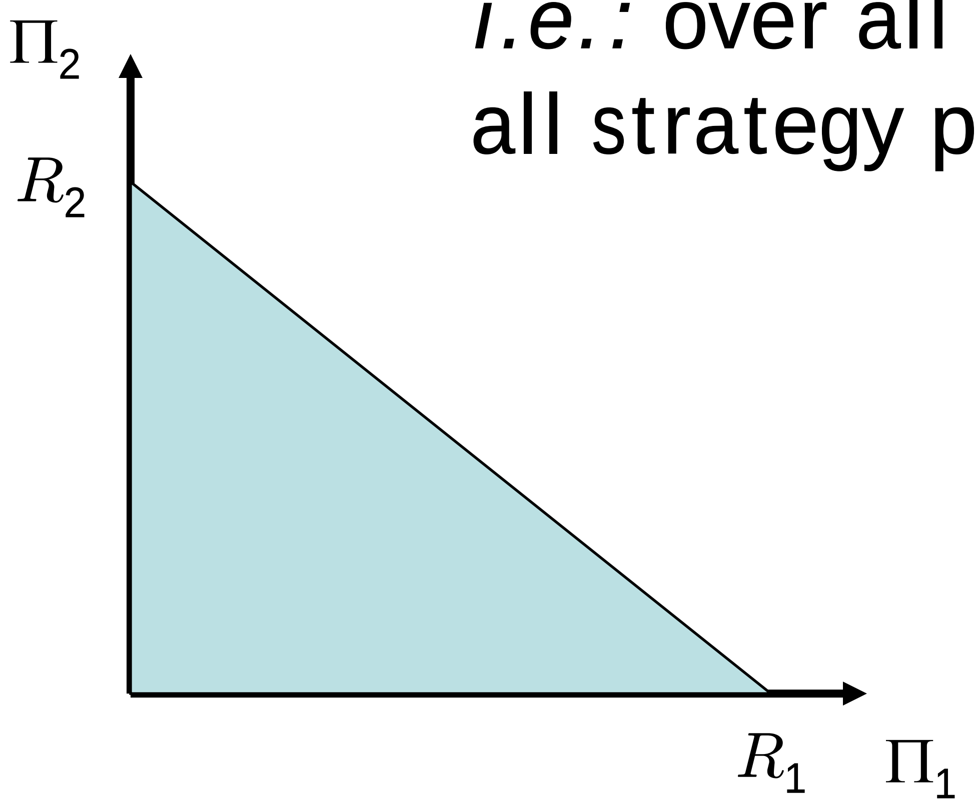

An interference model

Achievable payoffs via timesharing:

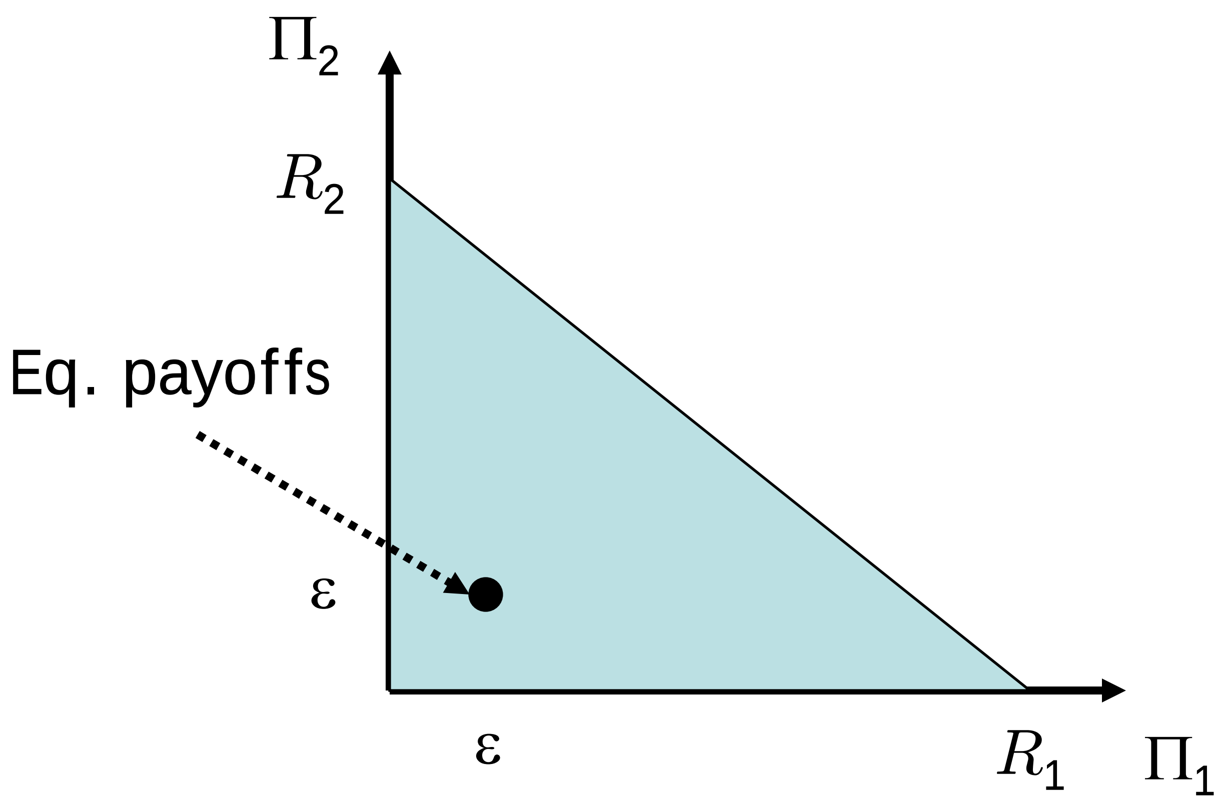

An interference model

- So when timesharing is used, the set of Pareto efficient payoffs becomes:
  $\{ (\Pi_1, \Pi_2) : \Pi_1 = q\, R_1, \Pi_2 = (1 - q)\, R_2 \}$

- For efficiency:
  When device n has control,
  it transmits at power P

An interference model

- So when timesharing is used, the set of Pareto efficient payoffs becomes:
  $\{ (\Pi_1, \Pi_2) : \Pi_1 = q\, R_1, \Pi_2 = (1 - q)\, R_2 \}$

- For efficiency:
  When device n has control, it transmits at power P

An interference model

- So when timesharing is used, the set of Pareto efficient payoffs becomes:
  $\{ (\Pi_1, \Pi_2) : \Pi_1 = q R_1, \Pi_2 = (1 - q) R_2 \}$

  (*Note:* in general, the set of achievable payoffs is the *convex hull* of entries in the payoff matrix)

Choosing an efficient point

Which q should the protocol choose?

- Choice 1: *Utilitarian solution*

  $\Rightarrow$ Maximize total throughput

  $\max_q q R_1 + (1- q) R_2 \quad \Rightarrow \quad q = 1$

  $\Pi_1 = R_1, \Pi_2 = 0$

  Is this "fair"?

Choosing an efficient point

Which q should the protocol choose?

- Choice 2: *Max-min fair solution*

  $\Rightarrow$ Maximize smallest Π_n

  $\max_q \min \{ q R_1, (1 - q) R_2 \}$

  $\Rightarrow q R_1 = (1 - q) R_2$,

  so $\Pi_1 = \Pi_2$ (i.e., equalize rates)

Fairness

Fairness corresponds to a rule for choosing between multiple efficient outcomes.

Unlike efficiency, there is no universally accepted definition of "fair."

Nash bargaining solution (NBS)

- Fix desirable properties of a “fair” outcome

- Show there exists a unique outcome satisfying those properties

NBS: Framework

- T = { (Π_1, Π_2) : (Π_1, Π_2) is achievable }
  - assumed closed, bounded, and convex

- $\Pi^* = (\Pi_1^*, \Pi_2^*)$: *status quo* point
  - each n can guarantee Π_n^* for himself through unilateral action

NBS: Framework

- $f(T, \Pi^*) = (f_1(T, \Pi^*),\ f_2(T, \Pi^*)) \in T$:
  a *"bargaining solution"*,
  i.e., a rule for choosing a payoff pair

- What properties (*axioms*) should f satisfy?

Axioms

Axiom 1: Pareto efficiency

The payoff pair $f(T, \Pi^*)$ must be Pareto efficient in T.

Axiom 2: Individual rationality

For all n, $f_n(T, \Pi^*) \geq \Pi_n^*$.

Axioms

Given $\mathbf{v} = (v_1, v_2)$, let

$T + \mathbf{v} = \{ (\Pi_1 + v_1, \ \Pi_2 + v_2) : (\Pi_1, \Pi_2) \in T \}$

(i.e., a change of origin)

Axiom 3: Independence of utility origins

Given any $\mathbf{v} = (v_1, v_2)$,

$$f(T + \mathbf{v}, \Pi^* + \mathbf{v}) = f(T, \Pi^*) + \mathbf{v}$$

Axioms

Given $\beta = (\beta_1, \beta_2)$, let

$\beta \cdot T = \{ (\beta_1 \Pi_1, \beta_2 \Pi_2) : (\Pi_1, \Pi_2) \in T \}$

(i.e., a change of utility units)

Axiom 4: Independence of utility units

Given any $\beta = (\beta_1, \beta_2)$, for each n we have

$$f_n(\beta \cdot T, (\beta_1 \Pi_1^*, \beta_2 \Pi_2^*)) = \beta_n f_n(T, \Pi^*)$$

Axioms

The set T is *symmetric* if it looks the same when the Π_1-Π_2 axes are swapped:

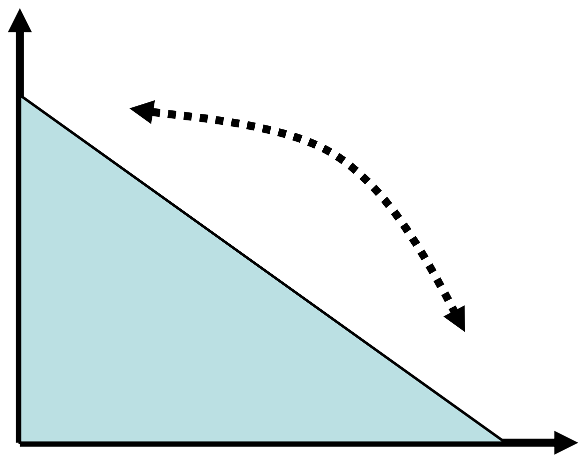

Axioms

The set T is *symmetric* if it looks the same when the Π_1-Π_2 axes are swapped:

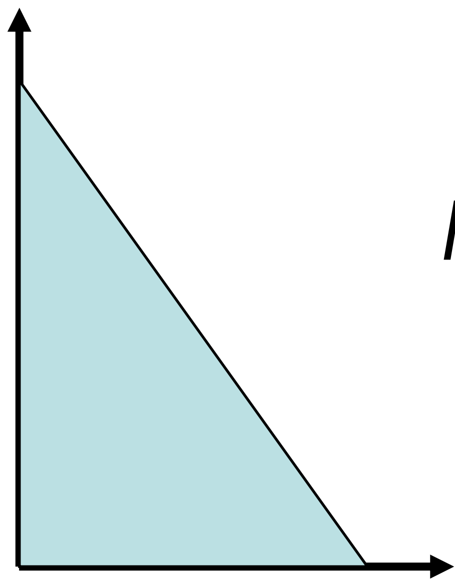

Not symmetric

Axioms

The set T is *symmetric* if it looks the same when the Π_1-Π_2 axes are swapped:

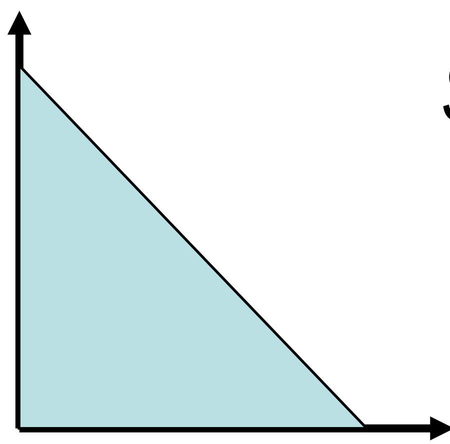

Symmetric

Axioms

The set T is *symmetric* if it looks the same when the Π_1-Π_2 axes are swapped.

Axiom 5: Symmetry

If T is symmetric and $\Pi_1^* = \Pi_2^*$, then $f_1(T, \Pi^*) = f_2(T, \Pi^*)$.

Axioms

Axiom 6: Independence of irrelevant alternatives

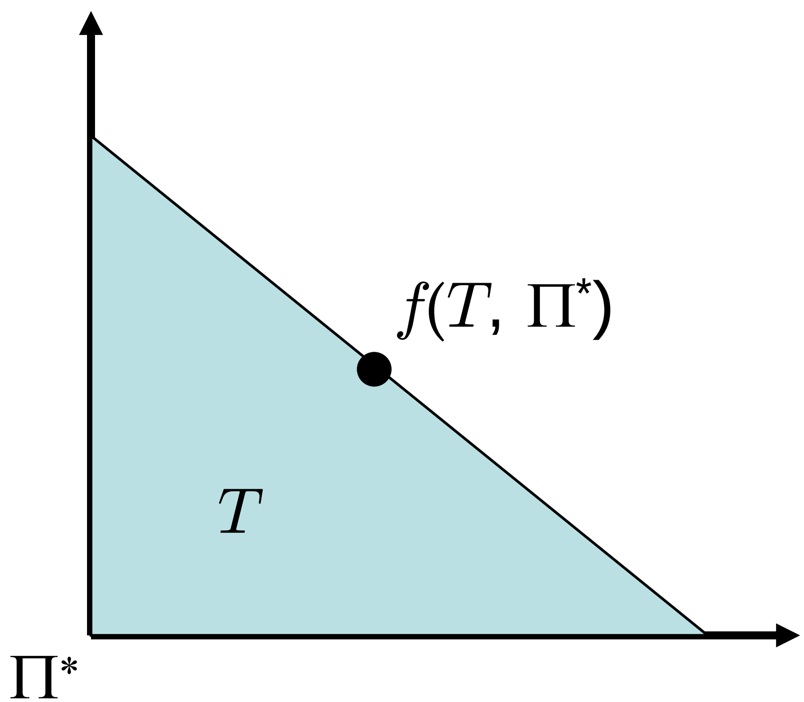

Axioms

Axiom 6: Independence of irrelevant alternatives

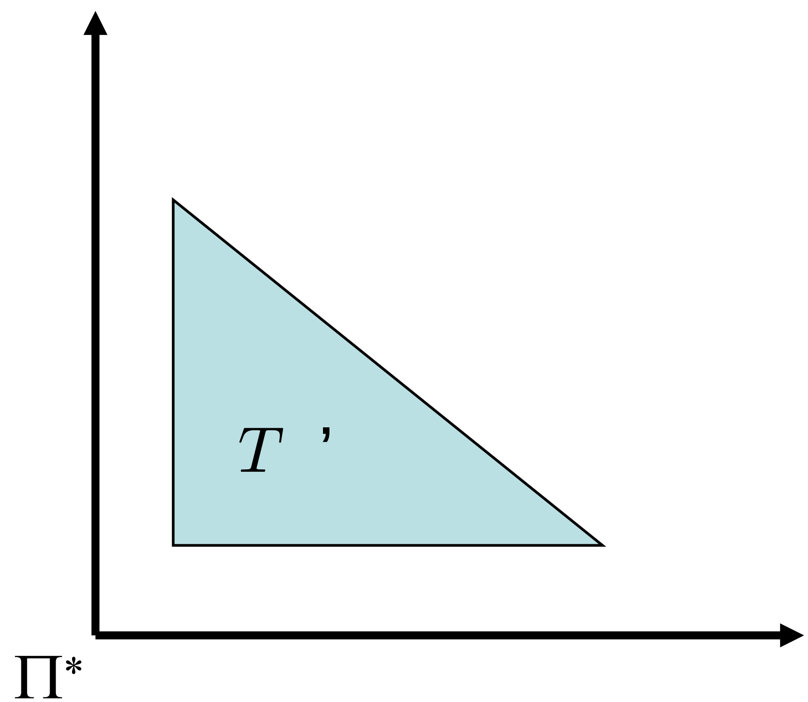

Axioms

Axiom 6: Independence of irrelevant alternatives

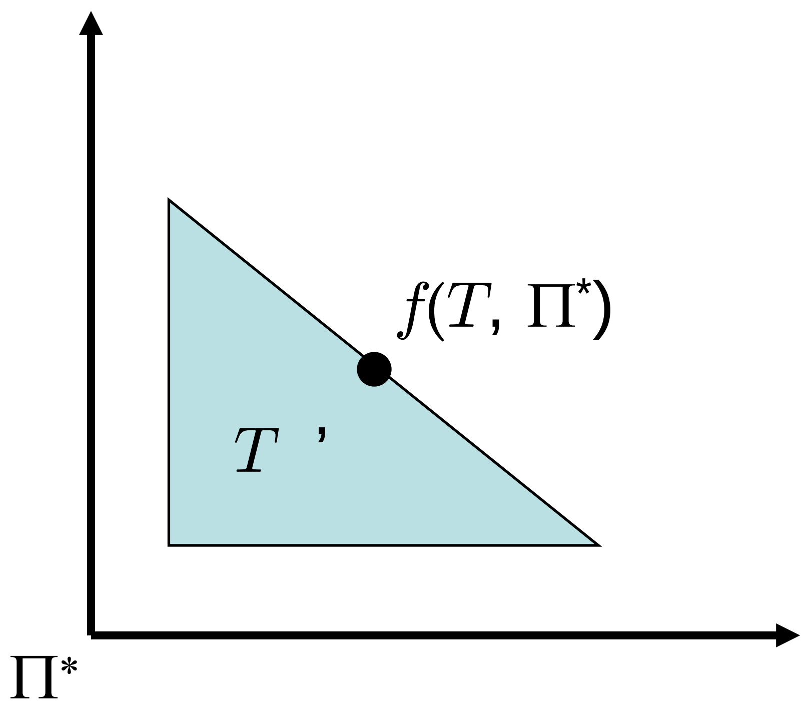

Axioms

Axiom 6: Independence of irrelevant alternatives

If $T' \subset T$ and $f(T, \Pi^*) \in T'$, then $f(T, \Pi^*) = f(T', \Pi^*)$.

Nash bargaining solution

Theorem (Nash):

There exists a unique f satisfying Axioms 1-6, and it is given by:

$f(T, \Pi^*)$

$= \arg\max_{\Pi \in T : \Pi \geq \Pi^*} (\Pi_1 - \Pi_1^*)(\Pi_2 - \Pi_2^*)$

$= \arg\max_{\Pi \in T : \Pi \geq \Pi^*} \sum_{n = 1,2} \log (\Pi_n - \Pi_n^*)$

(Sometimes called proportional fairness.)

Nash bargaining solution

- The proof relies on *all* the axioms
- The utilitarian solution and the max-min fair solution do not satisfy independence of utility units
- See course website for excerpt from MWG

Back to the interference model

- $T = \{ (\Pi_1, \Pi_2) \geq 0 :$
  $\Pi_1 \leq q\,R_1, \Pi_2 \leq (1 - q)\,R_2\ , 0 \leq q \leq 1\}$
- $\Pi^* = (\varepsilon, \varepsilon)$
- NBS: $\max_q \log(q\,R_1 - \varepsilon) + \log((1 - q)\,R_2 - \varepsilon)$

Solution: $q = 1/2 + (\varepsilon/2)(1/R_1 - 1/R_2)$

e.g., when $\varepsilon = 0$,

$\Pi_1^{NBS} = R_1/2,\ \ \Pi_2^{NBS} = R_2/2$

Comparisons

Assume $\varepsilon = 0$, $R_1 > R_2$

	q	Π_1	Π_2
Utilitarian	1	R_1	0
NBS	1/2	$R_1/2$	$R_2/2$
Max-min fair	$\frac{R_2}{R_1 + R_2}$	$\frac{R_1 R_2}{R_1 + R_2}$	$\frac{R_1 R_2}{R_1 + R_2}$

Summary

- When we say "efficient", we mean *Pareto efficient*.
- When we say "fair", we must make clear what we mean!
- Typically, Nash equilibria are not efficient
- The Nash bargaining solution is one axiomatic approach to fairness

MS&E 246: Lecture 6
Dynamic games of perfect and complete information

Ramesh Johari

Outline

- Dynamic games
- Perfect information
- Game trees
- Strategies
- Backward induction

Dynamic games

Instead of playing simultaneously,

the rules dictate *when* players play,

and *what they know about the past* when they play.

Example

Consider the following game:

		Player 2	
		L	*R*
Player 1	*l*	(4,1)	(1,2)
	r	(2,1)	(0,0)

Only pure NE is (*l*, *R*).

Example: dynamic game

Suppose player 1 moves first, and player 2 moves second:

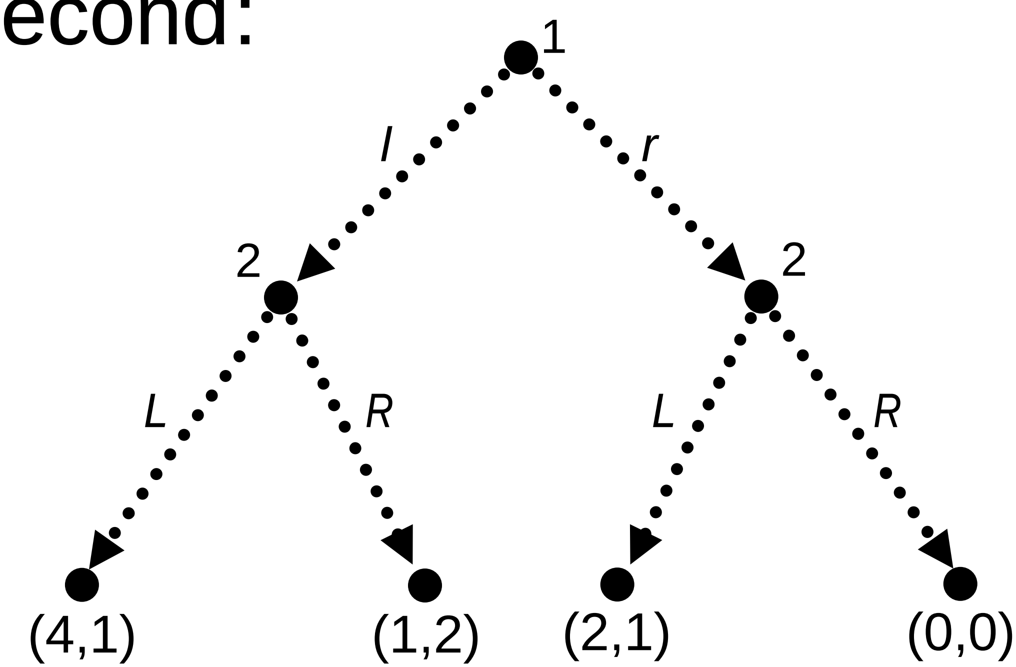

Example: dynamic game

If player 1 plays *l* ...

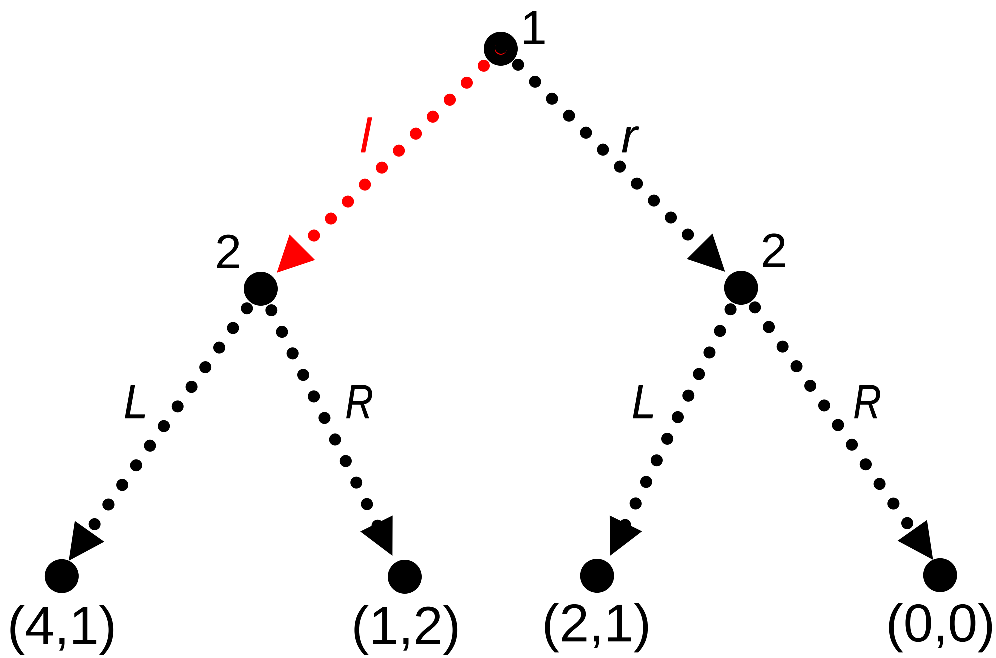

Example: dynamic game

...then player 2 plays *R* $\Rightarrow \Pi_1 = 1$.

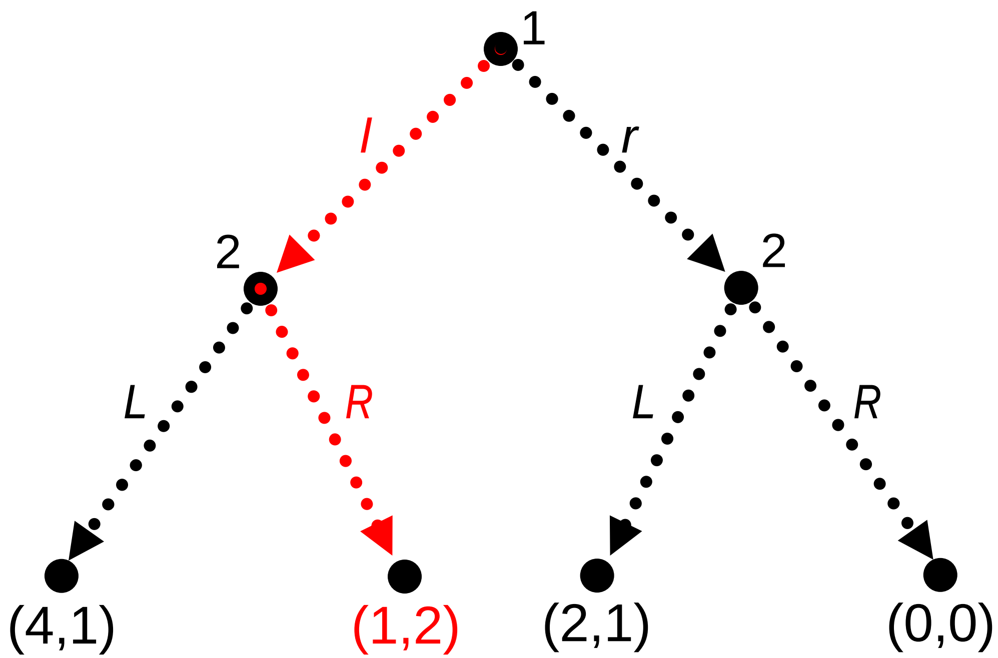

Example: dynamic game

If player 1 plays *r* ...

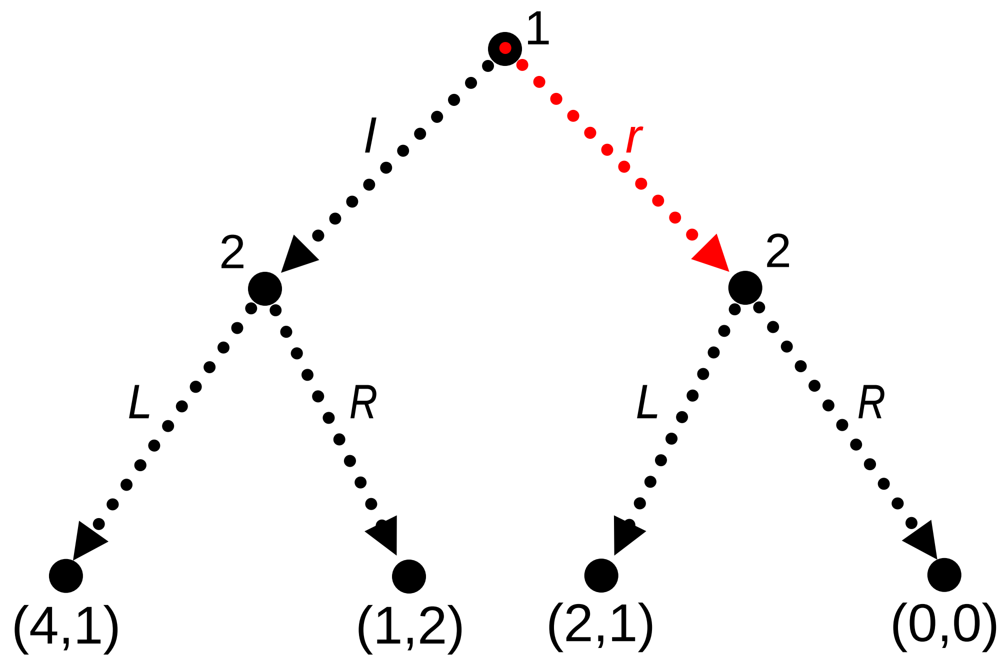

Example: dynamic game

...then player 2 plays *L* $\Rightarrow \Pi_1 = 2$!

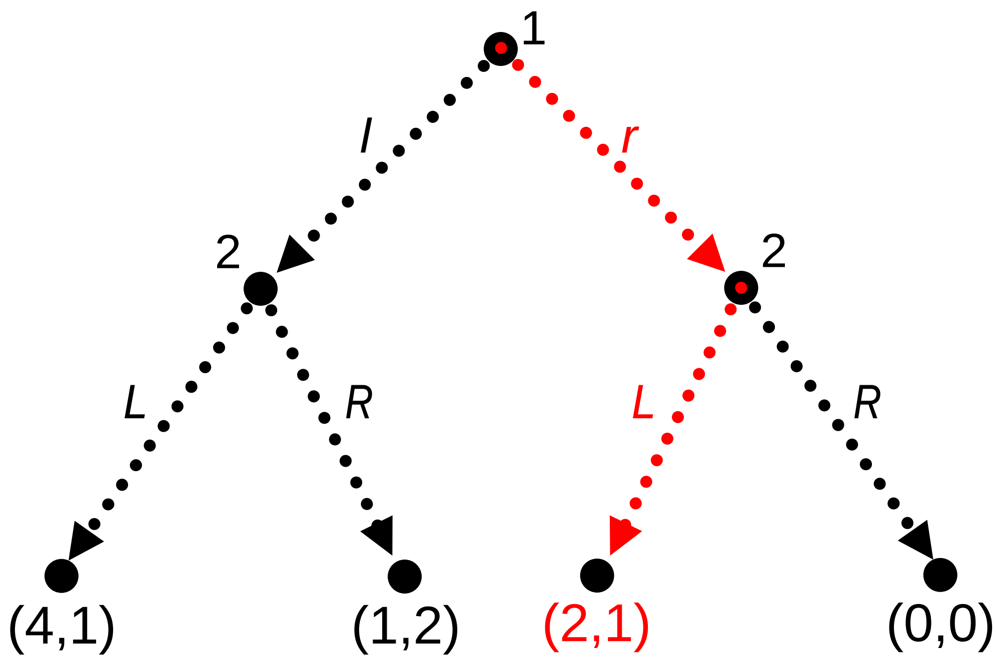

Example: dynamic game

What if player 2 does not observe player 1's decision?

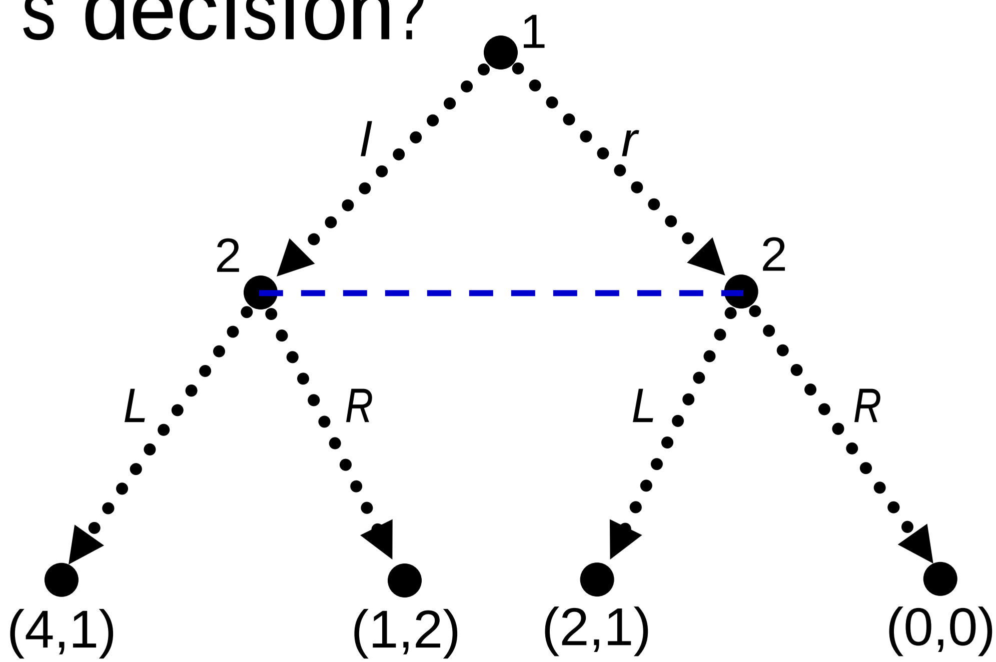

- - - -: Player 2 cannot distinguish between these nodes

Dynamics

What if player 2 does not observe player 1's decision?

Harder to predict what player 2 might do at stage 2.

Perfect information

In this lecture we will study (finite) dynamic games of *perfect information:*

These are games where all players observe the entire *history* of the game, and the game terminates in finitely many steps.

(NOTE: This is not complete information!)

Game tree

Fundamental structure: the *game tree*.

1) Each non-leaf node v is identified with a unique player $I(v)$.
2) All edges out from a node v correspond to *actions* available to $I(v)$.
3) All leaves are labeled with the *payoffs* for all players.

Game tree

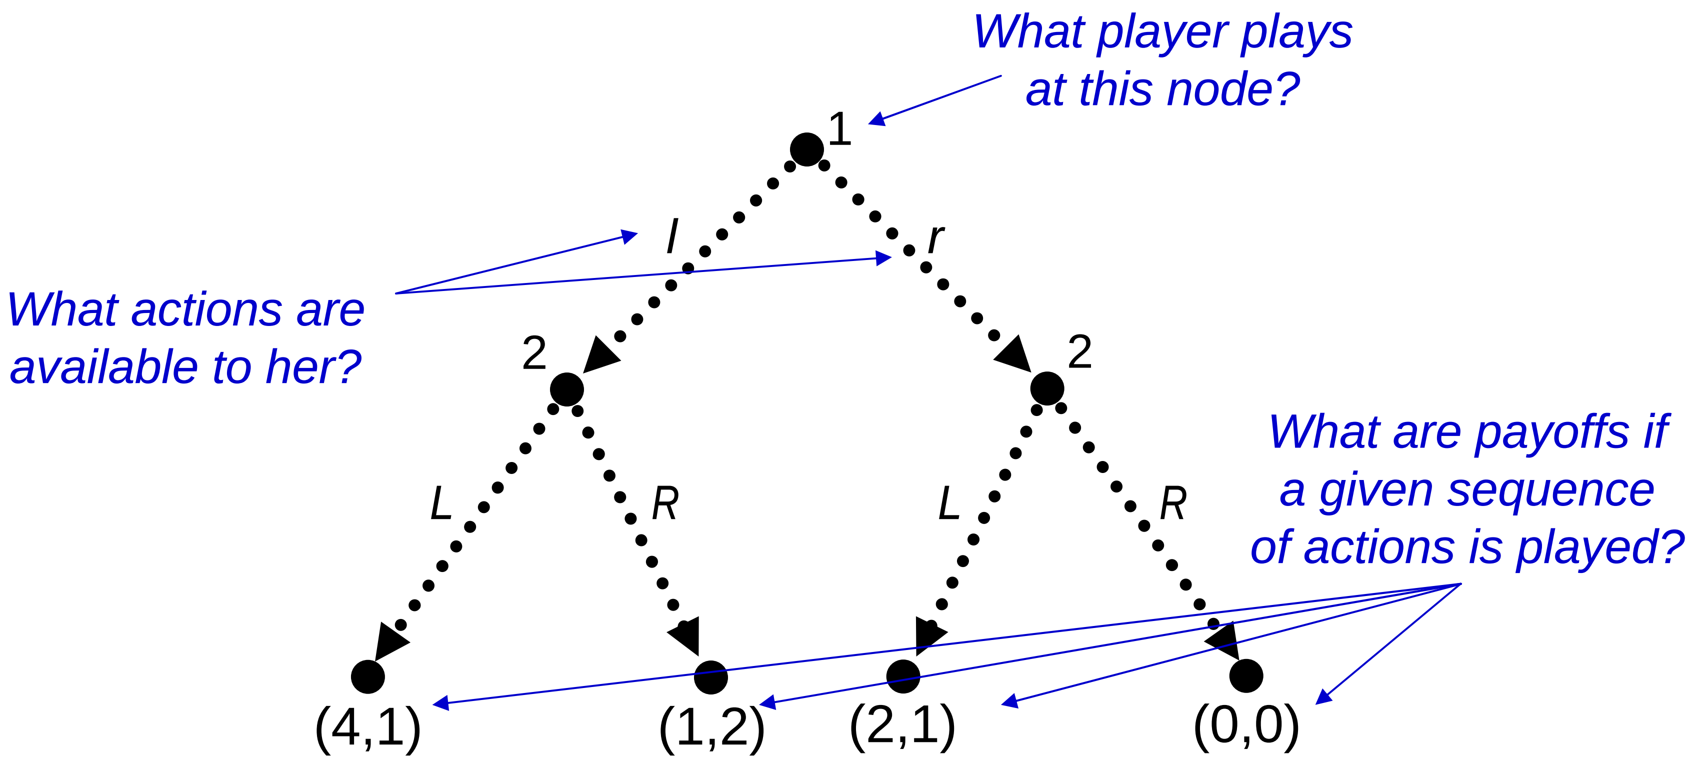
What player plays
at this node?
1
l
r
What actions are
available to her?
2
2
L
R
L
R
What are payoffs if
a given sequence
of actions is played?
(4,1)
(1,2)
(2,1)
(0,0)

Game trees and extensive form

Idea: At each node v, player $I(v)$ chooses an action; this leads to the next "stage."

The game tree is also called the *extensive form* of the game.

Strategies

For player 1 to “reason” about player 2, there must be a prediction of what player 2 would play *in any of his nodes.*

(See example at the beginning of lecture.)

Strategies

For player 1 to "reason" about player 2, there must be a prediction of what player 2 would play *in any of his nodes.*

Thus in a dynamic game, a strategy s_i is a *complete contingent plan*:

For each v such that $I(v) = i$, $s_i(v)$ specifies the *action* of player i at node v.

Backward induction

We solve finite games of perfect information using *backward induction.*

Idea: find "optimal" decisions for players from the bottom of the tree to the top.

Backward induction

Formally: Suppose the game has L stages.

- Find the set of optimal actions for player $I(v)$ at each node v in stage L (possibly including mixed actions).
- Label each node v in stage L with payoffs from optimal actions, and remove any children.
- Return to (1) and repeat.

Backward induction

- Note this specifies *complete strategies* for all players,
  as well as the *path(s) of actual play.*
- Any finite dynamic game of perfect information has a backward induction solution.

Example

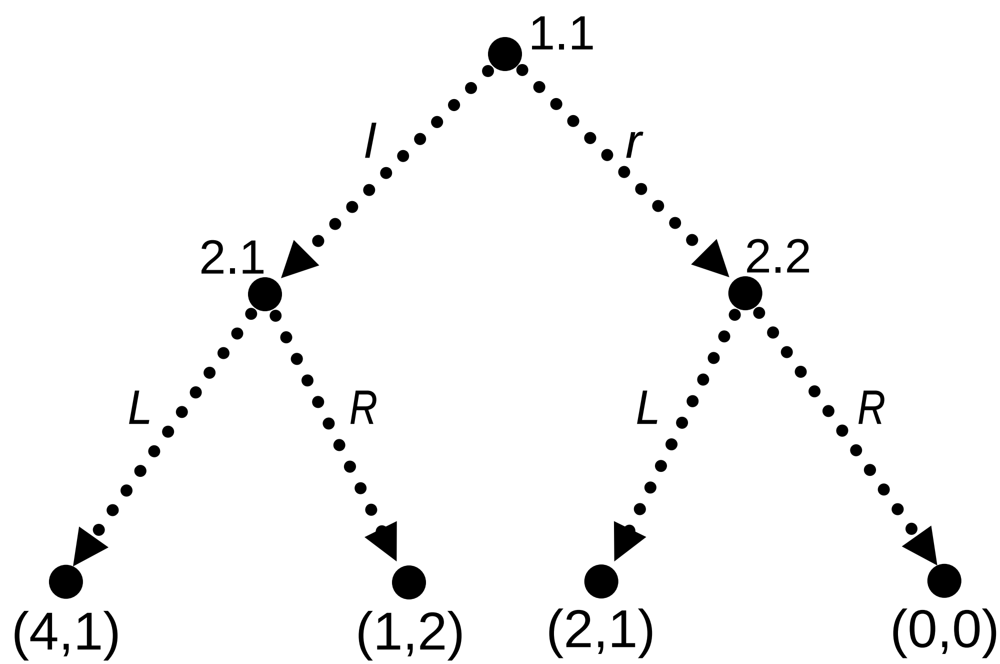

Example

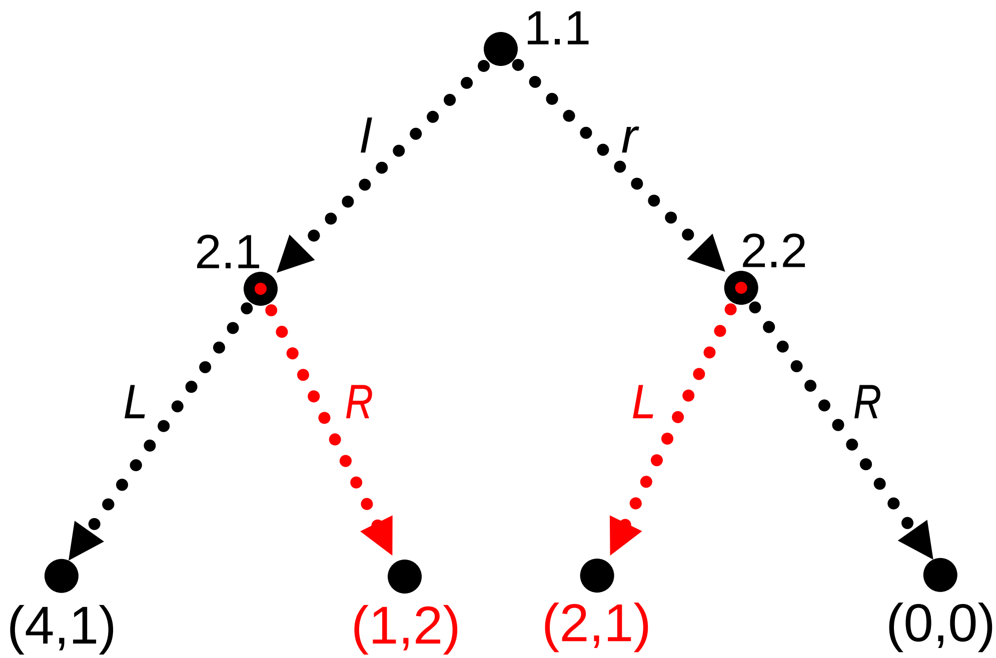

Example

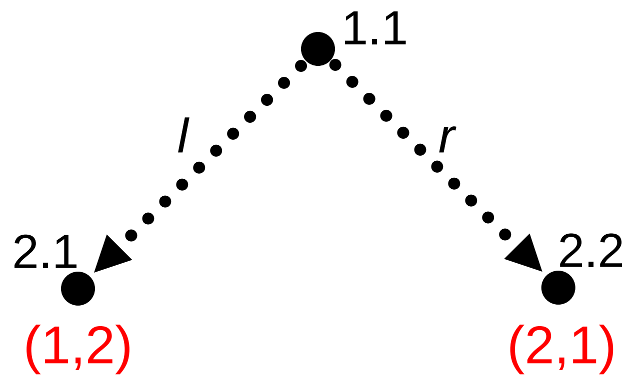

Example

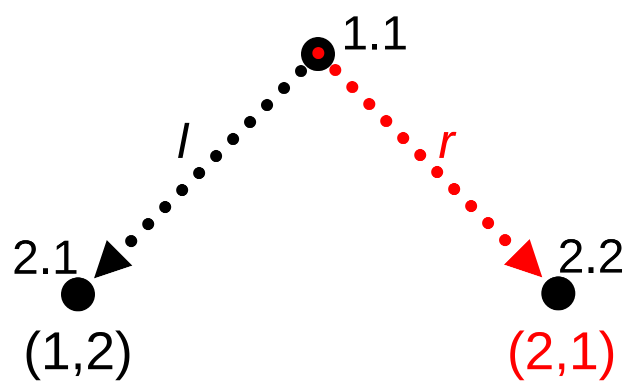

Example

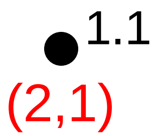

Backward induction solution

Strategies:

Player 1 plays *r* at node 1.1.

Player 2 plays *R* at node 2.1, and plays *L* at node 2.2.

The *equilibrium path of play* is (*r*, *L*).

Strategic form

The strategic (or normal) form is:

Player 2

Player 1		*LL*	*LR*	*RL*	*RR*
	l	(4,1)	(4,1)	(1,2)	(1,2)
	r	(2,1)	(0,0)	(2,1)	(0,0)

Strategic form

What are all pure NE of the strategic form?

(*r*, *RL*) and (*l*, *RR*)

But (*l*, *RR*) is *not credible:*

Player 1 knows a rational player 2 would never play *R* in 2.2.

Strategic form

Any backward induction solution must be a NE of the strategic form,

but the converse does not necessarily hold.

MS&E 246: Lecture 7
Stackelberg games

Ramesh Johari

Stackelberg games

In a Stackelberg game, one player (the “leader”) moves first,

and all other players (the “followers”) move after him.

Stackelberg competition

- Two firms (N = 2)
- Each firm chooses a quantity $s_n \geq 0$
- Cost of producing s_n : $c_n\ s_n$
- *Demand curve*:

  Price = $P(s_1 + s_2) = a - b\ (s_1 + s_2)$
- Payoffs:

  Profit = $\Pi_n(s_1,\ s_2) = P(s_1 + s_2)\ s_n - c_n\ s_n$

Stackelberg competition

In *Stackelberg competition*, firm 1 moves before firm 2.

Firm 2 observes firm 1's quantity choice s_1, then chooses s_2.

Stackelberg competition

We solve the game using *backward induction*.

Start with second stage:

Given s_1, firm 2 chooses s_2 as

$$s_2 = \arg\max_{s_2 \in S_2} \Pi_2(s_1, s_2)$$

But this is just the *best response* $R_2(s_1)$!

Best response for firm 2

Recall the best response given s_1:

$$\max_{s_2 \geq 0} [(a - bs_2 - bs_1)s_2 - c_2 s_2] \implies$$

Differentiate and solve:

$$a - c_2 - bs_1 - 2bs_2 = 0$$

So:

$$R_2(s_1) = \left[\frac{a - c_2}{2b} - \frac{s_1}{2}\right]^+$$

Firm 1's decision

Backward induction:

Maximize firm 1's decision, *accounting for firm 2's response at stage 2.*

Thus firm 1 chooses s_1 as

$$s_1 = \arg\max_{s_1 \in S_1} \Pi_1(s_1, R_2(s_1))$$

Firm 1's decision

Define $t_n = (a - c_n)/b$.

If $s_1 \leq t_2$, then payoff to firm 1 is:

$$\Pi_1 = \left(a - bs_1 - b\left(\frac{t_2}{2} - \frac{s_1}{2}\right)\right)s_1 - c_1 s_1$$

If $s_1 > t_2$, then payoff to firm 1 is:

$$\Pi_1 = (a - bs_1)\,s_1 - c_1 s_1$$

Firm 1's decision

For simplicity, we assume that

$$2c_2 \leq a + c_1$$

This assumption ensures that

$$(a - bs_1)\,s_1 - c_1 s_1$$

is *strictly decreasing* for $s_1 > t_2$.

Thus firm 1's optimal s_1 must lie in [0, t_2].

Firm 1's decision

If $s_1 \leq t_2$, then payoff to firm 1 is:

$$\Pi_1 = \left(a - bs_1 - b\left(\frac{t_2}{2} - \frac{s_1}{2}\right)\right)s_1 - c_1 s_1$$

Firm 1's decision

If $s_1 \leq t_2$, then payoff to firm 1 is:

$$\Pi_1 = \left(\frac{a}{2} - \frac{b}{2}s_1 + \frac{c_2}{2}\right)s_1 - c_1 s_1$$

Firm 1's decision

If $s_1 \leq t_2$, then payoff to firm 1 is:

$$\Pi_1 = -\frac{b}{2}s_1^2 + \left(\frac{a}{2} + \frac{c_2}{2} - c_1\right) s_1$$

Thus optimal s_1 is:

$$s_1 = \frac{a - 2c_1 + c_2}{2b}$$

Stackelberg equilibrium

So what is the Stackelberg equilibrium?

Must give complete strategies:

$s_1^* = (a - 2c_1 + c_2)/2b$

$s_2^*(s_1) = (t_2/2 - s_1/2)^+$

The *equilibrium outcome* is that
firm 1 plays s_1^*, and firm 2 plays $s_2^*(s_1^*)$.

Comparison to Cournot

Assume $c_1 = c_2 = c$.

In Cournot equilibrium:

(1) $s_1 = s_2 = t/3$.

(2) $\Pi_1 = \Pi_2 = (a - c)^2/(9b)$.

In Stackelberg equilibrium:

(1) $s_1 = t/2$, $s_2 = t/4$.

(2) $\Pi_1 = (a - c)^2/(8b)$, $\Pi_2 = (a - c)^2/(16b)$

Comparison to Cournot

So in Stackelberg competition:

-the *leader* has *higher* profits

-the *follower* has *lower* profits

This is called a *first mover advantage.*

Stackelberg competition: moral

Moral:

Additional information available can lower a player's payoff, if it is common knowledge that the player will have the additional information.

(*Here:* firm 1 takes advantage of knowing firm 2 knows s_1.)

MS&E 246: Lecture 8
Dynamic games of complete and imperfect information

Ramesh Johari

Outline

- Imperfect information
- Information sets
- Perfect recall
- Moves by nature
- Strategies
- Subgames and subgame perfection

Imperfect information

Informally:

In *perfect information* games, the history is common knowledge.

In *imperfect information* games, players move without necessarily knowing the past.

Game trees

We adhere to the same model of a *game tree* as in Lecture 6.

1) Each non-leaf node v is identified with a unique player $I(v)$.
2) All edges out from a node v correspond to *actions* available to $I(v)$.
3) All leaves are labeled with the *payoffs* for all players.

Information sets

We represent imperfect information by *combining* nodes into *information sets.*

An information set h is:

-a subset of nodes of the game tree

-all identified with the same player $I(h)$

-with the same actions available to $I(h)$ at each node in h

Information sets

Idea:

When player $I(h)$ is in information set h, *she cannot distinguish between the nodes of* h.

Information sets

Let H denote all information sets.

The union of all sets $h \in H$ gives all nodes in the tree.

(We use $h(v)$ to denote the information set corresponding to a node v.)

Perfect information

Formal definition of perfect information:

All information sets are singletons.

Example

Coordination game without observation:

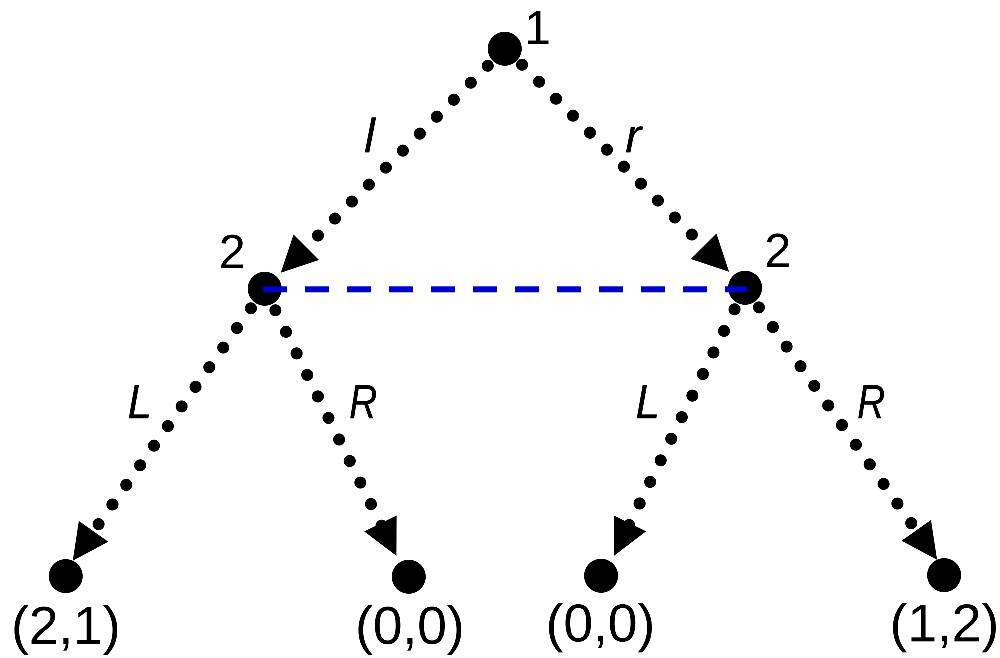

- - - - -: *Player 2 cannot distinguish between these nodes*

Perfect recall

We restrict attention to games of *perfect recall*:

These are games where the information sets ensure a player never forgets what she once knew, or what she played.

[Formally:

If v, v' are in the same information set h, neither is a predecessor of the other in the game tree.

Also, if v', v'' are in the same information set, and v is a predecessor of v', then there must exist a node $w \in h(v)$ that is a predecessor of v'', such that the action taken on the path from v to v' is the same as the action taken on the path from w to v''.]

Moves by nature

We allow one additional possibility:

At some nodes, *Nature* moves.

At such a node v, the edges are labeled with *probabilities* of being selected.

Any such node v models an *exogenous* event.

[*Note:* Players use *expected payoffs.*]

Strategies

The *strategy* of a player is a function from *information sets* of that player to an *action* in each information set.

(In perfect information games, strategies are mappings from nodes to actions.)

Example

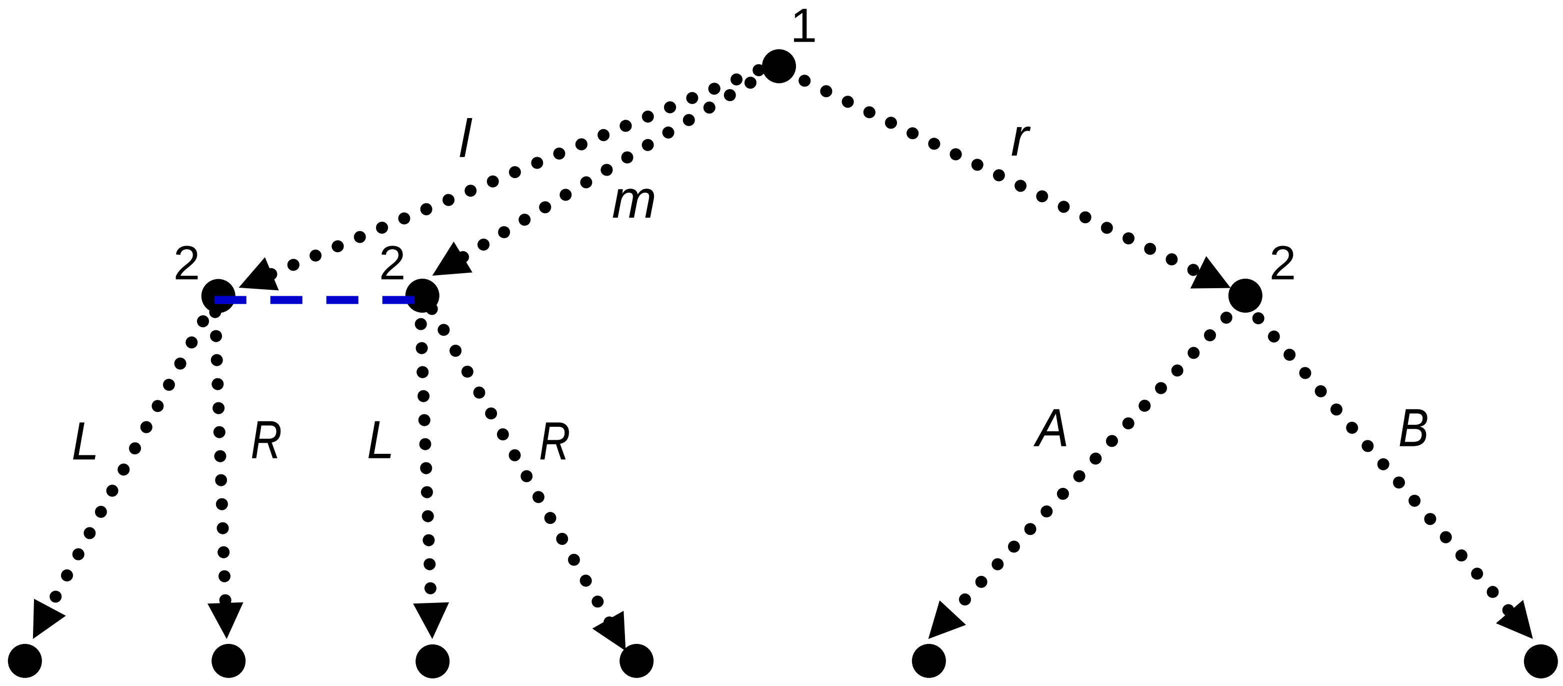

Example

Player 1: 1 information set
3 strategies – *l*, *m*, *r*

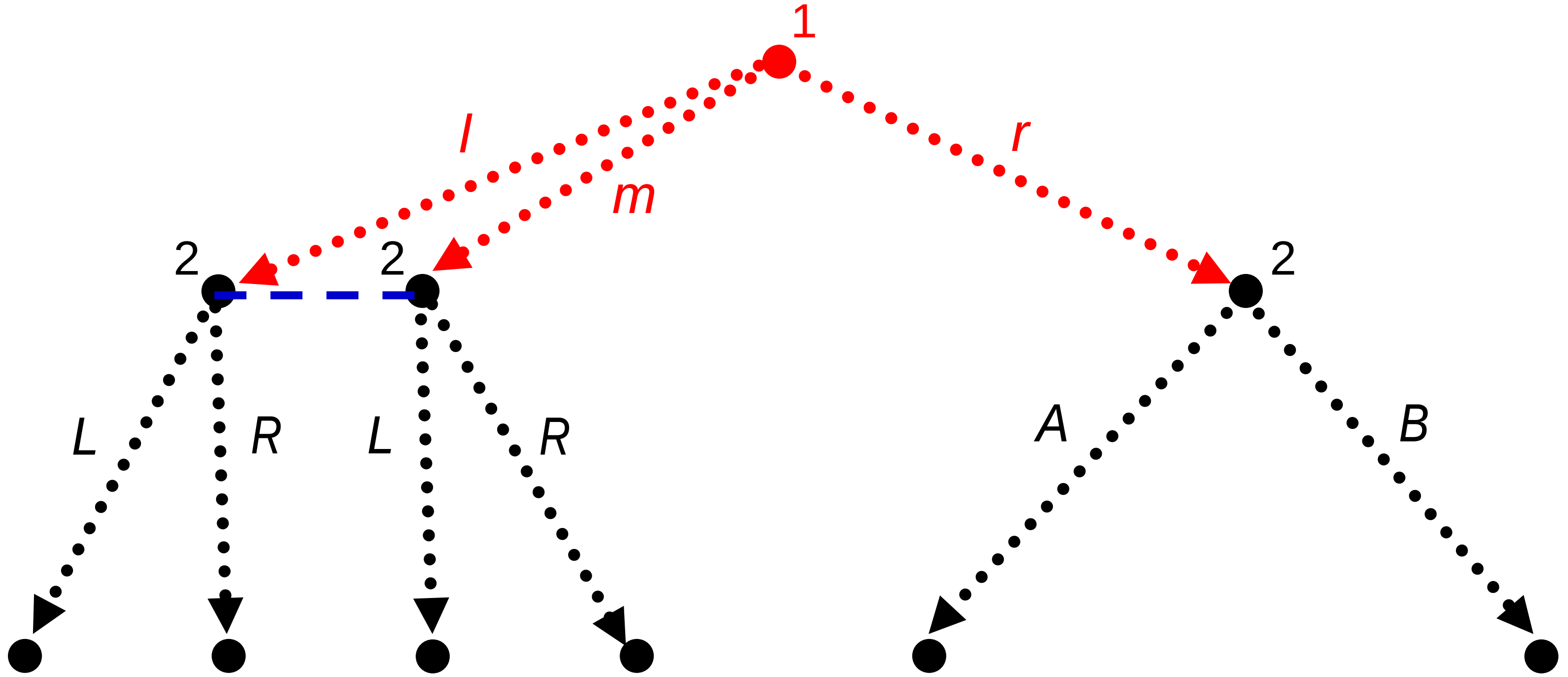

Example

Player 2: 2 information sets
4 strategies – *LA*, *LB*, *RA*, *RB*

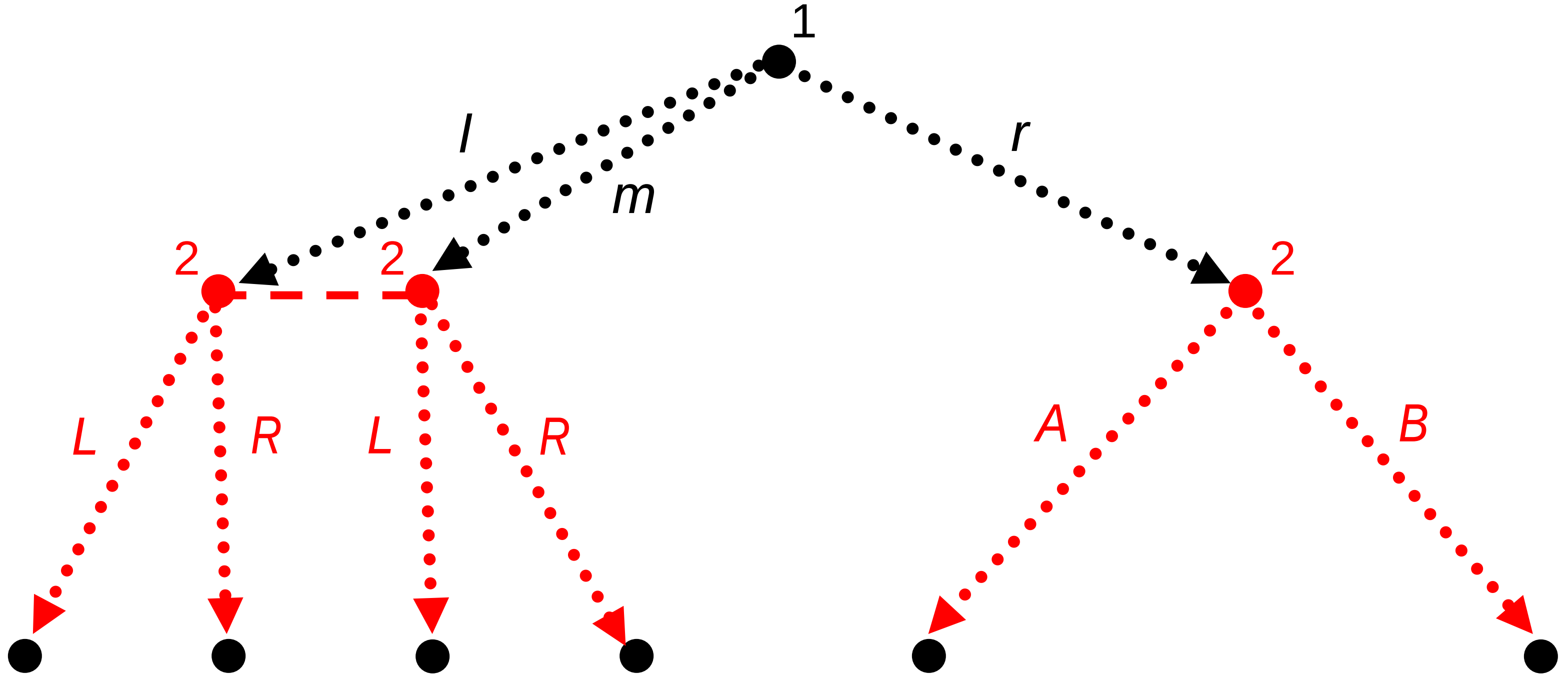

Subgames

A *(proper) subgame* is a subtree that:

-begins at a singleton information set;

-includes all subsequent nodes;

-and *does not cut any information sets*.

Idea:

Once a subgame begins, subsequent structure is *common knowledge.*

Example

This game has two subgames, rooted at 1.1 and 2.2.

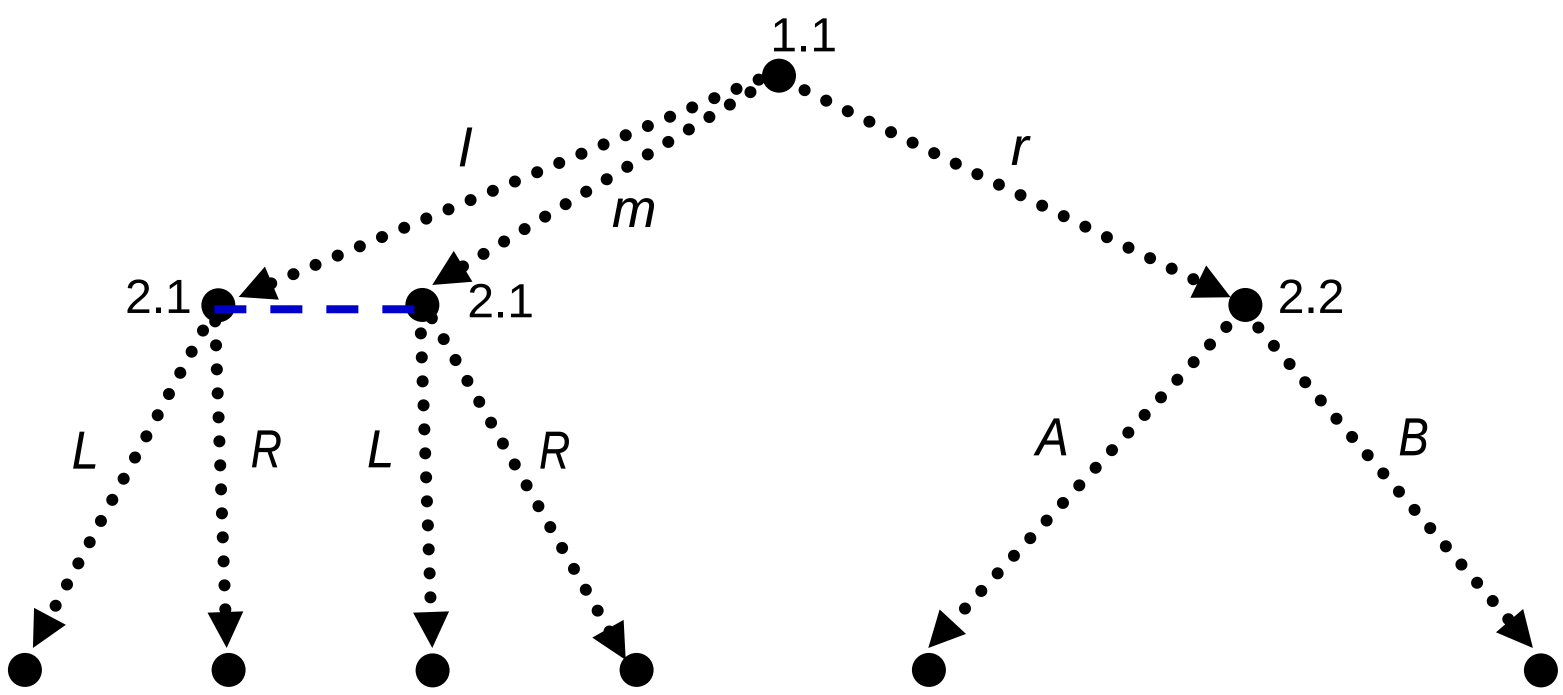

Extending backward induction

In games of perfect information, any subtree is a subgame.

What is the analog of backward induction?

Try to find “equilibrium” behavior from the “bottom” of the tree upwards.

Subgame perfection

A strategy vector (s_1, ..., s_N) is a *subgame perfect Nash equilibrium* (SPNE) if it induces a Nash equilibrium in (the strategic form of) every subgame.

(In games of perfect information, SPNE reduces to backward induction.)

Subgame perfection

"Every subgame" includes the game itself.

Idea:

-Find NE for "lowest" subgame

-Replace subgame subtree with equilibrium payoffs

-Repeat until we reach the root node of the original game

Subgame perfection

- As before, a SPNE specificies a *complete contingent plan* for each player.
- The *equilibrium path* is the actual play of the game under the SPNE strategies.
- As long as the game has finitely many stages, and finitely many actions at each information set, an SPNE always exists.

MS&E 246: Lecture 9
Sequential bargaining

Ramesh Johari

Nash bargaining solution

Recall Nash's approach to bargaining:

The planner is *given* the set of achievable payoffs and status quo point.

Implicitly:

The *process* of bargaining does not matter.

Dynamics of bargaining

In this lecture:

We use a dynamic game of perfect information to model the *process* of bargaining.

An interference model

Recall the interference model:

- Two devices
- Device 1 given channel a fraction q of the time
- For efficiency:
  When device n has control,
  it transmits at full power P

An interference model

- When timesharing is used, the set of Pareto efficient payoffs becomes:
  $\{ (\Pi_1, \Pi_2) : \Pi_1 = q\, R_1, \Pi_2 = (1 - q)\, R_2 \}$

- We now assume the devices bargain through a sequence of *alternating offers*.

Alternating offers

- At time 0:
  - Stage 0A:
    Device 1 *proposes* a choice of q (denoted q_1)
  - Stage 0B:
    Device 2 decides to *accept* or *reject* device 1's offer

Two period model

Assumption 1:

If device 2 *rejects* at stage 0B,
then predetermined choice $Q \in [0, 1]$
is implemented at time 1

Discounting

Assumption 2:

Devices care about delay:
Any payoff received by device i at time k is *discounted* by δ_i^k.

$0 < \delta_i < 1$: *discount factor* of device i

Game tree

Device 1 makes initial offer

Device 2 accepts or rejects

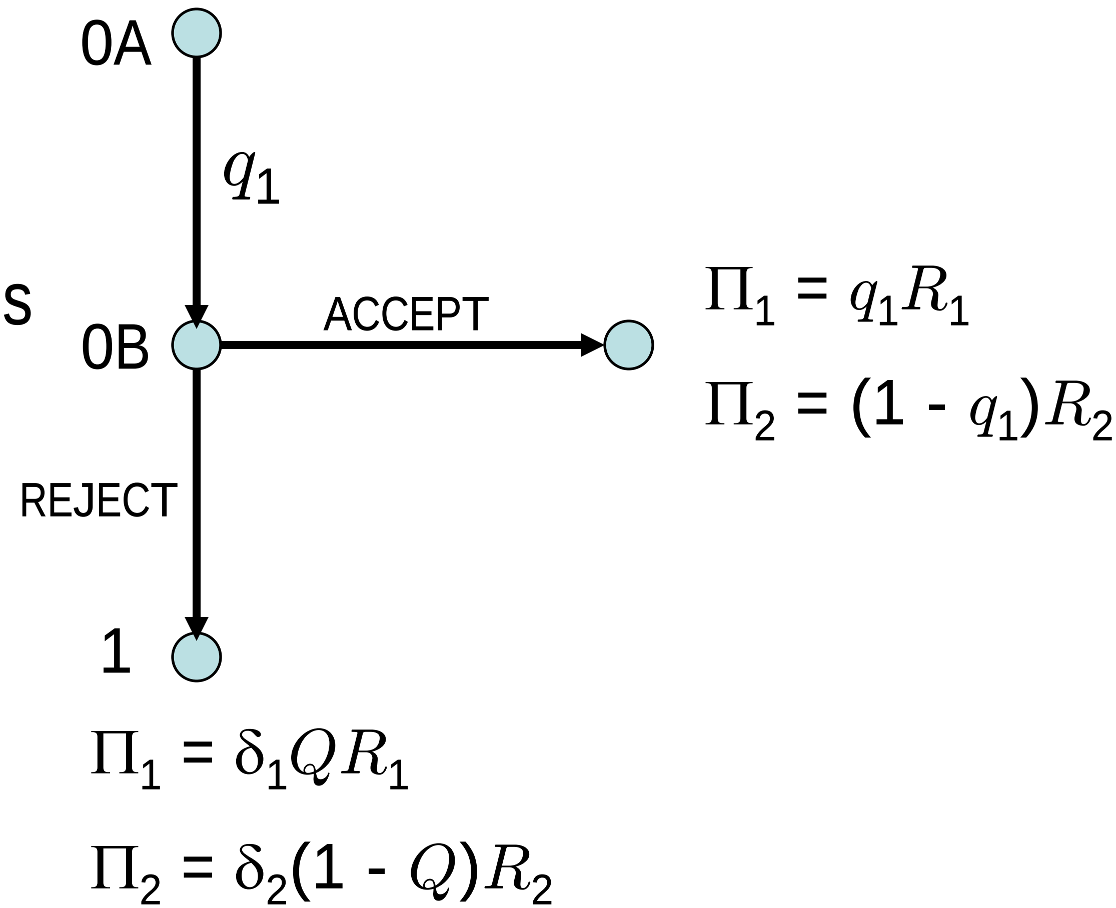

Game tree

- This is a *dynamic game of perfect information.*

- We solve it using *backward induction.*

Backward induction

1. Given q_1, at Stage 0B:

- Device 2 *rejects* if:
  $\delta_2 (1 - Q) > (1 - q_1)$

Backward induction

1. Given q_1, at Stage 0B:
   - Device 2 *rejects* ($s_2(q_1)$ = R) if:
     $q_1 > 1 - \delta_2 (1 - Q)$
   - Device 2 *accepts* ($s_2(q_1)$ = A) if:
     $q_1 < 1 - \delta_2 (1 - Q)$
   - Device 2 is *indifferent*
     ($s_2(q_1) \in \{$ A, R $\}$) if
     $q_1 = 1 - \delta_2 (1 - Q)$

Backward induction

2. At Stage 0A:

- Device 1 maximizes $\Pi_1(q_1, s_2(q_1))$ over offers ($0 \leq q_1 \leq 1$)
- *Claim:* Maximum value of Π_1 is

$$(1 - \delta_2 (1 - Q))\ R_1$$

Backward induction

2. At Stage 0A:

- *Claim:* Maximum value of Π_1 is

  $\Pi_1^{MAX} = (1 - \delta_2 (1 - Q)) R_1$

- *Proof:*

  (a) Maximum is achievable:

  If q_1 increases to $1 - \delta_2(1 - Q)$,

  then Π_1 increases to Π_1^{MAX}

Backward induction

2. At Stage 0A:

- *Claim:* Maximum value of Π_1 is

  $\Pi_1^{MAX} = (1 - \delta_2 (1 - Q)) R_1$

- *Proof:*

  (b) If $q_1 > 1 - \delta_2(1 - Q)$, then $\Pi_1 < \Pi_1^{MAX}$:

  Device 2 rejects $\Rightarrow$ $\Pi_1 = \delta_1 Q R_1$

Backward induction

2. At Stage 0A:

- *Claim:* Maximum value of Π_1 is

  $\Pi_1^{MAX} = (1 - \delta_2 (1 - Q)) R_1$

- *Proof:*

  (b) If $q_1 > 1 - \delta_2(1 - Q)$, then $\Pi_1 < \Pi_1^{MAX}$:

  But note that: $\delta_1 Q + \delta_2 (1 - Q) < 1$

Backward induction

2. At Stage 0A:

- *Claim:* Maximum value of Π_1 is

$\Pi_1^{MAX} = (1 - \delta_2 (1 - Q)) R_1$

- *Proof:*

(b) If $q_1 > 1 - \delta_2(1 - Q)$, then $\Pi_1 < \Pi_1^{MAX}$:

But note that: $\delta_1 Q < 1 - \delta_2 (1 - Q)$

Backward induction

2. At Stage 0A:

- *Claim:* Maximum value of Π_1 is

  $\Pi_1^{MAX} = (1 - \delta_2 (1 - Q)) R_1$

- *Proof:*

  (b) If $q_1 > 1 - \delta_2(1 - Q)$, then $\Pi_1 < \Pi_1^{MAX}$:

  So $\delta_1 Q R_1 < (1 - \delta_2 (1 - Q)) R_1$

Backward induction

2. At Stage 0A:

- *Best responses* for device 1 :

  All choices of q_1 that achieve Π_1^{MAX}

  The only possibility:

$$q_1^* = 1 - \delta_2 (1 - Q)$$

Backward induction

2. At Stage 0A:

- If $s_2(q_1^*)$ = reject,
  no best response exists for device 1!

- If $s_2(q_1^*)$ = accept,
  best response for device 1 is $q_1 = q_1^*$

Unique SPNE

What is the unique SPNE?

- Must give *strategies* for both players!

Unique SPNE

What is the unique SPNE?

- Device 1:

  At Stage 0A, offer $q_1 = q_1$*

- Device 2:

  At Stage 0B,

  accept if $q_1 \leq q_1$*, *reject* if $q_1 > q_1$*

Payoffs at unique SPNE

- So the offer of device 1 is *accepted immediately* by device 2.

- Device 1 gets: $\Pi_1 = (1 - \delta_2(1 - Q))\, R_1$

- Device 2 gets: $\Pi_2 = \delta_2(1 - q_0)\, R_2$

Infinite horizon

More realistic model:

Devices alternate offers indefinitely.

For simplicity: assume $\delta_1 = \delta_2 = \delta$

Finite horizon

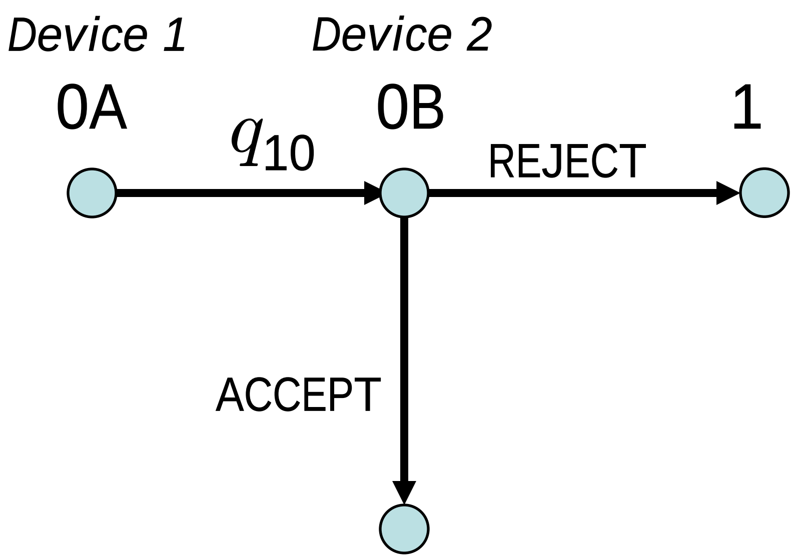

$$\Pi_1 = q_{10} R_1$$

$$\Pi_2 = (1 - q_{10}) R_2$$

Infinite horizon

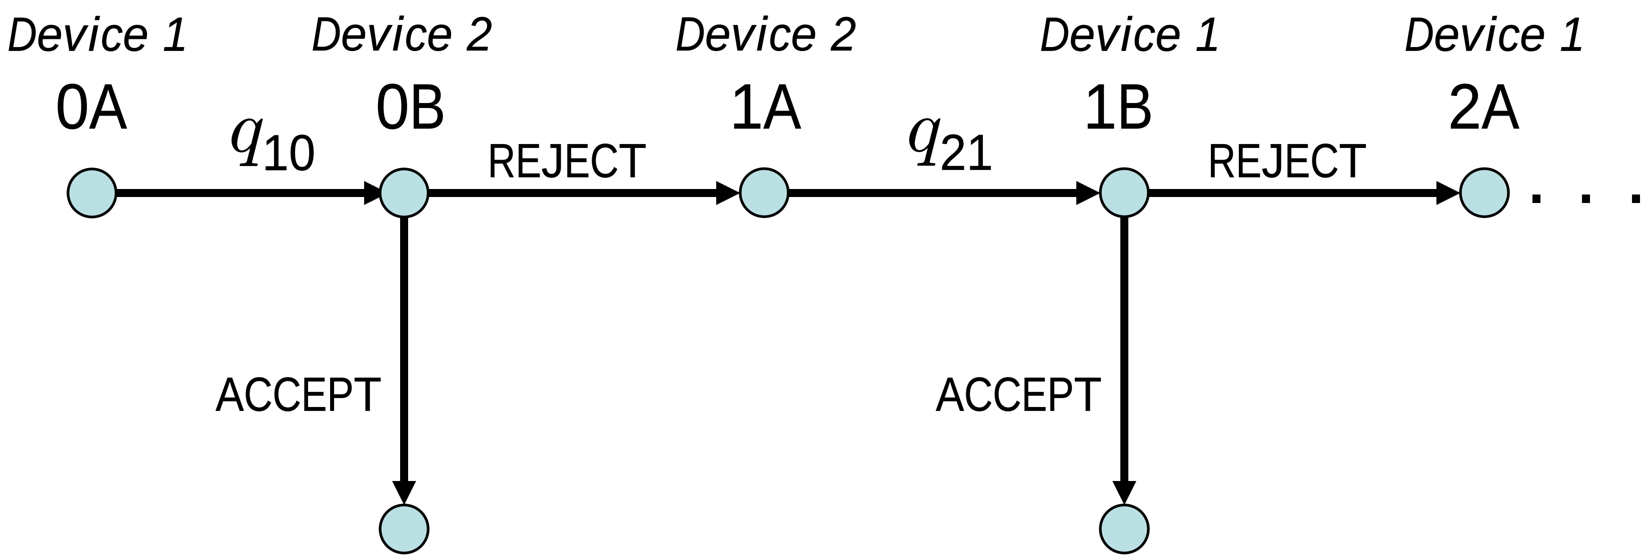

$\Pi_1 = q_{10}R_1$

$\Pi_2 = (1 - q_{10})R_2$

$\Pi_1 = \delta q_{21}R_1$

$\Pi_2 = \delta(1 - q_{21})R_2$

Infinite horizon: formal model

- Device 1 offers q_{1k} at stage kA, for k even
- Device 2 offers q_{2k} at stage kA, for k odd

- Device 2 accepts/rejects stage kA offer at stage kB, for k even
- Device 1 accepts/rejects stage kA offer at stage kB, for k odd

Infinite horizon: formal model

- Payoffs:

  $\Pi_1 = \Pi_2 = 0$ if no offer ever accepted (similar to status quo in NBS)

Infinite horizon: formal model

- Payoffs:

  If offer made at stage kA by player i accepted at stage kB :

$$\Pi_1 = \delta^k \, q_{ik} \, R_1$$

$$\Pi_2 = \delta^k \, (1 - q_{ik}) \, R_2$$

Infinite horizon

- Can't use backward induction!

- Use *stationarity:*

  Subgame rooted at 1A is
  the same as the original game,
  with roles of 1 and 2 reversed.

SPNE

Define V and v:

$V\,R_1$ = highest time 0 payoff to device 1 among *all* SPNE

$v\,R_1$ = lowest time 0 payoff to device 1 among *all* SPNE

SPNE

Then if device 2 rejects at 0B:

$V\ R_2$= highest time 1 payoff to device 2 among *all* SPNE

$v\ R_2$= lowest time 1 payoff to device 2 among *all* SPNE

SPNE: Two inequalities

- $v R_1 \geq (1 - \delta V) R_1$

At Stage 0B:

Device 2 will accept any $q_{10} < 1 - \delta V$

So at Stage 0A:

Device 1 must earn at least $(1 - \delta V) R_1$

SPNE: Two inequalities

- $V R_1 \leq (1 - \delta v) R_1$

  If offer q_{10} is *accepted* at stage 0B, device 2 must get a timeshare of at least δv

  $\Rightarrow \quad q_{10} \leq 1 - \delta v$

SPNE: Two inequalities

- $V R_1 \leq (1 - \delta v) R_1$

  If offer q_{10} is *rejected* at stage 0B,
  device 1 earns at most $\delta (1 - v) R_1$
  since device 2 earns at least $\delta v R_2$
  $\Rightarrow \Pi_1 \leq \delta (1 - v) R_1 \leq (1 - \delta v) R_1$

Combining inequalities

- $v \leq V$
- $v \geq 1 - \delta V$
- $V \leq 1 - \delta v$

Combining inequalities

- $v \leq V$
- $v + \delta V \geq 1$
- $V + \delta v \leq 1$

So: $V + \delta v \leq v + \delta V$

$\Rightarrow$ $(1 - \delta) V \leq (1 - \delta) v$

$\Rightarrow$ $V = v$

Unique SPNE

- So $V = 1 - \delta V \Rightarrow$

$$V = \frac{1}{1+\delta}$$

- SPNE strategies for device 1:

  At Stage kA , k even:

  Offer $q_{1k} = 1 - \delta V$

  At Stage kB , k odd:

  Accept if $q_{2k} \geq \delta V$

Unique SPNE

- So $V = 1 - \delta V \Rightarrow$

$$V = \frac{1}{1+\delta}$$

- SPNE strategies for device 2:

At Stage kA , k odd:

Offer $q_{2k} = \delta V$

At Stage kB , k even:

Accept if $q_{1k} \leq 1 - \delta V$

Unique SPNE: Payoffs

Stage 0A offer by device 1 is accepted in Stage 0B by device 2.

$$\Pi_1^{\mathrm{SPNE}} = \frac{R_1}{1+\delta}, \quad \Pi_2^{\mathrm{SPNE}} = \frac{\delta R_2}{1+\delta}$$

Infinite horizon: Discussion

- Outcome is *efficient:*
  No "lost utility" due to discounting

- *Stationary* SPNE strategies:
  Actions do not depend on time k

- *First mover* advantage:
  $\Pi_1^{\text{SPNE}} > \Pi_2^{\text{SPNE}}$

Shortening time periods

Shorten each time step to length $t < 1$ ...

... Same as changing discount factor to δ^t

$$\Pi_1^{\text{SPNE}} = \frac{R_1}{1+\delta^t}, \qquad \Pi_2^{\text{SPNE}} = \frac{\delta^t R_2}{1+\delta^t}$$

As $t \to 0$, note that $\Pi_i^{\text{SPNE}} \to R_i/2$.

Nash bargaining solution!

In general

If $\delta_1 \neq \delta_2$:

Find SPNE using two period model:

Note that Q must be SPNE payoff when device 2 offers first

Can show (for an appropriate limit) that *weighted NBS* obtained as $t \to 0$:

More patient player weighted higher

Summary

- Alternating offers: finite horizon

  Backward induction solution

- Alternating offers: infinite horizon

  Unique SPNE

  Relation to Nash bargaining solution

MS&E 246: Lecture 10
Repeated games

Ramesh Johari

What is a repeated game?

A *repeated game* is:

A dynamic game constructed by playing the same game over and over.

It is a dynamic game of imperfect information.

This lecture

- Finitely repeated games
- Infinitely repeated games
  - Trigger strategies
  - The folk theorem

Stage game

At each stage, the same game is played: the *stage game* G.

Assume:

- G is a simultaneous move game
- In G, player i has:
  - Action set A_i
  - Payoff $P_i(a_i, \mathbf{a}_{-i})$

Finitely repeated games

$G(K)$: G is repeated K times

Information sets:

All players observe outcome of each stage.

What are:

strategies? payoffs? equilibria?

History and strategies

Period t history h_t:

h_t = ($\mathbf{a}$(0), ..., $\mathbf{a}$(t-1)) where

$\mathbf{a}(\tau)$ = action profile played at stage τ

Strategy s_i:

Choice of stage t action $s_i(h_t) \in A_i$

for each history h_t

i.e. $a_i(t) = s_i(h_t)$

Payoffs

Assume payoff = *sum of stage game payoffs*

$$\Pi_i(\mathbf{s}) = \sum_{t=0}^{K-1} P_i(s_1(h_t), \ldots, s_N(h_t))$$

Example: Prisoner's dilemma

Recall the Prisoner's dilemma:

		Player 1	
		defect	cooperate
Player 2	defect	(1,1)	(4,0)
	cooperate	(0,4)	(2,2)

Example: Prisoner's dilemma

Two volunteers

Five rounds

No communication allowed!

Round	1	2	3	4	5	Total
Player 1	1	1	1	1	1	5
Player 2	1	1	1	1	1	5

SPNE

Suppose $\mathbf{a}^{\text{NE}}$ is a stage game NE.

Any such NE gives a SPNE:

Player i plays a_i^{NE} at every stage, regardless of history.

Question: Are there any other SPNE?

SPNE

How do we find SPNE of $G(K)$?

Observe:

Subgame starting after history h_t is identical to $G(K - t)$

SPNE: Unique stage game NE

Suppose G has a unique NE $\mathbf{a}^{\text{NE}}$

Then regardless of period K history h_K , last stage has unique NE $\mathbf{a}^{\text{NE}}$

$\Rightarrow$ At SPNE, $s_i(h_K) = a_i^{\text{NE}}$

SPNE: Backward induction

At stage K - 1, *given* $\mathbf{s}_{-i}(\cdot)$, player i chooses $s_i(h_{K-1})$ to maximize:

$$\underbrace{P_i(s_i(h_{K-1}),\ \mathbf{s}_{-i}(h_{K-1}))}_{\textit{payoff at stage } K-1} + \underbrace{P_i(\mathbf{s}(h_K))}_{\textit{payoff at stage } K}$$

SPNE: Backward induction

At stage K - 1, *given* $\mathbf{s}_{-i}(\cdot)$, player i chooses $s_i(h_{K-1})$ to maximize:

$$\underbrace{P_i(s_i(h_{K-1}),\ \mathbf{s}_{-i}(h_{K-1}))}_{\textit{payoff at stage } K-1} \quad + \quad \underbrace{P_i(\mathbf{a}^{NE})}_{\textit{payoff at stage } K}$$

We know: at last stage, $\mathbf{a}^{NE}$ is played.

SPNE: Backward induction

At stage K - 1, *given* $\mathbf{s}_{-i}(\cdot)$, player i chooses $s_i(h_{K-1})$ to maximize:

$$\underbrace{P_i(s_i(h_{K-1}), \mathbf{s}_{-i}(h_{K-1}))}_{\textit{payoff at stage } K-1}$$

$\Rightarrow$ *Stage game NE again!*

SPNE: Conclusion

Theorem:

If stage game has unique NE $\mathbf{a}^{NE}$, then finitely repeated game has unique SPNE:

$s_i(h_t) = a_i^{NE}$ for all h_t

Example: Prisoner's dilemma

Moral: "Cooperate" should never be played.

Axelrod's tournament (1980):

Winning strategy was "tit for tat":

Cooperate if and only if
your opponent did so at the last stage

SPNE: Multiple stage game NE

Note:

If multiple NE exist for stage game NE,
there may exist SPNE where actions are played that appear in *no stage game NE*

(See Gibbons, 2.3.A)

Infinitely repeated games

- History, strategy definitions same as finitely repeated games
- Payoffs:

  Sum might not be finite!

Discounting

Define payoff as:

$$\Pi_i(\mathbf{s}) = (1-\delta)\sum_{t=0}^{\infty}\delta^t P_i(s_1(h_t), \ldots, s_N(h_t))$$

i.e., *discounted sum* of stage game payoffs

This game is denoted $G(\delta, \infty)$

(*Note:* (1 - δ) is a normalization)

Discounting

Two interpretations:

1. Future payoffs worth less than today's payoffs

2. Total # of stages is a geometric random variable

Folk theorems

- Major problem with infinitely repeated games:

  If players are patient enough,
  SPNE can achieve "any" reasonable payoffs.

Prisoner's dilemma

Consider the following strategies, (s_1, s_2):

1. Play C at first stage.
2. If h_t = ((C,C), ..., (C,C)),
   then play C at stage t.

   Otherwise play D.

i.e., punish the other player for defecting

Prisoner's dilemma

Note: $G(\delta, \infty)$ is *stationary*

Case 1: Consider any subgame where at least one player has defected in h_t.

Then (D,D) played forever.

This is NE for subgame,
since (D,D) is stage game NE.

Prisoner's dilemma

Step 2: Suppose h_t = ((C,C), ..., (C,C)).

Player 1's options:

(a) Follow s_1 $\Rightarrow$ play C forever

(b) Deviate at time t $\Rightarrow$ play D forever

Prisoner's dilemma

Given s_2:

Playing C forever gives payoff:

$(1-\delta)\ (\ P_1(C,C) + \delta\ P_1(C,C) + \ldots\ ) = P_1(C,C)$

Playing D forever gives payoff:

$(1-\delta)\ (\ P_1(D,C) + \delta\ P_1(D,D) + \ldots\ )$

$= (1-\delta)\ P_1(D,C) + \delta\ P_1(D,D)$

Prisoner's dilemma

So cooperate if and only if:

$P_1(C,C) \geq (1 - \delta)\ P_1(D,C) + \delta\ P_1(D,D)$

Note: if $P_1(C,C) > P_1(D,D)$,

then this is always true for δ close to 1

Conclude:

If δ close to 1, then (s_1, s_2) is an SPNE

Prisoner's dilemma

In our game:

Need $2 \geq (1 - \delta)\, 4 + \delta \quad \Rightarrow \quad \delta \geq 2/3$

So cooperation can be sustained if
time horizon is *finite but uncertain*.

Trigger strategies

In a *(Nash) trigger strategy* for player i :

1. Play a_i at first stage.
2. If h_t = ($\mathbf{a}$, ..., $\mathbf{a}$),
   then play a_i at stage t.
   Otherwise play a_i^{NE}.

Trigger strategies

If $\mathbf{a}$ Pareto dominates $\mathbf{a}^{NE}$,

trigger strategies will be an SPNE for large enough δ

Formally: need

$P_i(\mathbf{a}) > (1 - \delta)\ P_i(a_i', \mathbf{a}_{-i}) + \delta\ P_i(\mathbf{a}^{NE})$

for all players i and actions a_i'.

Achievable payoffs

Achievable payoffs:

T = Convex hull of $\{ (P_1(\mathbf{a}), P_2(\mathbf{a})) : a_i \in S_i \}$

e.g., in Prisoner's Dilemma:

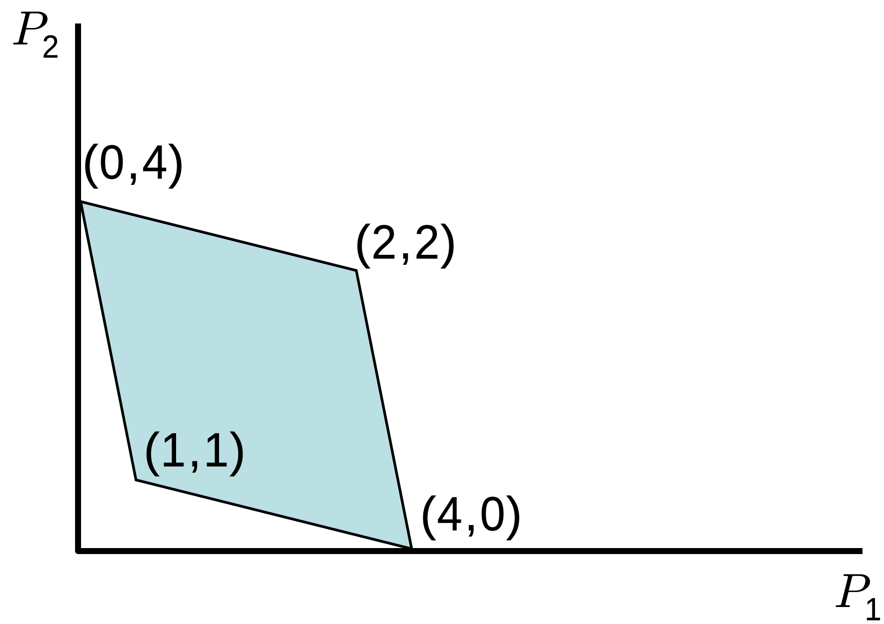

Achievable payoffs and SPNE

A key result in repeated games:

Any "reasonable" achievable payoff can be realized in an SPNE of the repeated game, if players are patient enough.

Simple proof: generalize prisoner's dilemma.

Randomization

- To generalize, suppose before stage t all players observe i.i.d. uniform r.v. U_t
- History:
  $h_t = (\mathbf{a}(0), \ldots, \mathbf{a}(t\text{-}1), U_0, \ldots, U_t)$
- Players can use U_t to *coordinate* strategies at stage t

Randomization

E.g., suppose players want to achieve

$\mathbf{P} = \alpha\, \mathbf{P}(\mathbf{a}) + (1 - \alpha)\, \mathbf{P}(\mathbf{a}')$

If $U_t \leq \alpha$: Player i plays a_i

If $U_t > \alpha$: Player i plays a_i'

We'll call this the $\mathbf{P}$-*achieving action* for i.

(Uniquely defined for all $\mathbf{P} \in T$.)

Randomization

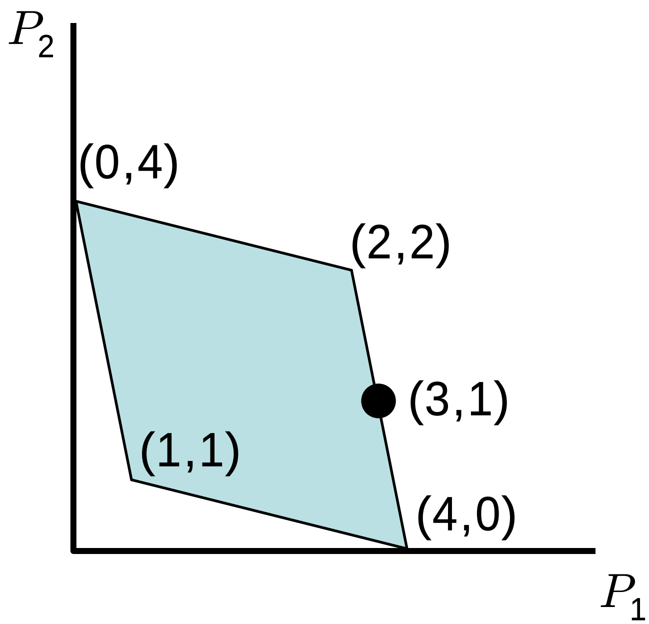

E.g., Prisoner's Dilemma

Let $\mathbf{P}$ = (3,1).

$\mathbf{P}$-achieving actions:

Player 1 plays C if $U_t \leq$ ½
  and D if U_t > ½

Player 2 plays C if $U_t \leq$ ½
  and C if U_t > ½

Randomization and triggering

So now suppose $\mathbf{P} \in T$ and:

$P_i > P_i(\mathbf{a}^{\text{NE}})$ for all i

Trigger strategy:

Punish forever (by playing a_i^{NE}) if opponent deviates from $\mathbf{P}$-achieving action

Randomization and triggering

Both players using this trigger strategy is again an SPNE for large enough δ.

Formally: need

$$(1 - \delta)\, P_i(p_i,\, \mathbf{p}_{-i}) + \delta\, P_i$$
$$> (1 - \delta)\, P_i(a_i',\, \mathbf{p}_{-i}) + \delta\, P_i(\mathbf{a}^{\text{NE}})$$

for all players i and actions a_i'.

(Here p is $\mathbf{P}$-achieving action for player i, and $\mathbf{p}_{-i}$ is $\mathbf{P}$-achieving action vector for all other players.)

Randomization and triggering

Both players using this trigger strategy is again an SPNE for large enough δ.

Formally: need

$$(1 - \delta)\ P_i(p_i,\ \mathbf{p}_{-i}) + \delta\ P_i$$
$$> (1 - \delta)\ P_i(a_i',\ \mathbf{p}_{-i}) + \delta\ P_i(\mathbf{a}^{NE})$$

for all players i and actions a_i'.

(At time t:
LHS is payoff if player i does not deviate after seeing U_t;
RHS is payoff if player i deviates to a_i' after seeing U_t)

Folk theorem

Theorem (Friedman, 1971):

Fix a Nash equilibrium $\mathbf{a}^{\text{NE}}$, and

$\mathbf{P} \in T$ such that

$$P_i > P_i(\mathbf{a}^{\text{NE}}) \text{ for all } i$$

Then for large enough δ,

there exists an SPNE $\mathbf{s}$ such that:

$$\Pi_i(\mathbf{s}) = P_i$$

Minimax payoffs

What is the *minimum* payoff Player 1 can guarantee himself?

$$\min_{a_2 \in A_2} \left\{ \quad \max_{a_1 \in A_1} P_1(a_1, a_2) \quad \right\}$$

Minimax payoffs

What is the *minimum* payoff Player 1 can guarantee himself?

$$\min_{a_2 \in A_2} \left\{ \underbrace{\max_{a_1 \in A_1} P_1(a_1, a_2)} \right\}$$

Given a_2, this is the highest payoff player 1 can get…

Minimax payoffs

What is the *minimum* payoff Player 1 can guarantee himself?

$$\underbrace{\min_{a_2 \in A_2} \left\{ \quad \max_{a_1 \in A_1} P_1(a_1, a_2) \quad \right\}}$$

...so Player 1 can guarantee himself this payoff if he knows how Player 2 is punishing him

Minimax payoffs

What is the *minimum* payoff Player 1 can guarantee himself?

$$\min_{a_2 \in A_2} \left\{ \max_{a_1 \in A_1} P_1(a_1, a_2) \right\}$$

This is m_1, the *minimax value* of Player 1.

Generalization

Theorem (Fudenberg and Maskin, 1986):

Folk theorem holds for all $\mathbf{P}$ such that $P_i > m_i$ for all i

(Technical note:

This result requires that dimension of T = # of players)

Finite vs. infinite

Theorem (Benoit and Krishna, 1985):

Assume: for each i, we can find two NE $\mathbf{a}^{\text{NE}}$, $\underline{\mathbf{a}}^{\text{NE}}$ such that $P_i(\mathbf{a}^{\text{NE}}) > P_i(\underline{\mathbf{a}}^{\text{NE}})$

Then as $K \to \infty$,
set of SPNE payoffs of $G(K)$
approaches $\{ \mathbf{P} \in T : P_i > m_i \}$

(Same technical note as Fudenberg-Maskin applies)

Finite vs. infinite

In the unique Prisoner's Dilemma NE,
only one NE exists

$\Rightarrow$ Benoit-Krishna result fails

Note at Prisoner's Dilemma NE,
each player gets minimax value.

Summary

Repeated games are a simple way to model interaction over time.

(1) In general, too many SPNE $\Rightarrow$ not very good predictive model

(2) However, can gain insight from *structure* of SPNE strategies

MS&E 246: Lecture 11 Concluding remarks on subgame perfection

Ramesh Johari

Dynamic games

In our discussion of dynamic games of complete information,
we studied two main types:

- Perfect information
- Imperfect information

In both cases, subgame perfect NE emerged as a natural way to capture “sequential rationality” (or credibility).

Possible problems

However, subgame perfection can give rise to two possible issues.

In some cases, it is overly restrictive as a predictive tool:
there are “not enough” SPNE.

In some cases, it is not useful as a predictive tool:
there are too many SPNE.

SPNE: overly restrictive?

Consider the following game:

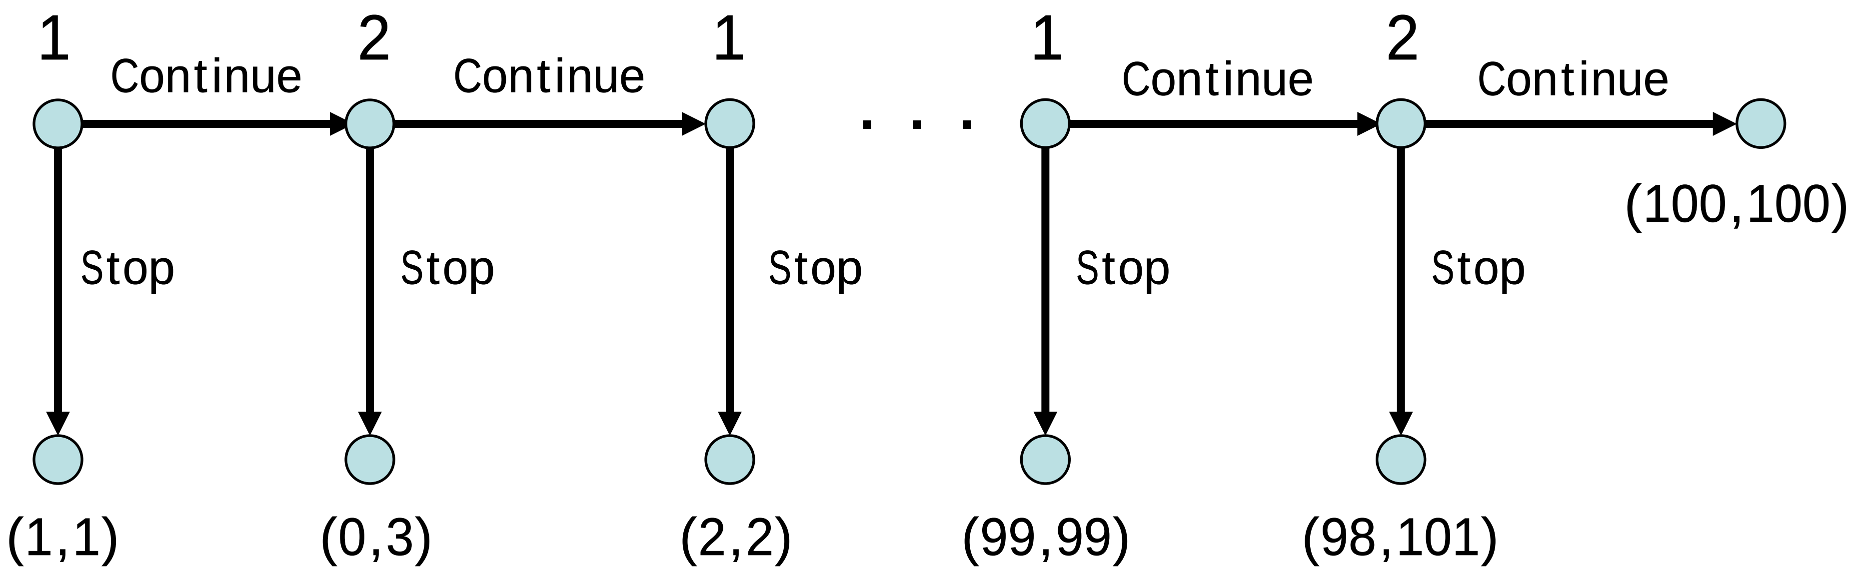

(This is called the “centipede” game.)

SPNE: overly restrictive?

- In last information set, player 2 prefers to “stop” instead of “continue”
- Inductively, in each information set each player prefers to “stop” instead of “continue”
- Equilibrium payoffs: (1,1)
- Is this a reasonable prediction of play?

SPNE: overly restrictive?

The centipede game reveals a key flaw in the definition of SPNE:

If play ever reaches a subgame *off* the equilibrium path of play,
then rationality must have failed already.

But SPNE assumes rational behavior in *every subgame!*

SPNE: not restrictive enough?

- In repeated games, we saw the folk theorem(s): with enough patience, any individually rational payoffs can be sustained by an SPNE.
- Too many equilibria for predictive use

SPNE: not restrictive enough?

Other problems can occur in situations where there are “not enough subgames” to rule out equilibria.

SPNE: not restrictive enough?

- Two firms
- First firm decides if/how to enter
- Second firm can choose to "fight"

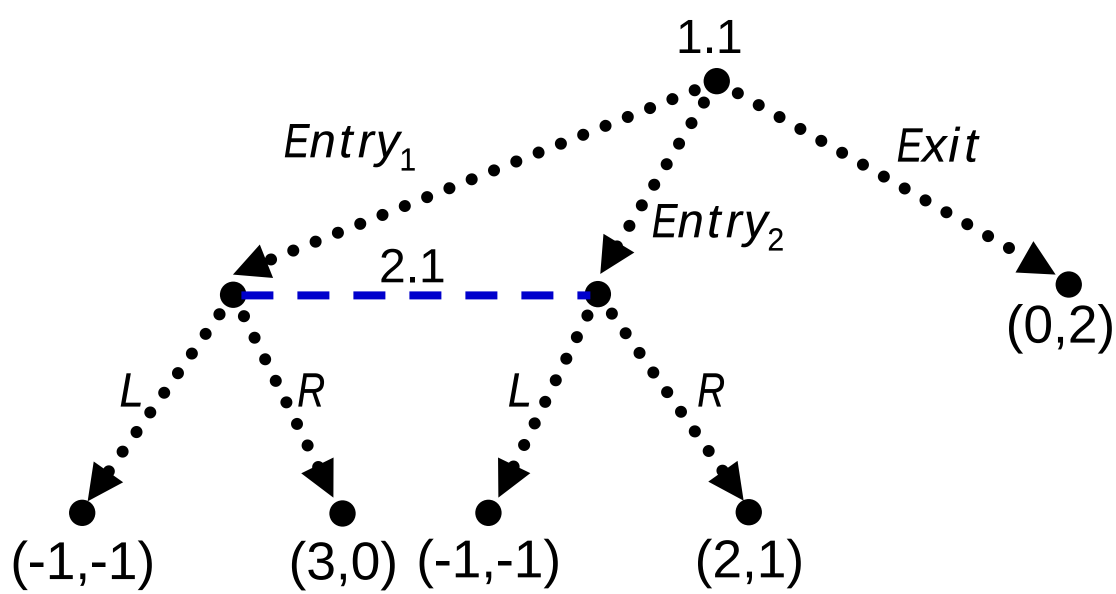

Entry example

Note that this game only has *one* subgame.

Thus SPNE are *any* NE of strategic form.

		Firm 2	
		L	*R*
	$Entry_1$	(-1,-1)	(3,0)
Firm 1	$Entry_2$	(-1,-1)	(2,1)
	Exit	(0,2)	(0,2)

Entry example

Two pure NE of strategic form:

($Entry_1$, *R*) and (*Exit*, *L*)

Firm 2

Firm 1

	L	*R*
$Entry_1$	(-1,-1)	(3,0)
$Entry_2$	(-1,-1)	(2,1)
Exit	(0,2)	(0,2)

SPNE: not restrictive enough?

But firm 1 should *“know”* that if
it chooses to enter,
firm 2 will never “fight.”

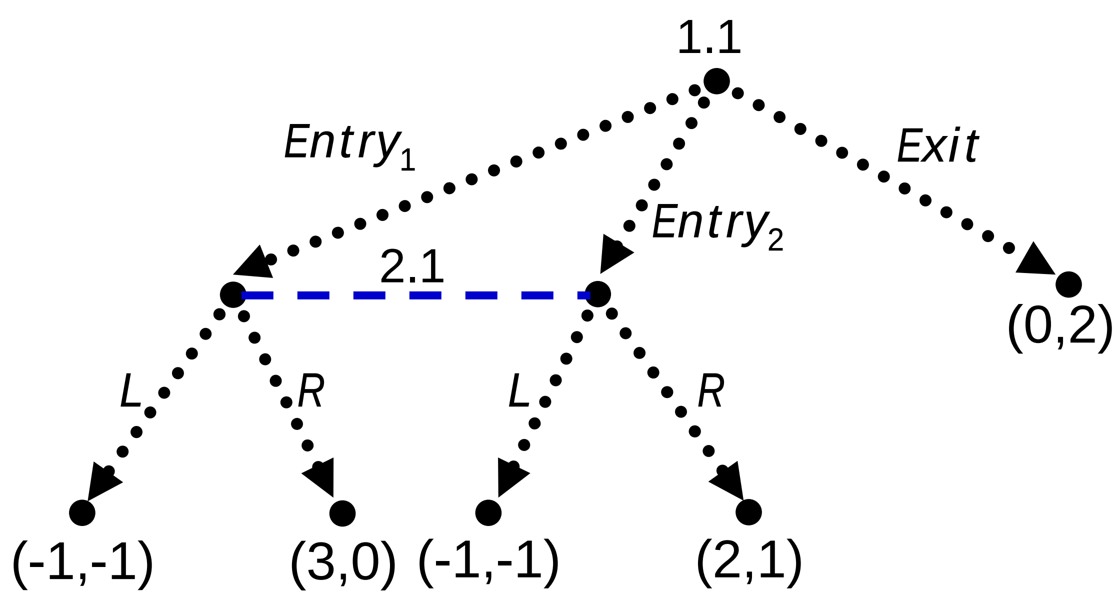

SPNE: not restrictive enough?

So in this situation, there are again too many SPNE.

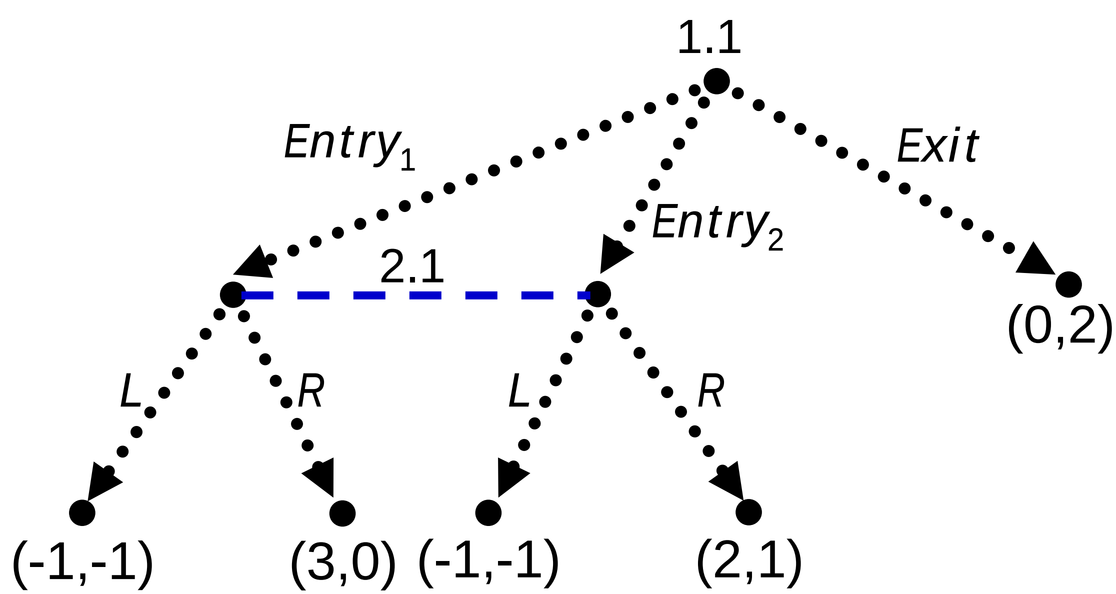

SPNE: not restrictive enough?

A solution to the problem of the entry game is to include *beliefs* as part of the solution concept:

Firm 2 should never fight, regardless of what it believes firm 1 played.

(We will study such an approach in the last part of the course.)

MS&E 246: Lecture 12
Static games of incomplete information

Ramesh Johari

Incomplete information

- *Complete information* means the entire structure of the game is common knowledge

- *Incomplete information* means "anything else"

Cournot revisited

Consider the Cournot game again:

- Two firms ($N = 2$)
- Cost of producing s_n : $c_n\, s_n$
- *Demand curve*:

  Price = $P(s_1 + s_2) = a - b\,(s_1 + s_2)$
- Payoffs:

  Profit = $\Pi_n(s_1, s_2) = P(s_1 + s_2)\, s_n - c_n\, s_n$

Cournot revisited

Suppose c_i is known only to firm i.

Incomplete information:

Payoffs of firms are not common knowledge.

How should firm i *reason* about what it expects the competitor to produce?

Beliefs

A first approach:

Describe firm 1's *beliefs* about firm 2.

Then need firm 2's beliefs about firm 1's beliefs…

And firm 1's beliefs about firm 2's beliefs about firm 1's beliefs…

Rapid growth of complexity!

Harsanyi's approach

Harsanyi made a key breakthrough in the analysis of incomplete information games:

He represented a *static game of incomplete information* as a *dynamic game of complete but imperfect information.*

Harsanyi's approach

- *Nature* chooses (c_1, c_2) according to some probability distribution.
- Firm i now knows c_i, but opponent only knows the *distribution* of c_i.
- Firms maximize *expected* payoff.

Harsanyi's approach

All firms share the *same* beliefs about incomplete information, determined by the *common prior*.

Intuitively:

Nature determines "who we are" at time 0, according to a distribution that is common knowledge.

Bayesian games

In a Bayesian game, player i's preferences are determined by his *type* $\theta_i \in T_i$; T_i is the *type space* for player i.

When player i has type θ_i, his payoff function is: $\Pi_i(\mathbf{s} \,;\, \theta_i)$

Bayesian games

A *Bayesian game* with N players is a dynamic game of $N + 1$ stages.

Stage 0: *Nature* moves first, and chooses a type θ_i for each player i, according to some joint distribution $P(\theta_1, \ldots, \theta_N)$ on $T_1 \times \cdots \times T_N$.

Bayesian games

Player i learns his own type θ_i,
but not the types of other players.

Stage i: Player i chooses an action
from S_i.

After stage N: payoffs are realized.

Bayesian games

An alternate, equivalent interpretation:

1) Nature chooses players' types.
2) All players *simultaneously* choose actions.
3) Payoffs are realized.

Bayesian games

- How many subgames are there?
- How many information sets does player i have?
- What is a strategy for player i?

Bayesian games

- How many subgames are there?

  ONE – the entire game.

- How many information sets does player i have?

  ONE per type θ_i.

- What is a strategy for player i?

  A function from $T_i \rightarrow S_i$.

Bayesian games

One subgame $\Rightarrow$
SPNE and NE are identical concepts.

A *Bayesian equilibrium*
(or *Bayes-Nash equilibrium*)
is a NE of this dynamic game.

Expected payoffs

How do players reason about uncertainty regarding *other players' types*?

They use *expected payoffs*.

Given $\mathbf{s}_{-i}(\cdot)$, player i chooses $a_i = s_i(\theta_i)$ to maximize

$$E[\ \Pi_i(\ a_i\ ,\ \mathbf{s}_{-i}(\theta_{-i})\ ;\ \theta_i\ )\ |\ \theta_i\ ]$$

Bayesian equilibrium

Thus $s_1(\cdot)$, ..., $s_N(\cdot)$
is a Bayesian equilibrium if and only if:

$s_i(\theta_i) \in \arg\max_{a_i \in S_i} E[\, \Pi_i(\, a_i \,, \mathbf{s}_{-i}(\theta_{-i}) \,;\, \theta_i) \mid \theta_i \,]$

for all θ_i, and for all players i.

[Notation: $\mathbf{s}_{-i}(\theta_{-i}) = (s_1(\theta_1), \ldots, s_{i-1}(\theta_{i-1}), s_{i+1}(\theta_{i+1}), \ldots, s_N(\theta_N))$]

Bayesian equilibrium

The conditional distribution $P(\theta_{-i} \mid \theta_i)$ is called player i's *belief*.

Thus in a Bayesian equilibrium, players maximize expected payoffs given their beliefs.

[*Note:* Beliefs are found using *Bayes' rule*:

$P(\theta_{-i} \mid \theta_i) = P(\theta_1, \ldots, \theta_N)/P(\theta_i)$]

Bayesian equilibrium

Since Bayesian equilibrium is a NE
of a certain dynamic game
of imperfect information,
it is guaranteed to exist
(as long as type spaces T_i are finite
and action spaces S_i are finite).

Cournot revisited

Back to the Cournot example:

Let's assume c_1 is common knowledge, but firm 1 does not know c_2.

Nature chooses c_2 according to:

c_2 = c_H, w/prob. p

= c_L, w/prob. 1 - p

(where $c_H > c_L$)

Cournot revisited

These are also informally called games of *asymmetric information*:

Firm 2 has information that firm 1 does not have.

Cournot revisited

- The *type* of each firm is their marginal cost of production, c_i.
- Note that c_1, c_2 are *independent.*
- Firm 1's *belief*:

  $P(c_2 = c_H \mid c_1) = 1 - P(c_2 = c_L \mid c_1) = p$
- Firm 2's *belief*: Knows c_1 exactly

Cournot revisited

- The *strategy* of firm 1 is the quantity s_1.
- The *strategy* of firm 2 is a function:

  $s_2(c_H)$: quantity produced if $c_2 = c_H$

  $s_2(c_L)$: quantity produced if $c_2 = c_L$

Cournot revisited

We want a NE of the dynamic game.

Given s_1 *and* c_2, Firm 2 plays a best response.

Thus given s_1 and c_2, Firm 2 produces:

$$R_2(s_1) = \left[\frac{a - c_2}{2b} - \frac{s_1}{2}\right]^+$$

Cournot revisited

Given $s_2(\cdot)$ and c_1,
Firm 1 maximizes *expected payoff*.

Thus Firm 1 maximizes:

$$E[\ \Pi_1 (s_1,\ s_2(c_2)\ ;\ c_1)\ |\ c_1\ ] =$$

$$[p\ P(s_1 + s_2(c_H)) + (1 - p)P(s_1 + s_2(c_L))]\ s_1 - c_1 s_1$$

Cournot revisited

- Recall demand is *linear:* $P(Q) = a - b\,Q$
- So expected payoff to firm 1 is:

  $[\, a - b(s_1 + p\, s_2(c_H) + (1 - p)s_2(c_L))\,]s_1 - c_1\, s_1$
- Thus firm 1 plays best response to *expected production of firm 2.*

$$R_1(s_2) = \left[\frac{a - c_1}{2b} - \frac{p s_2(c_H) + (1 - p) s_2(c_L)}{2}\right]^+$$

Cournot revisited: equilibrium

- A Bayesian equilibrium has 3 unknowns:

  s_1, $s_2(c_H)$, $s_2(c_L)$
- There are 3 equations:

  Best response of firm 1 given $s_2(c_H)$, $s_2(c_L)$

  Best response of firm 2 given s_1,
  when type is c_H

  Best response of firm 2 given s_1,
  when type is c_L

Cournot revisited: equilibrium

- Assume all quantities are positive at BNE (can show this must be the case)
- Solution:

$$s_1 = [\, a - 2c_1 + pc_{\mathrm{H}} + (1 - p)c_{\mathrm{L}} \,]/3$$

$$s_2(c_{\mathrm{H}}) = [a - 2c_{\mathrm{H}} + c_1]/3 + (1 - p)(c_{\mathrm{H}} - c_{\mathrm{L}})/6$$

$$s_2(c_{\mathrm{L}}) = [a - 2c_{\mathrm{L}} + c_1]/3 - p(c_{\mathrm{H}} - c_{\mathrm{L}})/6$$

Cournot revisited: equilibrium

When $p = 0$, *complete information*:

Both firms know $c_2 = c_L \Rightarrow$ NE $(s_1^{(0)}, s_2^{(0)})$

When $p = 1$, *complete information*:

Both firms know $c_2 = c_H \Rightarrow$ NE $(s_1^{(1)}, s_2^{(1)})$

For $0 < p < 1$, note that:

$s_2(c_L) < s_2^{(0)}$, $s_2(c_H) > s_2^{(1)}$, $s_1^{(0)} < s_1 < s_1^{(1)}$

Why?

Coordination game

Consider coordination game with incomplete information:

		Player 2	
		L	*R*
Player 1	*l*	$(2 + t_1, 1)$	(0,0)
	r	(0,0)	$(1, 2 + t_2)$

Coordination game

t_1, t_2 are independent uniform r.v.'s on $[0, x]$.

Player i learns t_i, but not t_{-i}.

Player 2

Player 1		L	R
	l	$(2 + t_1, 1)$	(0,0)
	r	(0,0)	$(1, 2 + t_2)$

Coordination game

Types: t_1, t_2

Beliefs:

Types are independent, so
player i believes t_{-i} is uniform[0, x]

Strategies:

Strategy of player i is a function $s_i(t_i)$

Coordination game

We'll search for a specific form of Bayesian equilibrium:

Assume each player i has a *threshold* c_i, such that the strategy of player 1 is:

$s_1(t_1)$ = *l* if $t_1 > c_1$; = *r* if $t_1 \leq c_1$.

and the strategy of player 2 is:

$s_2(t_2)$ = *R* if $t_2 > c_2$; = *L* if $t_2 \leq c_2$.

Coordination game

Given $s_2(\cdot)$ and t_1, player 1 maximizes expected payoff:

If player 1 plays *l*:

E [Π_1(*l*, $s_2(t_2)$; t_1) | t_1] =

$(2 + t_1) \cdot P(t_2 \leq c_2) + 0 \cdot P(t_2 > c_2)$

If player 1 plays *r*:

E [Π_1(*r*, $s_2(t_2)$; t_1) | t_1] =

$0 \cdot P(t_2 \leq c_2) + 1 \cdot P(t_2 > c_2)$

Coordination game

Given $s_2(\cdot)$ and t_1, player 1 maximizes expected payoff:

If player 1 plays *l*:

E [Π_1(*l*, $s_2(t_2)$; t_1) | t_1] =

(2 + t_1)(c_2/x)

If player 1 plays *r*:

E [Π_1(*r*, $s_2(t_2)$; t_1) | t_1] =

1 - c_2/x

Coordination game

So player 1 should play *l* if and only if:

$t_1 > x / c_2 - 3$

Similarly:

Player 2 should play *R* if and only if:

$t_2 > x / c_1 - 3$

Coordination game

So player 1 should play *l* if and only if:

$t_1 > x / c_2 - 3$

Similarly:

Player 2 should play *R* if and only if:

$t_2 > x / c_1 - 3$

The right hand sides must be the thresholds!

Coordination game: equilibrium

Solve: $c_1 = x / c_2 - 3$, $c_2 = x / c_1 - 3$

Solution:

$$c_i = \frac{\sqrt{9 + 4x} - 3}{2}$$

Thus *one* Bayesian equilibrium
is to play strategies $s_1(\cdot)$, $s_2(\cdot)$,
with thresholds c_1, c_2 (respectively).

Coordination game: equilibrium

What is the *unconditional* probability that player 1 plays *r* ?

= $c_1/x \to 1/3$ as $x \to 0$

Thus as $x \to 0$, the unconditional distribution of play matches the *mixed strategy NE of the complete information game*.

Purification

This phenomenon is one example of *purification*:

recovering a mixed strategy NE via Bayesian equilibrium of a perturbed game.

Harsanyi showed mixed strategy NE can "almost always" be purified in this way.

Summary

To find Bayesian equilibria,
provide strategies $s_1(\cdot)$, ..., $s_N(\cdot)$ where:

For each type θ_i, player i chooses $s_i(\theta_i)$ to maximize expected payoff given his belief P(θ_{-i} | θ_i).

MS&E 246: Lecture 13
Auctions: Imperfect information

Ramesh Johari

Auctions: Theory

- Basic definitions
- Revelation principle
- Truthtelling lemma
- Payoff equivalence theorem
- Revenue equivalence theorem
- Symmetric BNE
- Examples next lecture

A basic auction model

- Assume two players want the same item

- Type of player i : valuation $v_i \geq 0$

  Assume: $P(v_i \leq x_i) = \Phi_i(x_i)$

  F_i : continuous dist. on $[0, V]$, with pdf ϕ_i

  e.g. uniform: $\phi_i(x_i) = 1/V$, for $x_i \in [0, V]$

A basic auction model

Payoffs depend on *winning* and *payment*

- Let w = i if player i wins
- Let p_i = payment of player i
- Payoff to player i of type v_i:

$$\Pi_i(w, p_i; v_i) = \begin{cases} v_i - p_i, & \text{if } w = i \\ 0, & \text{otherwise} \end{cases}$$

A basic auction model

An *auction mechanism* is:

- action set for each player, B_i
- mapping from actions to:

winner: $w(b_1, b_2) \in \{ 1, 2 \}$

payments: $p_i(b_1, b_2) \in [0, \infty), \quad i = 1, 2$

Payoff to i: $Q_i(b_1, b_2 ; v_i) =$

$$\Pi_i(w(\mathbf{b}), p_i(\mathbf{b}) ; v_i)$$

Example: Second price auction

- Action space (bids): B_i = [0, ∞), i = 1, 2

- Winner: $w(b_1, b_2)$ = 1 if $b_1 > b_2$;
  = 2 if $b_1 \leq b_2$

- Payments: $p_i(b_1, b_2)$ = b_{-i} if $w(b_1, b_2) = i$;
  = 0 otherwise

Bayes-Nash equilibrium

Strategy of i : s_i : $[0,\infty) \to B_i$

s_1 is a *Bayesian best response* to s_2 if:

$$\int_0^\infty Q_1(s_1(v_1), s_2(v_2); v_1)\ \phi_2(v_2)\ dv_2 \geq \int_0^\infty Q_1(b_1, s_2(v_2); v_1)\ \phi_2(v_2)\ dv_2$$

for all $b_1 \in B_1$, and $v_1 \geq 0$

(similar definition for player 2)

Bayes-Nash equilibrium

Strategy of i : s_i : $[0,\infty) \to B_i$

s_1 is a *Bayesian best response* to s_2 if:

$$\mathbf{E}[Q_1(s_1(v_1), s_2(v_2); v_1)|v_1] \geq \mathbf{E}_{v_2}[Q_1(b_1, s_2(v_2); v_1)|v_1]$$

for all $b_1 \in B_1$, and $v_1 \geq 0$

(Similarly for player 2)

Bayes-Nash equilibrium

(s_1, s_2) is a BNE if:

- s_1 is a Bayesian best response to s_2
- s_2 is a Bayesian best response to s_1

Second price auction and BNE

We start by finding a BNE for the second price auction.

Recall: Given type v_i, *truthtelling* is a *weak dominant action* for i:

$$d_i(v_i) = v_i$$

Dominant actions and BNE

Consider *any* Bayesian game with type spaces T_1, T_2.

Suppose for each type t_i, player i has a *(weakly) dominant action* $d_i(t_i)$:

$$Q_i(d_i(t_i),\ a_{-i}\ ;\ t_i) \geq Q_i(a_i,\ a_{-i}\ ;\ t_i)$$

for any other action a_i

Dominant actions and BNE

Then $(d_1(\cdot),\ d_2(\cdot))$ is a BNE.

We know that for each player i:

$$Q_i(d_i(t_i),\ d_{-i}(t_{-i})\ ;\ t_i)$$
$$\geq Q_i(a_i,\ d_{-i}(t_{-i})\ ;\ t_i)$$

for all types t_i, t_{-i}, and actions a_i.

Dominant actions and BNE

Then $(d_1(\cdot), d_2(\cdot))$ is a BNE.

Take expectations:

$$E[\ Q_i(d_i(t_i), d_{-i}(t_{-i})\ ;\ t_i)\ |\ t_i\ ]$$
$$\geq E[\ Q_i(a_i, d_{-i}(t_{-i})\ ;\ t_i)\ |\ t_i\ ]$$

for all types t_i, and actions a_i.

This is exactly the condition for a BNE.

Second price auction and BNE

Conclusion:

In the second price auction, *truthtelling is a BNE :*

$$s_i(v_i) = d_i(v_i) = v_i$$

(Note that this requires a dominant action for *every* possible type!)

Incentive compatibility

Auctions where *truthtelling is a BNE*, i.e., where:

1. $B_i = [0, \infty)$ for $i = 1, 2$, and
2. $s_i(v_i) = v_i$ for $i = 1, 2$ is a BNE

are called *incentive compatible.*

Revelation principle

The *revelation principle* shows how to create an incentive compatible auction from any auction with a BNE.

The revelation principle

Given: B_1, B_2, $w(\cdot)$, $p_1(\cdot)$, $p_2(\cdot)$

and a BNE $s_1(\cdot)$, $s_2(\cdot)$

Create a new auction with:

$\underline{B}_1 = \underline{B}_2 = [0, \infty)$

$\underline{w}(\underline{b}_1, \underline{b}_2) = w(s_1(\underline{b}_1), s_2(\underline{b}_2))$

$\underline{p}_i(\underline{b}_1, \underline{b}_2) = p_i(s_1(\underline{b}_1), s_2(\underline{b}_2)), \quad i = 1, 2$

The revelation principle

At the BNE of the original auction:

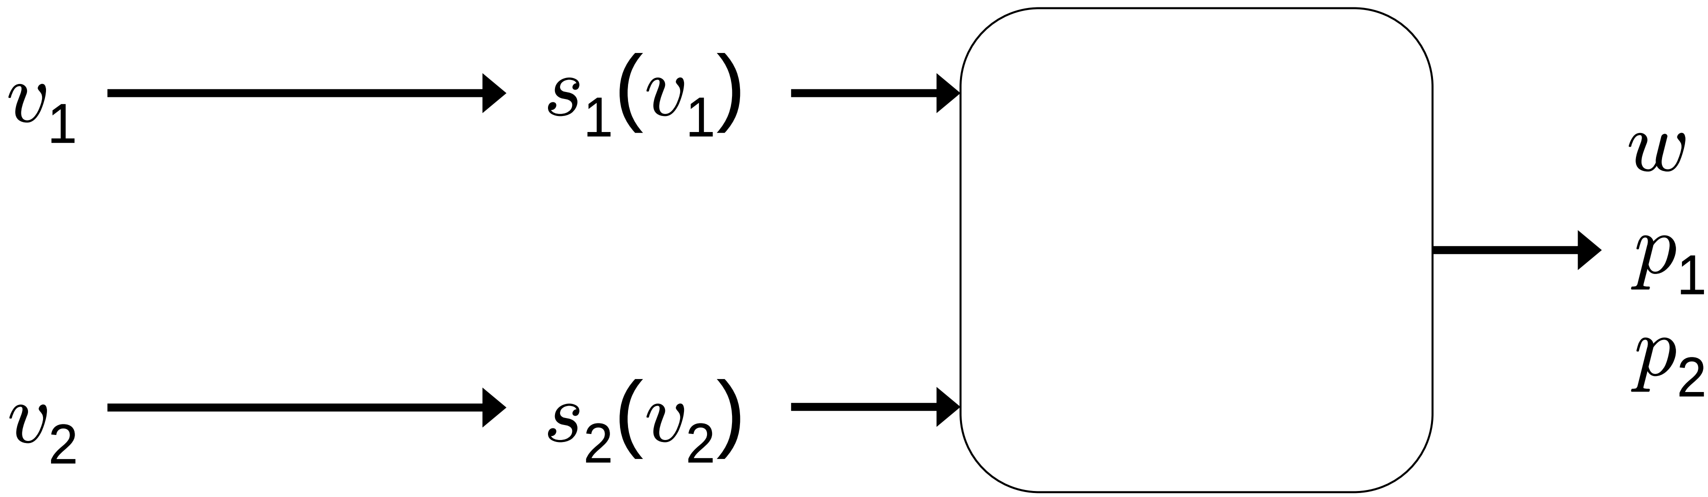

The revelation principle

In the new auction:

Ask players to *declare valuation*.

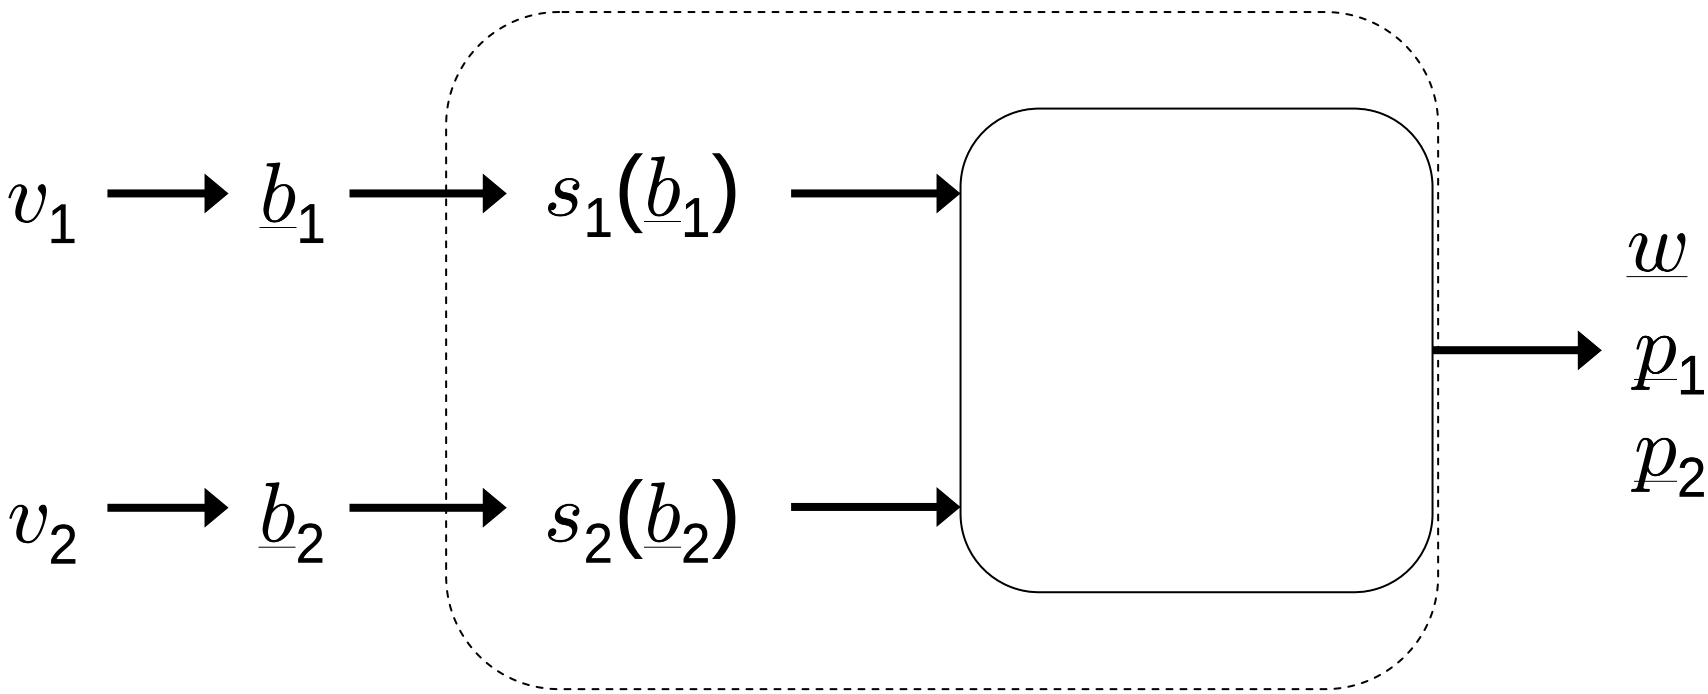

The revelation principle

Theorem:

The new auction is incentive compatible.

Further, the truthtelling strategies
in the new auction give exactly
the same outcomes as
the BNE of the original auction.

The revelation principle: Proof

For all v_1,

$$E[\, Q_1(v_1,\ v_2\ ;\ v_1)\ |\ v_1\ ]$$

$$=\quad E[\, Q_1(s_1(v_1),\ s_2(v_2)\ ;\ v_1)\ |\ v_1\ ]$$

$$\geq\quad E[\, Q_1(b_1,\ s_2(v_2)\ ;\ v_1)\ |\ v_1\ ]$$

for all $b_1 \in B_1$

The revelation principle: Proof

For all v_1,

$$E[\, \underline{Q}_1(v_1, v_2 \,;\, v_1) \mid v_1 \,]$$

$$= \; E[\, Q_1(s_1(v_1), s_2(v_2) \,;\, v_1) \mid v_1 \,]$$

$$\geq \; E[\, Q_1(s_1(\underline{b}_1), s_2(v_2) \,;\, v_1) \mid v_1 \,]$$

for all $\underline{b}_1 \in [0, \infty)$

The revelation principle: Proof

For all v_1,

$E[\, \underline{Q}_1(v_1, v_2 \,;\, v_1) \mid v_1 \,]$

$= \quad E[\, Q_1(s_1(v_1), s_2(v_2) \,;\, v_1) \mid v_1 \,]$

$\geq \quad E[\, \underline{Q}_1(\underline{b}_1, v_2 \,;\, v_1) \mid v_1 \,]$

for all $\underline{b}_1 \in [0, \infty)$

(Similarly for player 2)

The revelation principle: Proof

Given $v_1,\ v_2$:

Outcome at truthtelling strategies

$= (\underline{w}(v_1,\ v_2),\ \underline{p}_1(v_1,\ v_2),\ \underline{p}_2(v_1,\ v_2)\ )$

$= (\ w(s_1(v_1),\ s_2(v_2)),$

$p_1(s_1(v_1),\ s_2(v_2)),$

$p_2(s_1(v_1),\ s_2(v_2))\ )$

= outcome at original BNE

The revelation principle

The new auction is called a

direct revelation mechanism (DRM).

Note:

It may have other, undesirable equilibria!!

Two useful results

For a wide range of auctions (including first and second price),

we will show *payoff equivalence* and *revenue equivalence:*

These auctions all have the *same* payoffs and auctioneer revenue at BNE.

Definitions

Suppose we are given a DRM.

Truthtelling expected payoff to player 1:

$S_1(v_1) = \mathrm{E}[\ Q_1(\ v_1,\ v_2\ ;\ v_1)\ |\ v_1\ ]$

Truthtelling expected probability of winning for player 1:

$P_1(v_1) = \int_0^\infty I\{\ w(v_1,\ v_2) = 1\}\ \phi_2(v_2)\ dv_2$

The truthtelling lemma

Lemma: Truthtelling is a BNE if and only if for i = 1, 2:

(1) $S_i(v_i) = S_i(0) + \int_0^{v_i} \mathrm{P}_i(z)\, \mathrm{d}z$

(2) P_i is nondecreasing:

$$v_i \geq v_i' \Rightarrow \mathrm{P}_i(v_i) \geq \mathrm{P}_i(v_i')$$

The truthtelling lemma: Proof

Truthtelling is a BNE if and only if:

$$S_1(v_1) \geq \mathrm{E}[\ Q_1(\ v_1',\ v_2\ ;\ v_1)\ |\ v_1\ ] \quad \text{for all } v_1' \geq 0$$

(Similarly for player 2)

The truthtelling lemma: Proof

Truthtelling is a BNE if and only if:

$S_1(v_1) \geq$ E[Q_1 (v_1', v_2 ; v_1) | v_1]
for all $v_1' \geq 0$

E[Q_1 (v_1', v_2 ; v_1) | v_1] =

$\int_0^\infty [v_1 - p_1(v_1', v_2)]\, I\{ w(v_1', v_2) = 1\}\, \phi_2(v_2)\, dv_2$

The truthtelling lemma: Proof

Truthtelling is a BNE if and only if:

$S_1(v_1) \geq$ E[Q_1 (v_1', v_2 ; v_1) | v_1]
for all $v_1' \geq 0$

E[Q_1 (v_1', v_2 ; v_1) | v_1] =

$\int_0^\infty [v_1 - p_1(v_1', v_2)]\, I\{ w(v_1', v_2) = 1\}\, \phi_2(v_2)\, dv_2$

The truthtelling lemma: Proof

Truthtelling is a BNE if and only if:

$S_1(v_1) \geq$ E[Q_1 (v_1', v_2 ; v_1) | v_1]

for all $v_1' \geq 0$

E[Q_1 (v_1', v_2 ; v_1) | v_1] =

v_1 P$_1$ (v_1') -

$\int_0^\infty [p_1(v_1', v_2)]\, I\{ w(v_1', v_2) = 1\}\, \phi_2(v_2)\, dv_2$

The truthtelling lemma: Proof

Truthtelling is a BNE if and only if:

$S_1(v_1) \geq$ E[Q_1 (v_1', v_2 ; v_1) | v_1]
for all $v_1' \geq 0$

E[Q_1 (v_1', v_2 ; v_1) | v_1] =

v_1 P$_1$ (v_1') - v_1' P$_1$ (v_1')

$\int_0^\infty [v_1' - p_1(v_1', v_2)]\, I\{ w(v_1', v_2) = 1\}\, \phi_2(v_2)\, dv_2$

The truthtelling lemma: Proof

Truthtelling is a BNE if and only if:

$S_1(v_1) \geq$ E[Q_1 (v_1', v_2 ; v_1) | v_1]
for all v_1' ≥ 0

E[Q_1 (v_1', v_2 ; v_1) | v_1] =

v_1 P_1 (v_1') - v_1' P_1 (v_1') + S_1 (v_1')

The truthtelling lemma: Proof

Conclude:

Truthtelling is a BNE if and only if
for i = 1, 2, and for all v_i , v_i' $\geq$ 0:

$S_i(v_i) \geq S_i(v_i') + P_i(v_i')(v_i - v_i')$

i.e., S_i is *convex*.

The truthtelling lemma: Proof

Assume v_i' > v_i. Then:

$S_i(v_i) \geq S_i(v_i') + P_i(v_i')(v_i - v_i')$

$S_i(v_i') \geq S_i(v_i) + P_i(v_i)(v_i' - v_i)$

$\Rightarrow P_i(v_i)(v_i' - v_i) \leq P_i(v_i')(v_i' - v_i)$

So $P_i(v_i) \leq P_i(v_i') \Rightarrow P_i$ is *nondecreasing*

The truthtelling lemma: Proof

If $v_i > v_i'$:

$$\frac{S_i(v_i) - S_i(v_i')}{v_i - v_i'} \geq \mathsf{P}_i(v_i')$$

If $v_i < v_i'$:

$$\frac{S_i(v_i) - S_i(v_i')}{v_i - v_i'} \leq \mathsf{P}_i(v_i')$$

Take $v_i' \uparrow v_i$, $v_i' \downarrow v_i$ $\Rightarrow$ $S_i'(v_i) = \mathsf{P}_i(v_i)$

The truthtelling lemma

How to use the truthtelling lemma:

(1) Use a BNE of an auction
to create an incentive compatible DRM

(2) Apply the truthtelling lemma
to characterize the original BNE

Payoff equivalence

Given two auctions with BNE such that:

-in each BNE, if $v_i = 0$ then player i gets zero payoff; and

-in each BNE, item *always* goes to highest valuation player

Theorem: Both BNE yield the same *expected payoff* to each player.

Payoff equivalence: Proof

- Fix given BNE (s_1, s_2) of one of the auctions

- Construct incentive compatible DRM using revelation principle

- For this DRM:

  $\underline{S}_i(0) = 0$, and

  $\underline{P}_i(v_i) = \int_0^{v_i} \phi_{-i}(v_{-i})\ dv_{-i}$

Payoff equivalence: Proof

Expected payoff to player 1 of type v_1:

depends only on $S_1(0)$ and $P_1(\cdot)$,

by the truthtelling lemma

(Similarly for player 2)

Payoff equivalence: Proof

Expected payoff to player 1 of type v_1:

$$E[\ Q_1(s_1(v_1),\ s_2(v_2)\ ;\ v_1)\ |\ v_1\ ]$$

$$= E[\ \underline{Q}_1(\ v_1,\ v_2\ ;\ v_1)\ |\ v_1\ ]$$

$$= \underline{S}_1(v_1) = \int_0^{v_1} [\ \int_0^{v_1'}\ \phi_2(v_2)\ dv_2\ ]\ dv_1'$$

by the truthtelling lemma

(Similarly for player 2)

Payoff equivalence: Proof

So at given BNE of *either auction*, expected payoff to player 1 of type v_1 is:

$$\int_0^{v_1} \left[\int_0^{v_1'} \phi_2(v_2) \, dv_2 \right] dv_1'$$

This does not depend on the BNE!

(Similarly for player 2)

Revenue equivalence

Given two auctions with BNE such that:

-in each BNE, if v_i = 0 then player i gets zero payoff; and

-in each BNE, item *always* goes to highest valuation player

Theorem: Both BNE yield the same *expected revenue* to the auctioneer.

Revenue equivalence: Proof

Fix BNE (s_1, s_2) of one of the auctions

Note:

$$\sum^2_{i=1} Q_i\,(s_1(v_1),\ s_2(v_2)\ ;\ v_i)$$

$$= v_{w(s_1(v_1),\ s_2(v_2))} - p_{w(s_1(v_1),\ s_2(v_2))}$$

Revenue equivalence: Proof

Taking expectations:

Sum of expected payoffs to players

= $\mathbf{E}$ [max $\{v_1, v_2\}$] -

Expected revenue to auctioneer

Revenue equivalence: Proof

Taking expected values:

Expected revenue to auctioneer

= E[max $\{v_1, v_2\}$] -

Sum of expected payoffs to players

Revenue equivalence: Proof

Taking expected values:

Expected revenue to auctioneer

= E[max $\{v_1, v_2\}$] -

Sum of expected payoffs to players

Right hand side is same for BNE of both auctions (by payoff equivalence)

Revenue equivalence

Note that at BNE,
expected payoffs to players are ≥ 0.

So:
Revenue to auctioneer $\leq$ E[max $\{v_1, v_2\}$]

[Aside: optimal auction theory]

The problem of maximizing the equilibrium revenue to the auctioneer is called *optimal auction design.*

For this problem to be well-defined, an additional constraint is needed, *individual rationality*:

E[$Q_i(s_i(v_i), \mathbf{s}_{\text{-}i}(\mathbf{v}_{\text{-}i}) \; ; \; v_i) \mid v_i$] $\geq$ 0 for all i.

(Otherwise bidder i would not participate.)

[Aside: optimal auction theory]

The framework defined here can be used to characterize the optimal auction design for any distribution of players' valuations.

(See Myerson 1979)

Symmetric BNE

From now on, assume:

- $B_1 = B_2 = [0, \infty)$
- $\phi_1 = \phi_2 = \phi$ *(same distribution)*
  (Assume ϕ is positive on its entire domain)

A BNE is *symmetric* if:

$s_1(v) = s_2(v)$ for all $v \geq 0$

$s(v) = s_i(v)$ is called the *bid function*.

Symmetric BNE theorem

Theorem:

If highest bidder wins,

then in a symmetric BNE with bid function s,

s is strictly increasing,

so the winning bidder also has *the highest valuation*.

(In case of tie, assume player 2 wins)

Symmetric BNE: Proof

Apply revelation principle to build new incentive compatible auction:

$\underline{w}(\underline{b}_1, \underline{b}_2) = w(s(\underline{b}_1), s(\underline{b}_2))$

$\underline{p}_i(\underline{b}_1, \underline{b}_2) = p_i(s(\underline{b}_1), s(\underline{b}_2)), \quad i = 1, 2$

In this auction:

$\underline{\mathrm{P}}_1(v_1) = \int_0^\infty I\{ \underline{w}(v_1, v_2) = 1\} \, \phi(v_2) \, \mathrm{d}v_2$

Symmetric BNE: Proof

Apply revelation principle to build new incentive compatible auction:

$\underline{w}(\underline{b}_1, \underline{b}_2) = w(s(\underline{b}_1), s(\underline{b}_2))$

$\underline{p}_i(\underline{b}_1, \underline{b}_2) = p_i(s(\underline{b}_1), s(\underline{b}_2)), \quad i = 1, 2$

In this auction:

$\underline{P}_1(v_1) = \int_0^\infty I\{ w(s(v_1), s(v_2)) = 1\} \phi(v_2) \, dv_2$

Symmetric BNE: Proof

Apply revelation principle to build new incentive compatible auction:

$\underline{w}(\underline{b}_1, \underline{b}_2) = w(s(\underline{b}_1), s(\underline{b}_2))$

$\underline{p}_i(\underline{b}_1, \underline{b}_2) = p_i(s(\underline{b}_1), s(\underline{b}_2)), \quad i = 1, 2$

In this auction:

$\underline{\mathrm{P}}_1(v_1) = \int_0^\infty I\{\, s(v_1) > s(v_2) \,\}\, \phi(v_2)\, \mathrm{d}v_2$

Symmetric BNE: Proof

By truthtelling lemma,

$$\int_0^\infty I\{\ s(v_1) > s(v_2)\ \}\ \phi(v_2)\ \mathrm{d}v_2$$

is nondecreasing in v_1.

Only possible if $s(v)$ is nondecreasing in v.

We only need to show s is *strictly increasing.*

Symmetric BNE: Proof

We will show s is *strictly increasing* in the special case of the *first price auction:*

$p_i(b_1, b_2) = b_i$ if $w(b_1, b_2) = i$;

$= 0$ otherwise

However, the result holds more generally for the other auctions we consider.

Symmetric BNE: Proof

Suppose s is not strictly increasing:

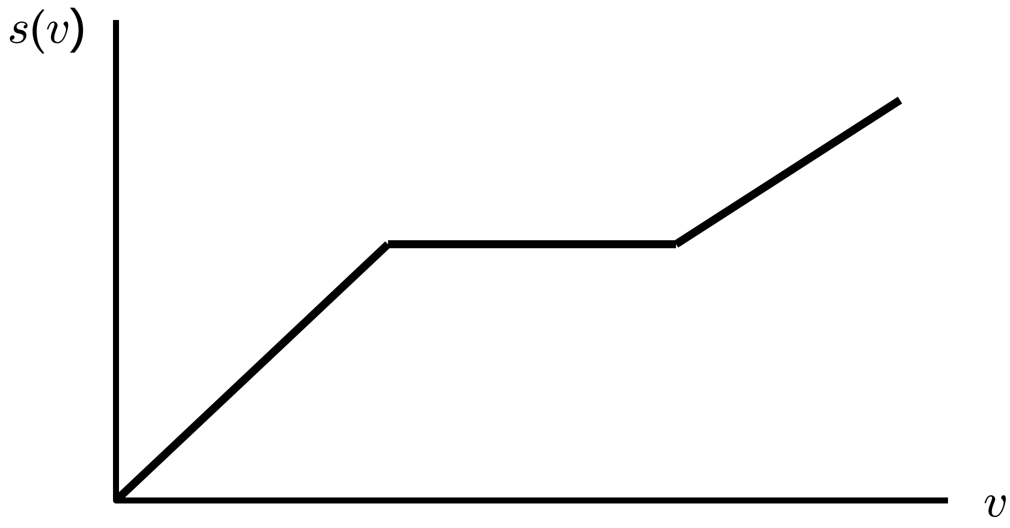

Symmetric BNE: Proof

Suppose s is not strictly increasing:

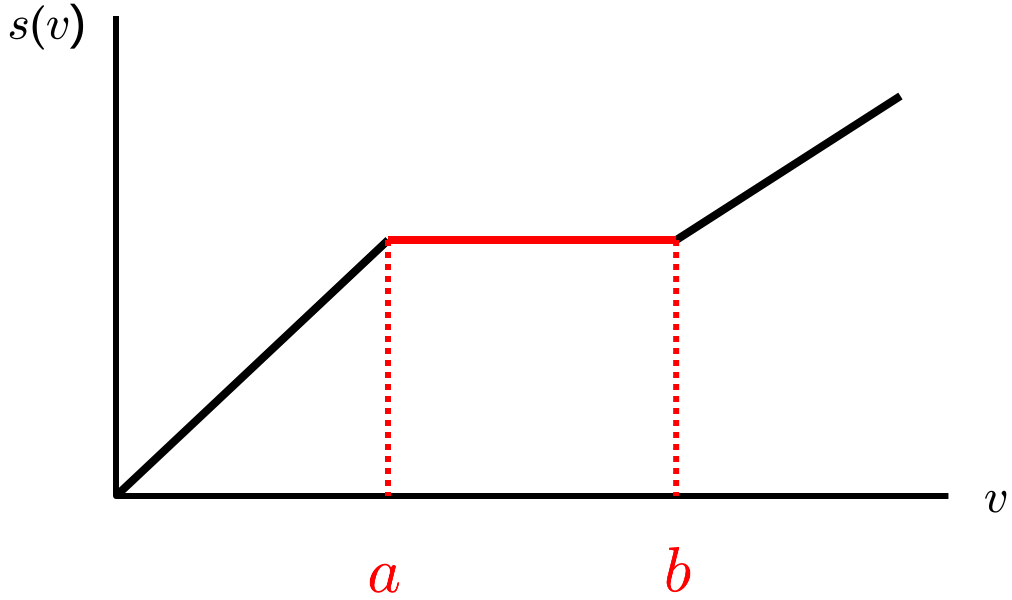

Symmetric BNE: Proof

Suppose s is not strictly increasing.

Fix a < b such that:

$$s(v) = s(a), \quad a \leq v \leq b$$

We can assume: a > $s(a)$.

(If not, just increase a slightly.)

Symmetric BNE: Proof

Given player 2 is using $s_2 = s$,

suppose player 1 bids

$$b_1 = s(a) + \varepsilon \text{ when } v_1 = a.$$

Then when $v_1 = a$:

- Expected *payment* by player 1 increases by at most ε
- Player 1 wins if $v_2 \in [a, b]$

Symmetric BNE: Proof

Given player 2 is using $s_2 = s$,

suppose player 1 bids

$$b_1 = s(a) + \varepsilon \text{ when } v_1 = a.$$

Player 1's change in expected payoff

$$\geq (a - s(a) - \varepsilon)(F(b) - F(a)) - \varepsilon$$

$$> 0 \quad \text{for small enough } \varepsilon$$

Profitable deviation!

Symmetric BNE: Proof

Conclude:

- s is strictly increasing

So:

- Winner must have highest valuation

Symmetric BNE

In general, can show the same result if:

(1) $p_i(b_1, b_2) \geq 0$ for all b_1, b_2;

(2) the winner's payment is positive when at least one of b_1, b_2 is positive; and

(3) $p_1(b_1, b_2) = p_2(b_2, b_1)$ for all b_1, b_2 *(permutation invariance)*

Moral

Symmetric BNE of “standard auctions”

(first price, second price, etc.)

have the *same* expected payoffs and auctioneer revenue.

In particular,

Expected revenue = E[second highest bid]

(Why?)

MS&E 246: Lecture 14
Auctions: Examples

Ramesh Johari

Example: Second price auction

- Action space (bids): $B_i = [0, \infty)$, $i = 1, 2$
- Winner: $w(b_1, b_2) = 1$ if $b_1 > b_2$;
  $= 2$ if $b_1 \leq b_2$
- Payments: $p_i(b_1, b_2) = b_{-i}$ if $w(b_1, b_2) = i$;
  $= 0$ otherwise
- Assume: $\phi_1 = \phi_2 = \phi$, continuous, positive everywhere on its domain, with distribution F

Example: Second price auction

Since $d_i(v_i) = v_i$ is a dominant action for each player,

$s_i(v_i) = v_i$ for $i = 1, 2$ is a BNE.

It is a *symmetric* and *truthtelling* BNE.

So the second price auction is *incentive compatible.*

Example: Second price auction

The truthtelling lemma tells us:

$S_1(v_1) = S_1(0) + \int_0^{v_1} \mathrm{P}_1(z)\, \mathrm{d}z$

But: $S_1(0) = 0$

$\mathrm{P}_1(v_1) = \mathrm{P}(v_1 > v_2 \mid v_1) = \int_0^{v_1} \phi(v_2)\, \mathrm{d}v_2$

So $S_1(v_1) = \int_0^{v_1} \left[\int_0^{z} \phi(v_2)\, \mathrm{d}v_2 \right] \mathrm{d}z$

(Similarly for player 2)

Example: Second price auction

The truthtelling lemma tells us:

$S_1(v_1) = S_1(0) + \int_0^{v_1} \mathrm{P}_1(z)\, \mathrm{d}z$

But: $S_1(0) = 0$

$\mathrm{P}_1(v_1) = \mathrm{P}(v_1 > v_2 \mid v_1) = \int_0^{v_1} \phi(v_2)\, \mathrm{d}v_2$

So $S_1(v_1) = \int_0^{v_1} F(z)\, \mathrm{d}z$

(Similarly for player 2)

Example: Second price auction

Expected payoff to player 1, given type v_1

$= S_1(v_1) = \int_0^{v_1} F(z)\,\mathrm{d}z$

(Similarly for player 2)

Example: Second price auction

Are there any *other* symmetric BNE of the second price auction?

Suppose $s_1 = s_2 = s$ is such a BNE.

Example: Second price auction

By symmetric BNE theorem,

s is strictly increasing, and

player 1 wins if and only if $v_1 > v_2$

Expected payoff to player 1 of type v_1:

$\underline{S}_1(v_1) = \int_0^{v_1} (v_1 - s(v_2))\, \phi(v_2)\, dv_2$

Example: Second price auction

Observe that:

- $\underline{S}_i(0) = 0$ for $i = 1, 2$
- Highest valuation player always wins

So by payoff equivalence theorem:

$$S_1(v_1) = \underline{S}_1(v_1)$$

Example: Second price auction

By payoff equivalence:

$$S_1(v_1) = \underline{S}_1(v_1)$$

Example: Second price auction

By payoff equivalence:

$$\int_0^{v_1} F(z)\, \mathrm{d}z = \int_0^{v_1} (v_1 - s(v_2))\, \phi(v_2)\, \mathrm{d}v_2$$

Example: Second price auction

By payoff equivalence:

$$\int_0^{v_1} F(z)\,\mathrm{d}z = v_1 F(v_1) - \int_0^{v_1} s(v_2))\,\phi(v_2)\,\mathrm{d}v_2$$

Differentiate:

$$F(v_1) = F(v_1) + v_1\,\phi(v_1) - s(v_1)\,\phi(v_1)$$

Example: Second price auction

By payoff equivalence:

$$\int_0^{v_1} F(z)\,\mathrm{d}z = v_1 F(v_1) - \int_0^{v_1} s(v_2))\,\phi(v_2)\,\mathrm{d}v_2$$

Differentiate:

$$0 = v_1\,\phi(v_1) - s(v_1)\,\phi(v_1)$$

Example: Second price auction

By payoff equivalence:

$$\int_0^{v_1} F(z)\,\mathrm{d}z = v_1 F(v_1) - \int_0^{v_1} s(v_2))\,\phi(v_2)\,\mathrm{d}v_2$$

Differentiate:

$$0 = (\ v_1 - s(v_1)\ )\ \phi(v_1)$$

So: $v_1 = s(v_1)$

Example: Second price auction

Conclude: truthtelling is unique symmetric BNE.

Expected revenue to auctioneer:
$\mathbf{E}$[second highest valuation]

Example: First price auction

- Action space (bids): $B_i = [0, \infty)$, $i = 1, 2$
- Winner: $w(b_1, b_2) = 1$ if $b_1 > b_2$;
  $= 2$ if $b_1 \leq b_2$
- Payments: $p_i(b_1, b_2) = b_i$ if $w(b_1, b_2) = i$;
  $= 0$ otherwise
- Assume: $\phi_1 = \phi_2 = \phi$, continuous, positive everywhere on its domain, with distribution F

Example: First price auction

What are the symmetric BNE of the first price auction?

Suppose $s_1 = s_2 = s$ is a symmetric BNE.

Example: First price auction

By symmetric BNE theorem,

s is strictly increasing, and

player 1 wins if and only if $v_1 > v_2$

Expected payoff to player 1 of type v_1:

$$S_1^{FP}(v_1) = \int_0^{v_1} (v_1 - s(v_1))\, \phi(v_2)\, dv_2$$

$$= (v_1 - s(v_1))\, F(v_1)$$

Example: First price auction

Observe that:

- $S_i(0) = 0$ for $i = 1, 2$
- Highest valuation player always wins

So by payoff equivalence theorem,

$S_1^{FP}(v_1)$ = expected payoff to type v_1 player in second price auction

Example: First price auction

By payoff equivalence:

$$S_1^{\mathrm{FP}}(v_1) = \int_0^{v_1} F(z)\, \mathrm{d}z$$

Example: First price auction

By payoff equivalence:

$(v_1 - s(v_1))\, F(v_1) = \int_0^{v_1} F(z)\, dz$

So:

$$s(v) = v - \frac{\int_0^v F(z)dz}{F(v)}$$

(e.g., when Φ is uniform: $s(v) = v/2$)

Example: First price auction

$$s(v) = v - \frac{\int_0^v F(z)dz}{F(v)}$$

Observe that $s(v) < v$.

This practice is called *bid shading*.

Example: First price auction

Revenue equivalence also holds, so:

Expected revenue to auctioneer =

expected revenue under second price auction =

$\mathbf{E}$[second highest valuation]

< $\mathbf{E}$[max{ v_1, v_2 }]

Revenue

The shortfall between $\mathbf{E}[\max\{v_1, v_2\}]$ and expected revenue is called an *information rent:*

At an equilibrium the buyers must make a profit if they reveal their private valuation.

Revenue

However, this relies on *independent private valuations.*

If valuations are correlated, the auctioneer can get expected revenue = $\mathbf{E}[\max\{v_1, v_2\}]$

See Problem Set 6.

(Theorem: Cremer and McLean, 1985)

MS&E 246: Lecture 15
Perfect Bayesian equilibrium

Ramesh Johari

Dynamic games

In this lecture, we begin a study of *dynamic games of incomplete information.*

We will develop an analog of Bayesian equilibrium for this setting, called *perfect Bayesian equilibrium*.

Why do we need beliefs?

Recall in our study of subgame perfection that problems can occur if there are “not enough subgames” to rule out equilibria.

Entry example

- Two firms
- First firm decides if/how to enter
- Second firm can choose to "fight"

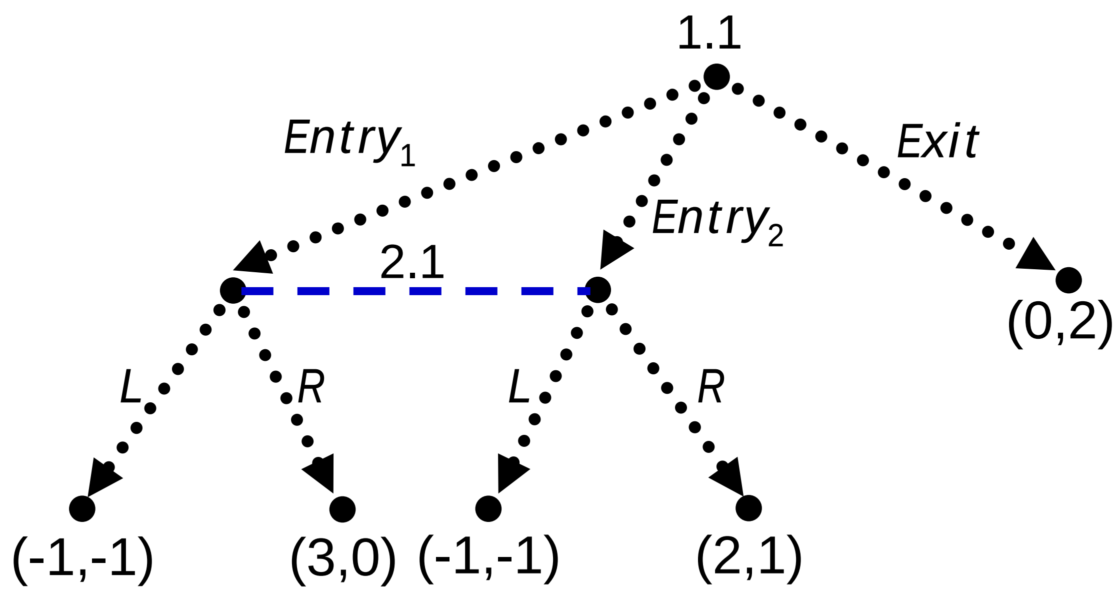

Entry example

Note that this game only has *one* subgame.

Thus SPNE are *any* NE of strategic form.

		Firm 2	
		L	*R*
	$Entry_1$	(-1,-1)	(3,0)
Firm 1	$Entry_2$	(-1,-1)	(2,1)
	Exit	(0,2)	(0,2)

Entry example

Two pure NE of strategic form:

($Entry_1$, R) and ($Exit$, L)

Firm 2

Firm 1		L	R
	$Entry_1$	(-1,-1)	(3,0)
	$Entry_2$	(-1,-1)	(2,1)
	$Exit$	(0,2)	(0,2)

Entry example

But firm 1 should *"know"* that if it chooses to enter, firm 2 will never "fight."

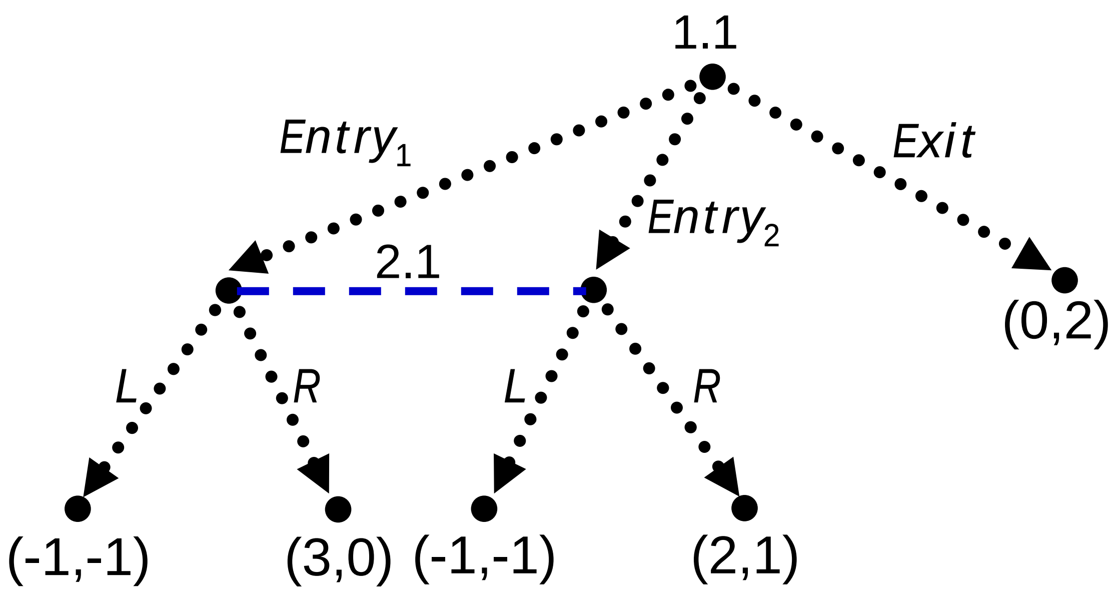

Entry example

So in this situation,
there are too many SPNE.

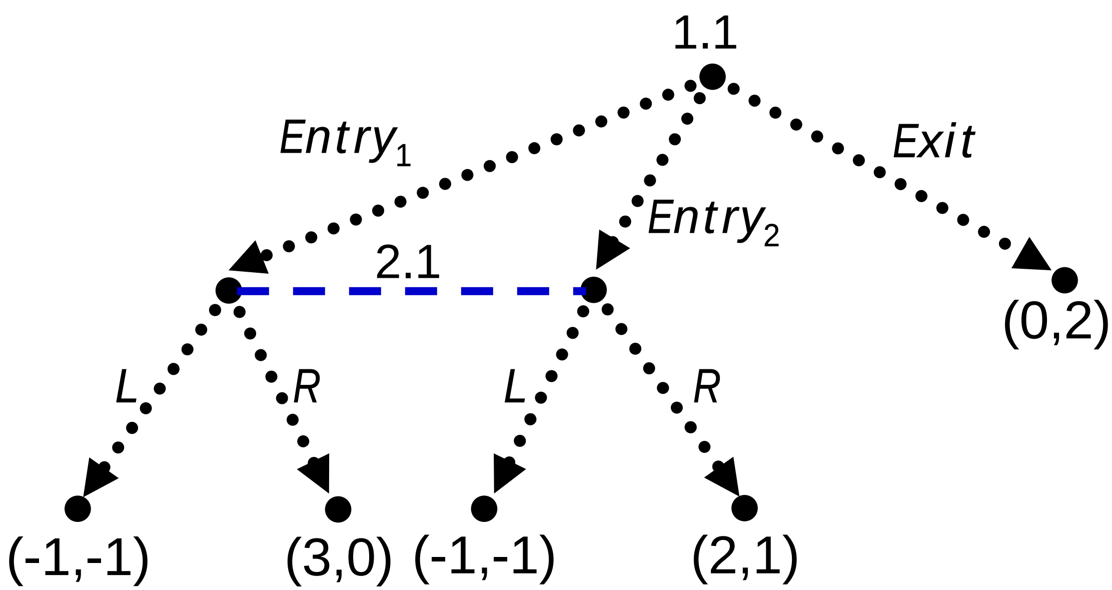

Beliefs

A solution to the problem of the entry game is to include *beliefs* as part of the solution concept:

Firm 2 should never fight, regardless of what it believes firm 1 played.

Beliefs

In general, the beliefs of player i are:

a conditional distribution over everything player i *does not know*,

given everything that player i *does know.*

Beliefs

In general, the beliefs of player i are:

a conditional distribution over the nodes of the information set i is in,

given player i is at that information set.

(When player i is in information set h, denoted by $P_i(v \mid h)$, for $v \in h$)

Beliefs

One example of beliefs:

In static Bayesian games, player i's *belief* is $P(\theta_{-i} \mid \theta_i)$ (where θ_j is type of player j).

But *types* and *information sets* are in 1-to-1 correspondence in Bayesian games, so this matches the new definition.

Perfect Bayesian equilibrium

Perfect Bayesian equilibrium (PBE) strengthens subgame perfection by requiring two elements:

- a complete strategy for each player i (mapping from info. sets to mixed actions)
- beliefs for each player i ($P_i(v \mid h)$ for all information sets h of player i)

Entry example

In our entry example, firm 1 has only one information set, containing one node.

His belief just puts probability 1 on this node.

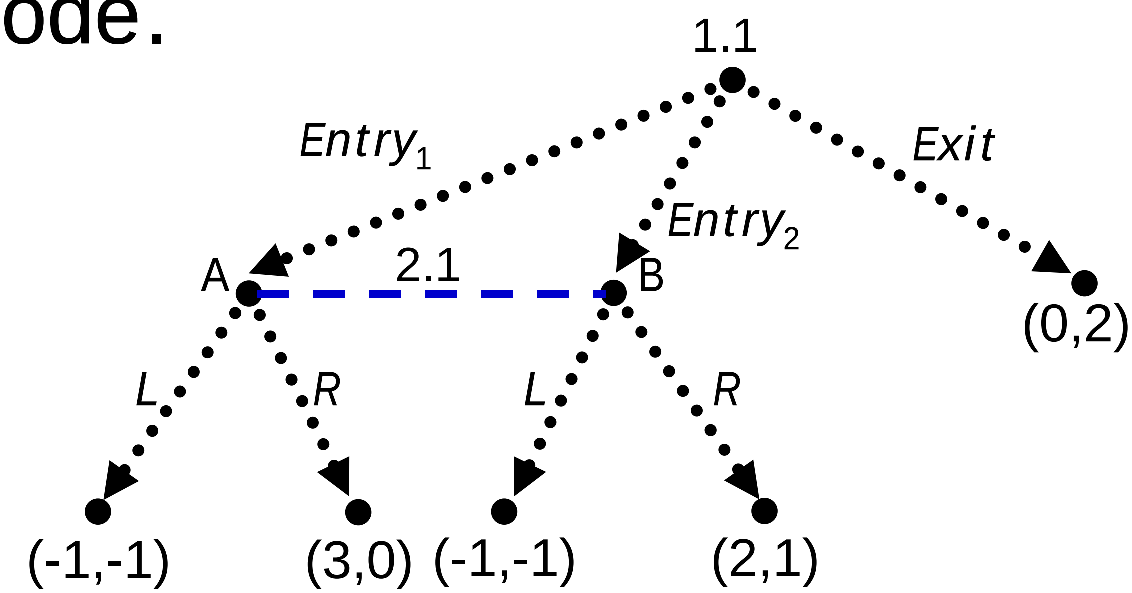

Entry example

Suppose firm 1 plays a mixed action with probabilities $(p_{\text{Entry}_1}, p_{\text{Entry}_2}, p_{\text{Exit}})$, with $p_{\text{Exit}} < 1$.

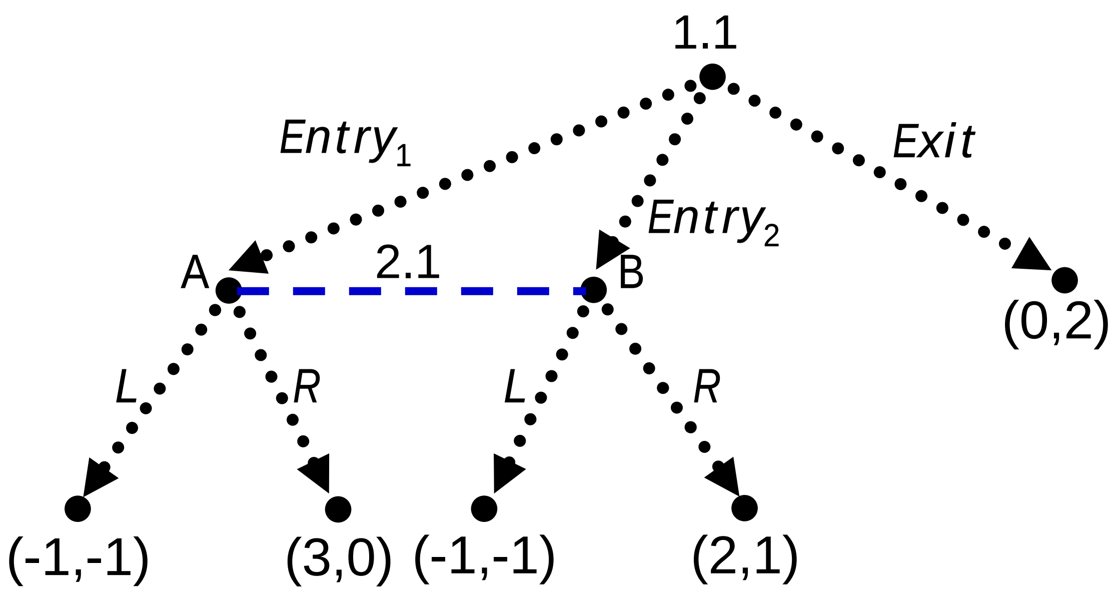

Entry example

What are firm 2's *beliefs* in 2.1?

Computed using Bayes' Rule!

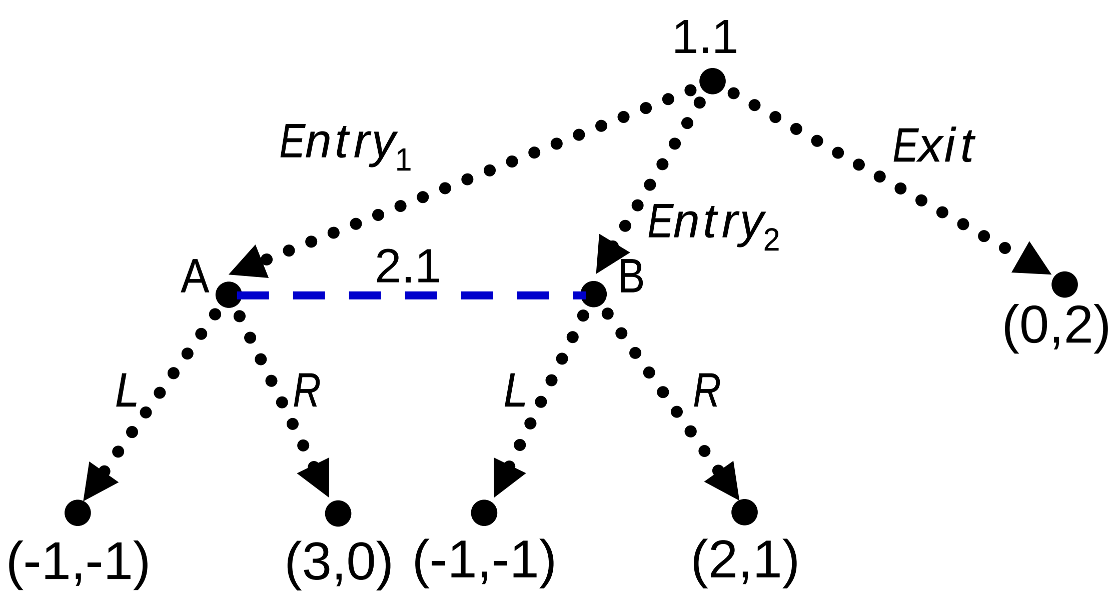

Entry example

What are firm 2's *beliefs* in 2.1?

$P_2(A \mid 2.1) = p_{\text{Entry}_1} / (p_{\text{Entry}_1} + p_{\text{Entry}_2})$

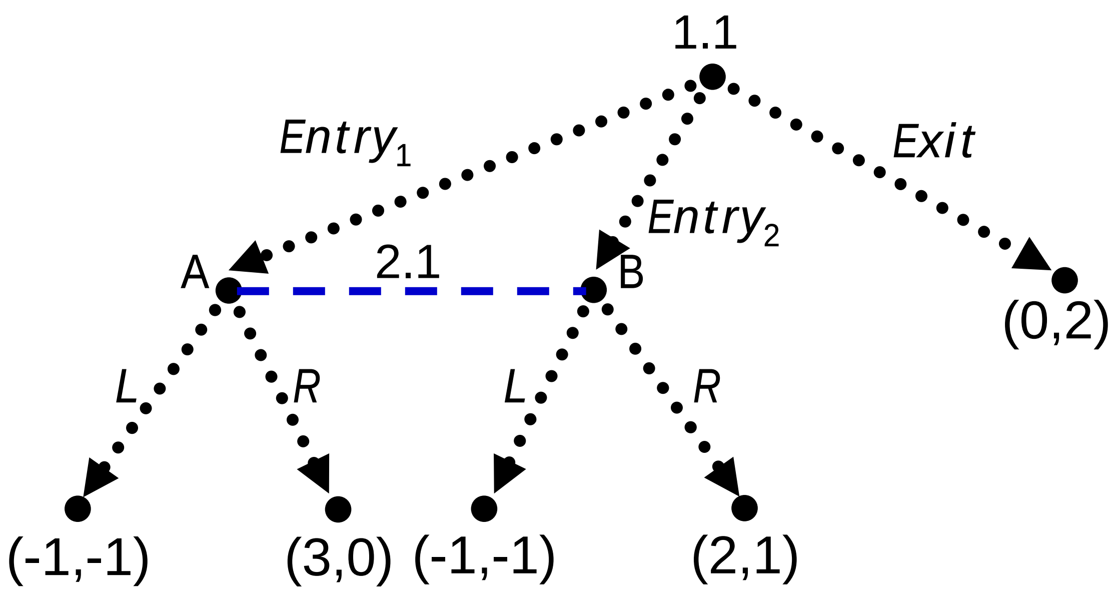

Entry example

What are firm 2's *beliefs* in 2.1?

$P_2(B \mid 2.1) = p_{\text{Entry}_2} / (p_{\text{Entry}_1} + p_{\text{Entry}_2})$

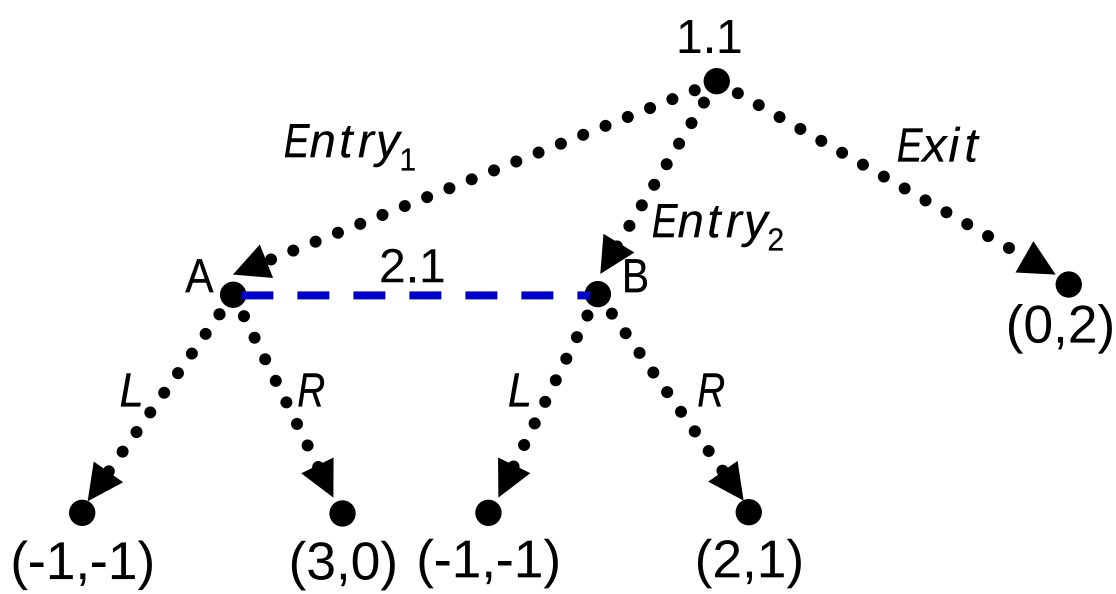

Beliefs

In a perfect Bayesian equilibrium,

"*wherever possible*",

beliefs must be computed using Bayes' rule and the strategies of the players.

(At the very least, this ensures information sets that can be reached with positive probability have beliefs assigned using Bayes' rule.)

Rationality

How do player's choose strategies?

As always, they do so to
maximize payoff.

Formally:
Player i's strategy $s_i(\cdot)$ is such that
in any information set h of player i,
$s_i(h)$ maximizes player i's expected payoff,
given his beliefs and others' strategies.

Entry example

For *any* beliefs player 2 has in 2.1,

he maximizes expected payoff by playing *R.*

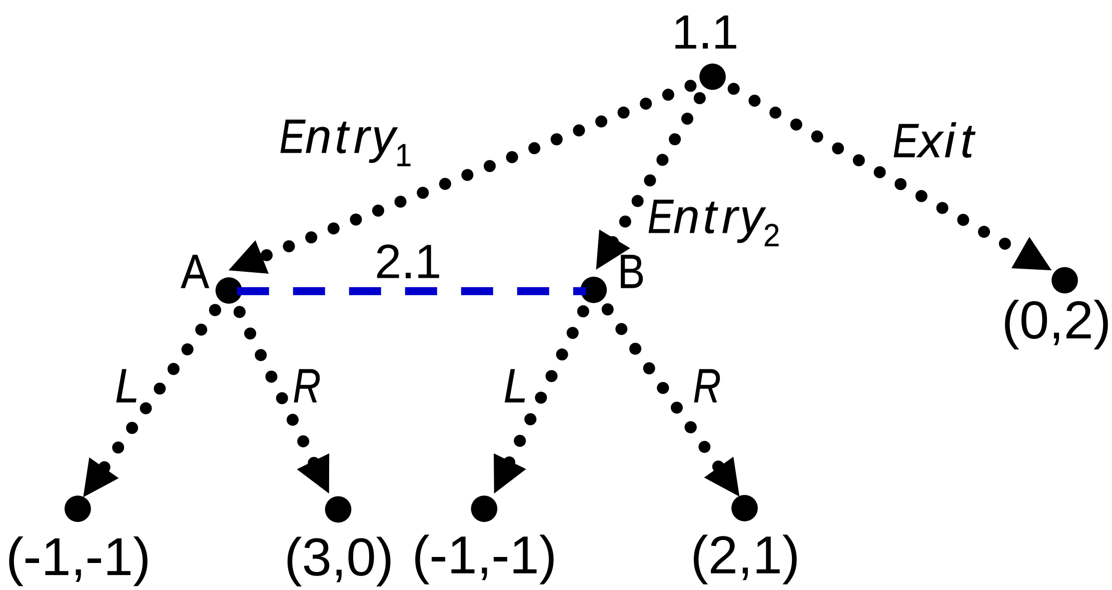

Entry example

Thus, in any PBE, player 2 must play *R* in 2.1.

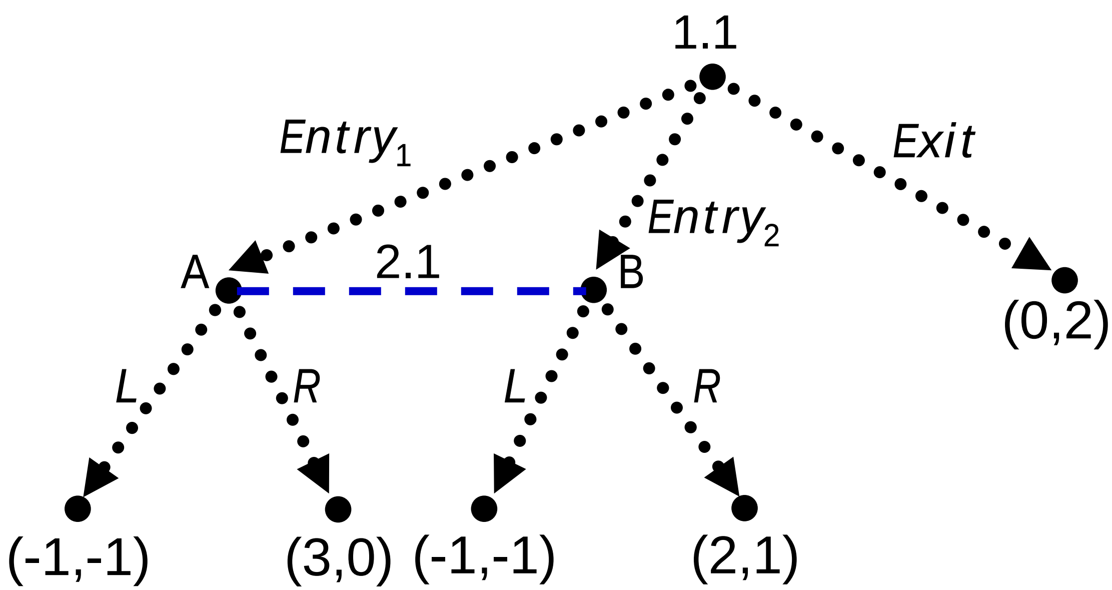

Entry example

Thus, in any PBE, player 2 must play *R* in 2.1.

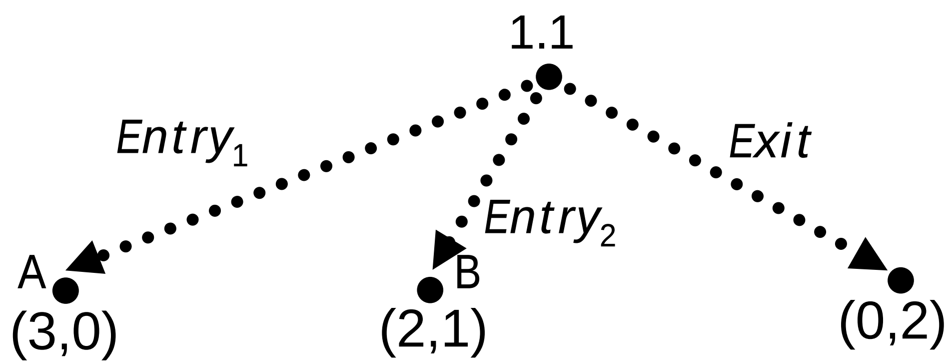

Entry example

So in a PBE, player 1 will play $Entry_1$ in 1.1.

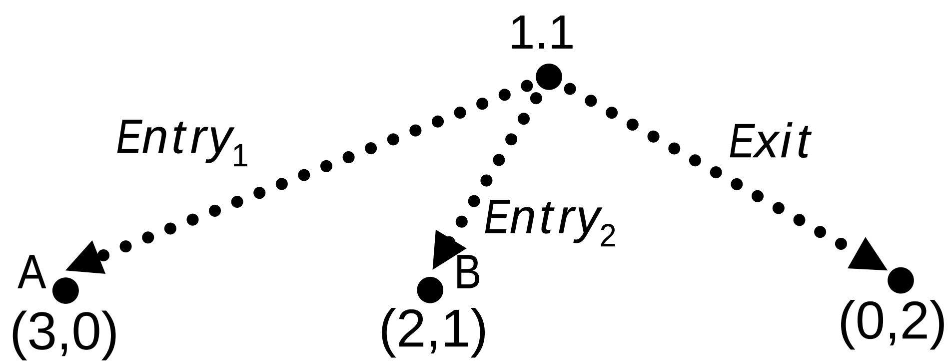

Entry example

Conclusion: unique PBE is ($Entry\_1$, R).

We have eliminated the NE ($Exit$, L).

PBE vs. SPNE

Note that a PBE is *equivalent to* SPNE for dynamic games of complete and perfect information:

All information sets are singletons, so beliefs are trivial.

In general, PBE is *stronger* than SPNE for dynamic games of complete and imperfect information.

Summary

- Beliefs: conditional distribution at every information set of a player
- Perfect Bayesian equilibrium:
  1. Beliefs computed using Bayes' rule and strategies (when possible)
  2. Actions maximize expected payoff, given beliefs and strategies

An ante game

- Let t_1, t_2 be uniform[0,1], independent.
- Player i observes t_i;
  each player puts $1 in the pot.
- Player 1 can force a "showdown",
  or player 1 can "raise" (and add $1 to the pot).
- In case of a showdown, both players show t_i;
  the highest t_i wins the entire pot.
- In case of a raise, Player 2 can "fold" (so player 1 wins) or "match" (and add $1 to the pot).
- If Player 2 matches, there is a showdown.

An ante game

To find the perfect Bayesian equilibria of this game:

Must provide strategies $s_1(\cdot)$, $s_2(\cdot)$; and beliefs $P_1(\cdot \mid \cdot)$, $P_2(\cdot \mid \cdot)$.

An ante game

Information sets of player 1:

t_1 : His *type*.

Information sets of player 2:

(t_2, a_1) : type t_2,
and action a_1 played by player 1.

An ante game

Represent the beliefs by *densities*.

Beliefs of player 1:

$p_1(t_2 \mid t_1) = t_2$ (as types are independent)

Beliefs of player 2:

$p_2(t_1 \mid t_2, a_1)$ = density of player 1's type, *conditional* on having played a_1
= $p_2(t_1 \mid a_1)$ (as types are indep.)

An ante game

Using this representation,
can you find a perfect Bayesian equilibrium of the game?

MS&E 246: Lecture 16
Signaling games

Ramesh Johari

Signaling games

Signaling games are two-stage games where:

- Player 1 (with private information) moves first.
  His move is observed by Player 2.
- Player 2 (with no knowledge of Player 1's private information) moves second.
- Then payoffs are realized.

Dynamic games

Signaling games are a key example of *dynamic games of incomplete information.*

(i.e., a dynamic game where the entire structure is *not* common knowledge)

Signaling games

The formal description:

Stage 0:

Nature chooses a random variable t_1,
observable only to Player 1,
from a distribution P(t_1).

Signaling games

The formal description:

Stage 1:

Player 1 chooses an *action* a_1 from the set A_1.

Player 2 observes this choice of action.

(The action of Player 1 is also called a "message.")

Signaling games

The formal description:

Stage 2:

Player 2 chooses an action a_2 from the set A_2.

Following Stage 2, payoffs are realized:

$\Pi_1(a_1, a_2 ; t_1)$; $\Pi_2(a_1, a_2 ; t_1)$.

Signaling games

Observations:

- The modeling approach follows Harsanyi's method for static Bayesian games.
- Note that Player 2's payoff depends on the type of player 1!
- When Player 2 moves first, and Player 1 moves second, it is called a *screening* game.

Application 1: Labor markets

A key application due to Spence (1973):

Player 1: worker

t_1 : intrinsic ability

a_1 : education decision

Player 2: firm(s)

a_2 : wage offered

Payoffs: Π_1 = net benefit

Π_2 = productivity

Application 2: Online auctions

Player 1: seller

t_1 : true quality of the good

a_1 : advertised quality

Player 2: buyer(s)

a_2 : bid offered

Payoffs: Π_1 = profit

Π_2 = net benefit

Application 3: Contracting

A model of Cachon and Lariviere (2001):

Player 1:	manufacturer
t_1 :	demand forecast
a_1 :	declared demand forecast, contract offer
Player 2:	supplier
a_2 :	capacity built
Payoffs:	Π_1 = profit of manufacturer Π_2 = profit of supplier

A simple signaling game

Suppose there are two types for Player 1, and two actions for each player:

- t_1 = *H* or t_1 = *L*

  Let p = P(t_1 = *H*)
- A_1 = { *a*, *b* }
- A_2 = { *A*, *B* }

A simple signaling game

Nature moves first:

Nature

A simple signaling game

Nature moves first:

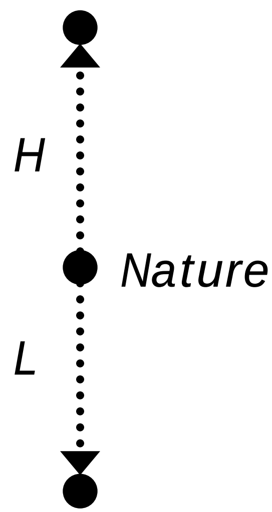

A simple signaling game

Player 1 moves second:

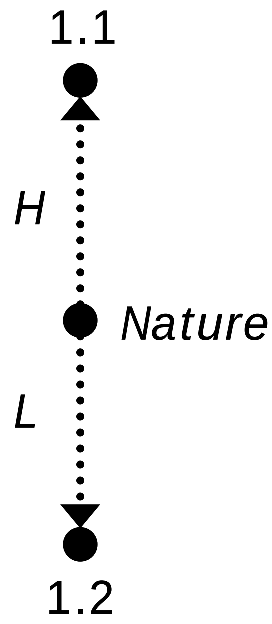

A simple signaling game

Player 1 moves second:

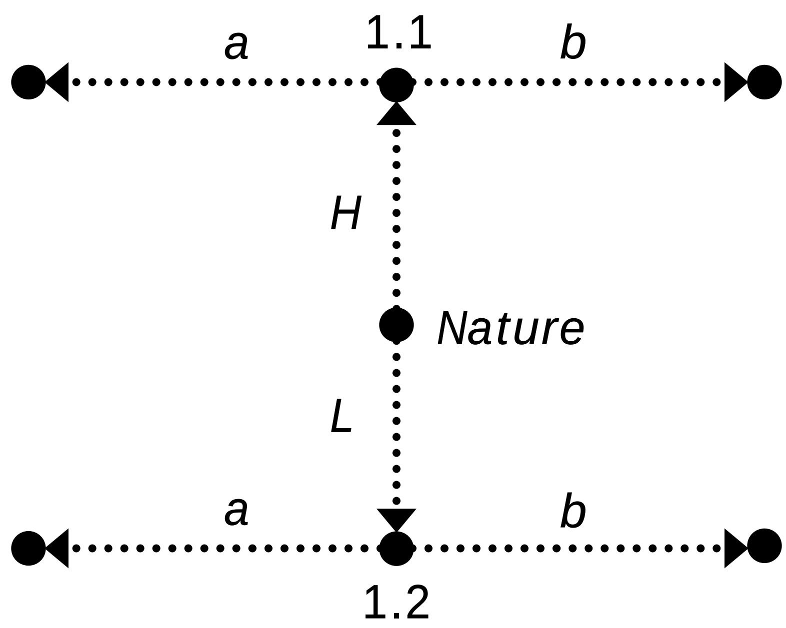

A simple signaling game

Player 2 observes Player 1's action:

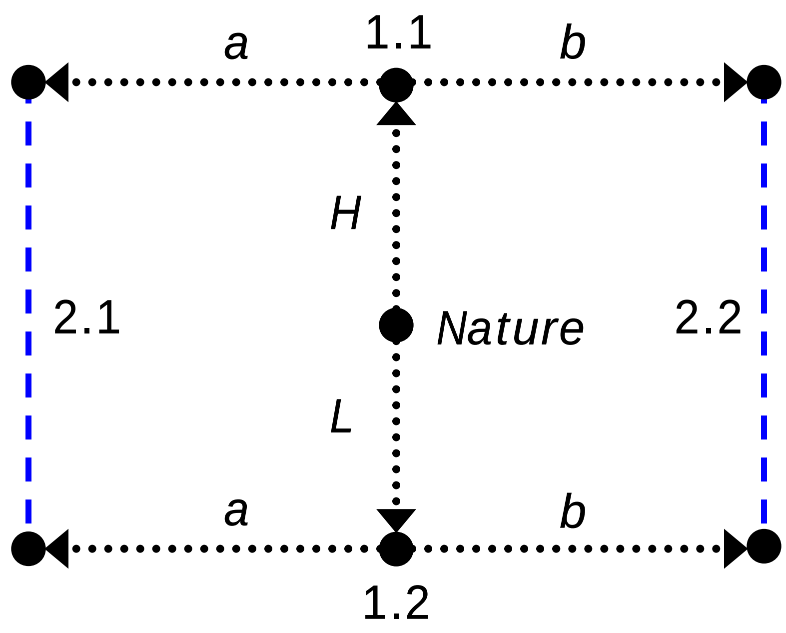

A simple signaling game

Player 2 moves:

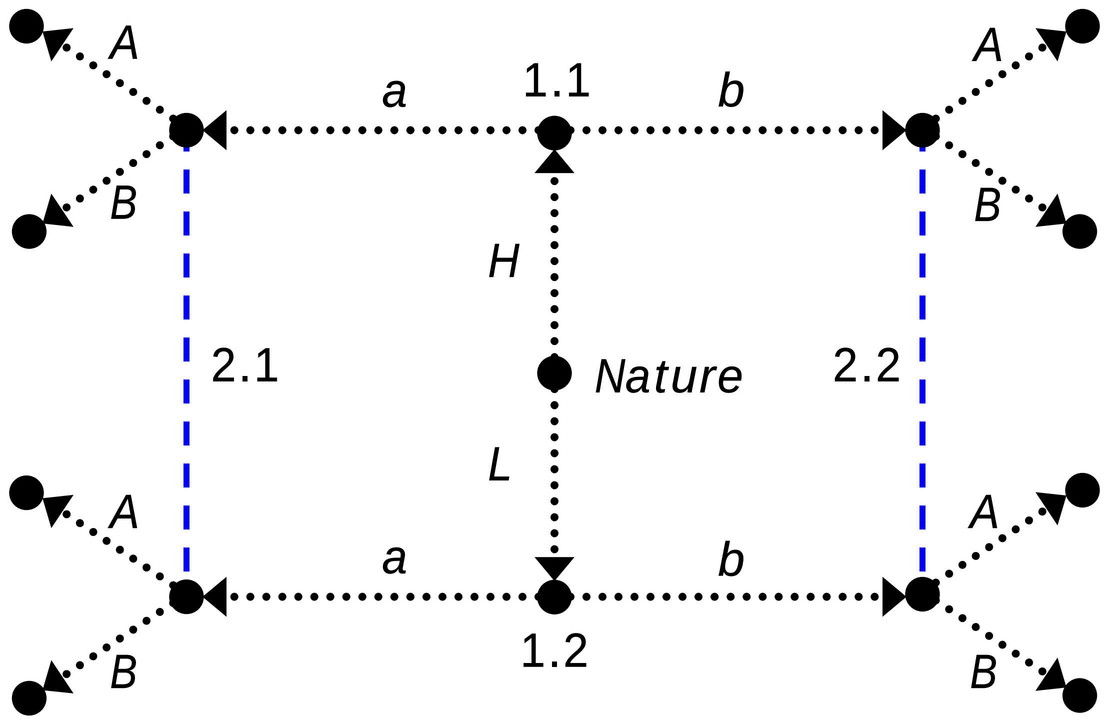

A simple signaling game

Payoffs are realized: $\Pi_i(a_1, a_2; t_1)$

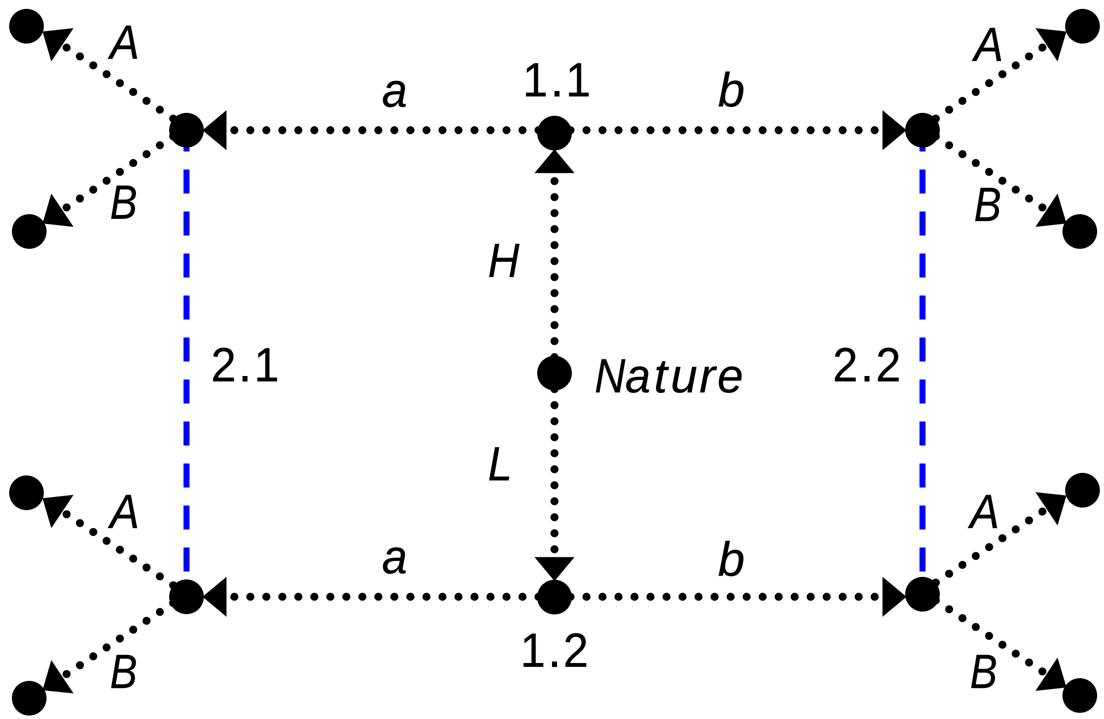

Perfect Bayesian equilibrium

Each player has 2 information sets, and 2 actions in each, so 4 strategies.

A PBE is a pair of strategies and beliefs such that:

-each players' beliefs are derived from strategies using Bayes' rule (if possible)

-each players' strategies maximize expected payoff given beliefs

Pooling vs. separating equilibria

When player 1 plays the *same* action, regardless of his type, it is called a *pooling strategy.*

When player 1 plays *different* actions, depending on his type, it is called a *separating strategy.*

Pooling vs. separating equilibria

In a *pooling equilibrium*,

Player 2 gains no information about t_1 from Player 1's message

$\Rightarrow P_2(t_1 = H \mid a_1) = P(t_1 = H) = p$

In a *separating equilibrium,*

Player 2 knows Player 1's type *exactly* from Player 1's message

$\Rightarrow P_2(t_1 = H \mid a_1) = 0$ or 1

An eBay-like model

- Suppose seller has an item with quality either *H* (prob. p) or *L* (prob. 1 - p).
- Seller can advertise either *H* or *L*.
- Assume there are two bidders.
- Suppose that bidders always bid truthfully, given their beliefs.
  - (This would be the case if the seller used a second price auction.)

An eBay-like model

Suppose seller *always* advertises *high.*

Then: buyers will never “trust” the seller, and always bid expected valuation.

This is the pooling equilibrium:

$s_1(H) = s_1(L) = H$.

$s_B(H) = s_B(L) = p\, H + (1 - p)\, L$.

An eBay-like model

Is there any equilibrium where $s_1(t_1) = t_1$?

(In this case the seller is *truthful*.)

In this case the buyers bid:

$s_B(H) = H$, $s_B(L) = L$.

But if the buyers use this strategy,
the seller prefers to *always* advertise H!

An eBay-like model

Now suppose that if the seller *lies* when the true value is L,
there is a cost c (in the form of lower reputation in future transactions).

If $H - c < L$,
then the seller prefers to tell the truth
$\Rightarrow$ separating equilibrium.

An eBay-like model

This example highlights the importance of *signaling costs*:

To achieve a separating equilibrium, there must be a difference in the costs of different messages.

(When there is no cost, the resulting message is called "cheap talk.")