Game Theory

Catherine

Prerequisites:
Game Theory from Micro

Stategic-Form Game (2x2 Nash Equilibrium)
Extensive-Form Game, Subgam-Perfect Equilibrium
Cournot Game, Stackelberg Game
Repeated Game, Bertrand Game

Strategic-Form Games

Game Theory: multi-agent optimization problems

Cournot (1838): quantity-setting game in oligopolistic competition
John von Neuman (1928): a 'minimax' equilibrium in zero-sum games
John Nash (1950): Nash equilibrium
Reinhard Selten (1965): perfection in dynamic games
John Harsanyi (1967): Bayesian equilibrium under incomplete info

A Game in Strategic Form

- players: maximize own payoff knowing others do the same
- possible strategies for each player $\xrightarrow{\text{simultaeneously}}$ strategy profile
- payoff for each player at every outcome

1. strictly dominant strategy
2. iterated strict dominance (ISD)
   iterated elimination of strictly dominated strategies (IESDS)
   (*) iterated weak dominance
   (the order of deletion may affect prediction)

3. Nash equilibrium: mutual best-replies, a self-fulfilling prophecy

Example 3: Discrete Cournot
- best-reply payoffs underlined
- unique survivor of ISD is (M,M)
- unique Nash is (M,M)

	H	M	L
H	4, 4	8, $\underline{6}$	$\underline{16}$, 4
M	$\underline{6}$, 8	$\underline{9}$, $\underline{9}$	15, 5
L	4, $\underline{16}$	5, 15	7, 7

(*) mixed strategy Nash equilibrium { self-indifferent / make others indifferent

Nash's Theorem (1950): Every game with finitely many strategies has at least one Nash equilibrium in pure or mixed strategies.

fun fact: if players have strict preferences over the outcomes then the number of Nash equilibria is odd! (but finding all Nash equilibria is computationally difficult)

	H	T
H	1, -1	-1, 1
T	-1, 1	1, -1

If Row plays each strategy with 50% chance then the mathematical expected value of Column's payoff is $0.5 \cdot (-1) + 0.5 \cdot 1 = 0$ from H, $0.5 \cdot 1 + 0.5 \cdot (-1) = 0$ from T, Column is indifferent, may mix
If Column mixes 50-50 then Row is indifferent, may mix as assumed
$\Rightarrow$ $(0.5 \cdot H + 0.5 \cdot T,\ 0.5 \cdot H + 0.5 \cdot T)$ is mixed strategy Nash equilibrium

	B	S
B	5, 3	2, 2
S	0, 0	3, 5

mixed Nash:
$(\frac{5}{6}B + \frac{1}{6}S,\ \frac{1}{6}B + \frac{5}{6}S)$
$\Rightarrow$ payoffs $= (\frac{15}{6}, \frac{15}{6})$

2 pure Nash equilibriums + 1 Nash equilibrium in mixed strategies:
if Ann plays B with probability p , S with probability $(1 - p)$ then Bob's payoff from B is $3p$, from S it is $2p + 5(1 - p)$ Bob is indifferent iff $3p = 2p + 5 - 5p \Leftrightarrow p = 5/6$

if Bob plays B with prob q and S with $(1 - q)$, then Ann is indifferent iff $5q + 2(1 - q) = 3(1 - q)$ equivalently $q = 1/6$

Nash in 2x2 Games

Example: Hawk-Dove Game

	L $(1-r)$	R (r)
(t) **T**	-3, -3	8, 0
$(1-t)$ **B**	0, 8	2, 2

best reply $\Longrightarrow$

Row's

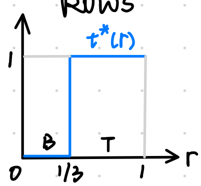

Column's

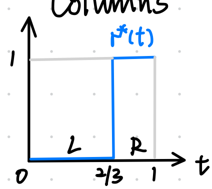

Row's expected payoff: $\begin{cases} u_1(T) = -3(1-r) + 8r \\ u_1(B) = 2r \end{cases}$

Column's expected payoffs: $\begin{cases} u_2(L) = -3t + 8(1-t) \\ u_2(R) = 2(1-t) \end{cases}$

$\Rightarrow$ three Nash equilibria

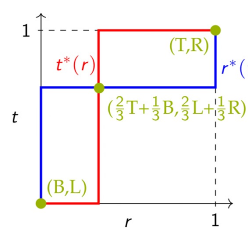

	L	R
T	0, -3	8, 0
B	0, 8	2, 2

best reply $\Longrightarrow$

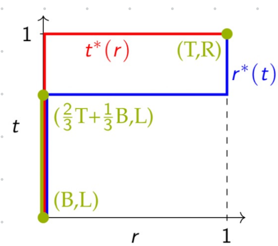

(T, R) and
$\{(tT + (1-t)B, L) \mid 0 \le t \le \frac{2}{3}\}$

Modified Example

Dynamic Games

extensive form : tree + info sets
strategic form : x timing of moves

From strategic form to Extensive Form

- a directed tree graph (nodes, directed edges, no cycles)
- a player assigned to each nodes, action corresponding to edges
- payoff written at terminal nodes
- "who knows what" indicated via information sets

Backward Induction ⇒ Subgame-Perfect Equilibrium (SPE)

Firm 1 (entrant) decides whether to enter (action E) or not (action N)

Observing this, firm 2 (incumbent) decides whether to fight (action F) or acquiesce (A) 默许

Payoffs displayed at terminal nodes in the order (firm 1, firm 2)

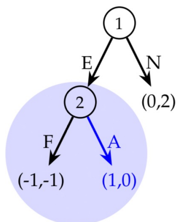

optimal
continuation play

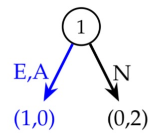

F: sequentially irrational
(not credible)

1. Backward Induction can be used in every sequential-moves game with perfect information (= players observing all earlier moves)
2. Subgame-Perfect Equilibrium (SPE)
   ⇒ is Nash Equilibrium (SPNE)
3. SPE induces Nash equilibrium in every continuation (subgame) of the perfect-information, sequential-moves game.

	F	A
N	$\underline{0}$, $\underline{2}$	0, $\underline{2}$
E	-1, -1	$\underline{1}$, $\underline{0}$

Firm 1 (first payoff) chooses row

Best-reply payoffs underlined

Two pure Nash: (N,F) and (E,A)

Mixed: (N,pF+$(1-p)$A) with $p \geq \frac{1}{2}$

(use the "recipe" to find them all)

Among all Nash equilibria **only (E,A) is subgame-perfect**

SPE rules out Nash equilibria sustained by non-credible threats
⇒ SPNE (via backward induction) is a more robust prediction than Nash equilibrium in dynamic games

4. simultaneous moves in the extensive form

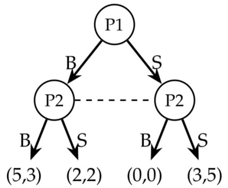

Battle of the Sexes

	B	S
B	5, 3	2, 2
S	0, 0	3, 5

- P1 wants to go to beach (B)
- P2 wants to go to store (S)
- both wish to meet

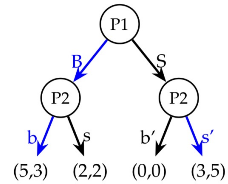

(sequential BoS, P_1 first)

Dashed line between P2's decision nodes indicates P2 cannot directly observe which node he is at; these nodes form an **information set**

- the same action must be taken at all nodes within an information set because the player cannot observe which node he or she is at

→ imperfect information

Imperfect Competition in an Industry

Cournot (1838) Game

$\begin{cases} \text{firm 1: produce } q_1 \\ \text{firm 2: produce } q_2 \end{cases}$, price $p = 1 - q_1 - q_2$

$\Rightarrow$ $\left.\begin{array}{l} \text{firm 1: } \max\ (1-q_1-q_2)q_1 \text{ given } q_2 : \ q_1^*(q_2) = \frac{1-q_2}{2} \\ \text{firm 2: } \max\ (1-q_1-q_2)q_2 \text{ given } q_1 : \ q_2^*(q_1) = \frac{1-q_1}{2} \end{array}\right\} q_1^* = q_2^* = \frac{1}{3}$

—— Cournot-Nash equilibrium

Stackelberg Game (Sequential quantity setting)

firm 1 picks q_1 first (first-mover)

$\Rightarrow \max\ (1-q_1-q_2)q_1 = (1-q_1-\frac{1-q_1}{2})q_1 \Rightarrow q_1^* = \frac{1}{2},\ q_2^* = \frac{1}{4}$

—— The first mover is strictly better off in quantity-setting duopoly

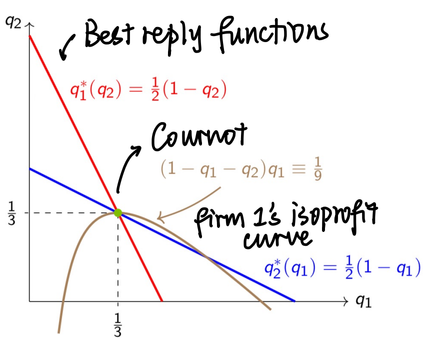

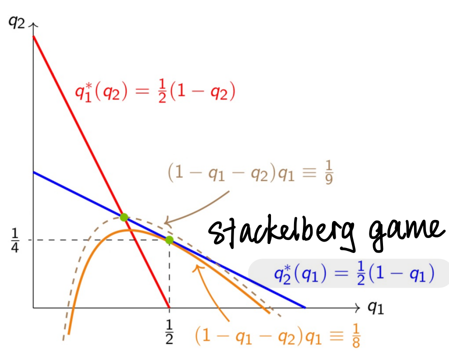

Repeated Games

Five Pirates Puzzle

Five pirates, A, B, C, D and E, find 100 gold coins. They take turns, in alphabetical order, to propose a division of the loot.
All pirates that have not been eliminated have to vote for or against the current proposal. If the proposal gets weak majority then the coins are distributed accordingly and the game ends. If it is voted down (strictly) then the proposer is eliminated (has to "walk the plank"), and the next pirate makes a proposal. Each pirate's main goal is not to be eliminated; conditional on that, he or she maximizes the number of gold coins received. All else equal, each pirate prefers to eliminate as many other pirates as possible.

Solution:

Start at the end of the game when only {D, E} remain: D offers 0 to E, E votes against but D votes for the proposal, hence {100, 0} passes

Knowing this, when {C, D, E} remain, it is sufficient for C to offer 1 coin to E to get E's vote, hence C offers {99, 0, 1} which is accepted

When {B, C, D, E} remain B just needs one vote: offers 1 coin to D (who would get 0 in the continuation) and {99, 0, 1, 0} is accepted

At the start of the game A needs two additional votes (C and E are cheapest): hence **A offers {98, 0, 1, 0, 1}, it passes, game ends**

Centipede (Bob Rosenthal, 1981)

players 1 and 2 have £1 each, player 1 gets to move first
a player can either Stop (ending the game) or Pass (decreasing her own payoff by £1, increasing the other's by £2, and letting the other move next), each can Pass a fixed number of times (say, 3)

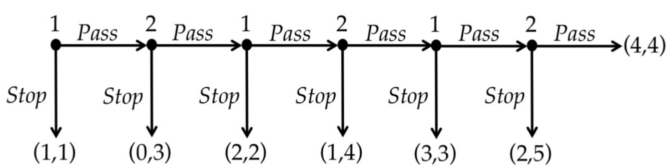

Finite Repetition

Prisoner's Dilemma

Prisoner's Dilemma **played exactly once**

Pareto-dominates

	C	D
C	1, 1	-1, 2
D	2, -1	0, 0

strict Nash equilibrium

- D is strictly dominant for each player
- Hence (D,D) is the unique (and strict) Nash equilibrium of the stage game
- (C,C) strictly Pareto-dominates it but cannot be attained in equilibrium

⇒ finitely repeated prisoner's dilemma

unique SPE: play D in every period

Evidence by Robert Axelrod (political scientist, 1981):

First tournament: each program paired with all others to play 200 iterations of the PD; the highest total cumulative score wins

The winner was "TIT FOR TAT" submitted by Anatol Rapoport (psychologist, 1911-2007), which played C in the first game against a new opponent, then played the opponent's most recent move

Axelrod published the results, then called another tournament; the winner was again "TIT FOR TAT"

Axelrod's conclusions: be nice but provocable, fair, and not tricky

Hawk-Dove

Hawk-Dove

	H	D
H	-1, -1	3, 0
D	0, 3	2, 2

One-shot Hawk-Dove (not PD!)

- (H,D) and (D,H) both Nash
- (D,D) yields higher total payoff but unattainable for rational players

Play it exactly $T+2$ times and consider the following "social norm"

- each should play D in all periods $t = 1, \dots, T$ as long as nobody has deviated and then (H, D) and (D, H) in the final two periods
- if anyone deviates the other plays H in all subsequent periods and the deviator (rationally) plays D against that

This is SPE in the finitely-repeated Hawk-Dove game!

Infinite Repetition

Theorem: in infinitely-repeated Prisoner's Dilemma, the outcome "(C, C) forever" can be sustained in SPE for δ sufficiently close to 1.

Proof:

- Consider the strategy "play C in the first period then C as long as the outcome has been (C,C), otherwise switch to D forever"
- If anyone deviates from C then mutual "D forever" is of course Nash equilibrium in the continuation (subgame) for any δ ✓
- Check deviation from C: playing D pays 2 once then 0 forever, whereas staying on-path pays 1 forever; the latter is better iff $2 < 1 + \delta + \delta^2 + \ldots = \frac{1}{1-\delta} \iff \delta > \frac{1}{2}$, i.e., true for δ close to 1 ✓

grim trigger (Nash reversion)
threatening with one-shot Nash forever

Bertrand Game (tacit collusion in oligopoly)

Assumptions:
1. N firms, a homogeneous product, price competition
2. whichever charges the lowest price takes entir market
3. constant unit cost c
4. monopoly price p^M, monopoly profit $\pi^M > 0$

Equilibrium in one-shot Bertrand game: $p^* = c$ (unique Nash)

—— unlike Cournot's quantity competition model, Bertrand argued that competition should drive price down to the marginal cost even with just two firms

{ Quantity Competition: e.g. airplane manufacturing, strict capacity
{ Price Competition: e.g. copyrighted software / music, sales scaled up

Infinitely-repeated Bertrand Model:

sustain cooperation (p^M) in SPE in the long run $\xrightarrow{\text{threat}}$ revert to Marginal-Cost pricing forever if one deviates

$\Rightarrow \left(\frac{\pi^M}{N}\right) \cdot \frac{1}{1-\delta} \geqslant \pi^M \Leftrightarrow \delta \geqslant \frac{N-1}{N}$

(more firms, more impatient firms, find it harder to collude)

Weeks 1:
Static Games of Complete Info

Formalization of a game
- Dominance
- Rationalizability

Nash Equilibrium (pure, mixed, continuous actions)
Correlated Equilibrium

Formalization of a Game

For a static game of complete information (strategic-form game):

- set of players $N = \{1, 2, \cdots N\}$, $N \in$ finity who play simoultaneously
- set of actions available to each player, A_i, $i \in N$
  - strategy: a complete (contingent) plan which specifies how the player will act in every possible distinguishable circumstance, a_i
  - Mixed Strategy: α_i, a probability distribution over actions $a_i \in A_i$.
- individual payoffs / preferences of each action. $u_i(a_i, a_{-i})$, $a_{-i} \in \prod_{j \neq i} A_j$

  $a_{-i} = (a_1, a_2, \cdots a_{i-1}, a_{i+1}, \cdots a_n)$

matrix form (normal form)
and equivalent extensive form representation

Dominance

strict dominance

expected payoffs

$\forall a_{-i}, \exists \alpha_i', v_i(\alpha_i', a_{-i}) > v_i(a_i, a_{-i})$: strategy a_i strictly dominated

$\forall a_i' \in A_i$ strictly dominated by a_i: strategy a_i strictly dominant

	l (p)	c (1-p)	r
T	$\underline{1}, \underline{6}$	$\underline{2}, 0$	$1, 1$
B	$0, 0$	$1, \underline{6}$	$\underline{3}, 1$

r strictly dominated by any strategy $s = pl + (1-p)c$ with $p \in (1/6, 5/6)$

- No strategy is dominated by another pure stragy.
- However, strategy r looks 'weak'.
- Consider a mixture of two 'strong' strategies: $s = p \times l + (1-p) \times c$.
- If player 1 plays T, then s gives a higher payoff than r if

$$v_2(T, s) = 6p + 0 > v_2(T, r) = 1 \Rightarrow p > 1/6.$$

- If player 1 plays B, then s gives a higher payoff than r if

$$v_2(B, s) = 0 + 6(1-p) > v_2(B, r) = 1 \Rightarrow p < 5/6.$$

Iterative Elimination of Strictly Dominated Strategies (IESDS)

(all based on every players are and are know to be rational)

	a	b	c	d
A	5, 5	4, 0	2, 2	2, 4
B	0, 4	6, 6	4, 5	0, 3
C	6, 2	5, 1	3, 4	0, 1
D	4, 2	3, 0	1, 0	1, 6

$\xrightarrow{A \succ D}$

	a	b	c	d
A	5, 5	4, 0	2, 2	2, 4
B	0, 4	6, 6	4, 5	0, 3
C	6, 2	5, 1	3, 4	0, 1

$\xrightarrow{a \succ d}$

	a	b	c
A	5, 5	4, 0	2, 2
B	0, 4	6, 6	4, 5
C	6, 2	5, 1	3, 4

$\xrightarrow{C \succ A}$

	a	b	c
B	0, 4	6, 6	4, 5
C	6, 2	5, 1	3, 4

$\xrightarrow{c \succ a}$

	b	c
B	6, 6	4, 5
C	5, 1	3, 4

$\xrightarrow{B \succ C}$

	b	c
B	6, 6	4, 5

$\xrightarrow{b \succ c}$

	b
B	6, 6

IESDS → Dominance Solvable

1. Let $\Sigma_i(R)$ be a set of mixed strategies with support R. e.g. $\Sigma(\{A,C\}) = \{pA + (1-p)C\}$
2. For all players i, define $R_i^0 = A_i$: set of pure strategies
3. Let $R_i^k \subseteq R_i^{k-1}$ be a subset of undominated actions, i.e., such that for any $a_i \in R_i^k$, there is no such $\alpha_i \in \Sigma(R_i^{k-1})$ that $u(\alpha_i, a_{-i}) > u(a_i, a_{-i})$ for all $a_{-i} \in R_{-i}^{k-1}$
   —— R_i^k is a set of pure strategies which are not strictly dominated by another (mixed) strategy from R_i^{k-1}
4. The set of pure strategies which survive the iterated elimination of strictly dominated strategies is denoted $R_i^\infty = \bigcap_k R_i^k$
5. If R_i^∞ is singleton, than the game is called **dominance solvable**

property : In finit games the order of elimination doesn't matter

In games with infinite action spaces the order-independence of IESDS solution is not guaranteed unless the strategy space is compact and payoff functions are continuous, see Dufwenberg and Stegeman (Econometrica, 2002).

紧集 = 有界闭集

A Cournot Oligopoly Game

- Players : 2 firms
- Actions : $A_i = [0,1)$, firm chooses $q_i \in A_i$
- Payoff : $A_1 \times A_2 \to \mathbb{R}^2$, $\pi_i(q_1, q_2) = \begin{cases} q_i[1 - q_i - q_{-i}] & \text{if } q_1 + q_2 \le 1 \\ 0 & \text{if } q_1 + q_2 > 1 \end{cases}$

1. By maximising $q_i(1 - q_i - q_{-i})$ we get a reaction function of firm i is $q_i = \frac{1-q_{-i}}{2}$.
2. Denote $R^k = (\underline{q}^k, \overline{q}^k)$, $\underline{q}^0 = 0, \overline{q}^0 = 1$.
3. Then, given the reaction function, $\overline{q}^1 = (1-0)/2 = 1/2$.
4. Then, given the reaction function, $\underline{q}^2 = (1-1/2)/2 = 1/4$.
5. Then, given the reaction function, $\overline{q}^3 = (1-1/4)/2 = 3/8$.
6. Note that (either upper or lower) bound on step k is given by

$$q^k = \frac{1-q^{k-1}}{2} = \frac{1}{2}\left(1 - \frac{1}{2} + \frac{1}{4} + \ldots + \left(-\frac{1}{2}\right)^{k-1}\right)$$

7. By taking the limit $k \to \infty$

$$\lim_{k\to\infty} q_k = \frac{1}{2}\left(1 - \frac{1}{2} + \frac{1}{4} - \frac{1}{8} + \frac{1}{16} - \frac{1}{32}\ldots\right) = \frac{1}{2}\left(\frac{1}{2} + \frac{1}{8} + \frac{1}{32}\ldots\right) =$$
$$\frac{1}{4}\left(1 + \frac{1}{4} + \frac{1}{16}\ldots\right) = \frac{1}{4} \times \frac{1}{1-\frac{1}{4}} = \frac{1}{3}$$

8. Thus, Cournot Duopoly is dominance solvable.

Iterative Elimination of Weakly Dominated Strategies (IEWDS)

$\forall a_{-i}, \exists \alpha_i', v_i(\alpha_i', a_{-i}) \geq v_i(a_i, a_{-i})$: strategy a_i weakly dominated

	L	M	R
T	1, 1	1, 1	0, 0
B	0, 0	1, 2	1, 2

$\xrightarrow{M \succ R}$

	L	M
T	1, 1	1, 1
B	0, 0	1, 2

$\xrightarrow{T \succ B}$

	L	M
T	1, 1	1, 1

	L	M	R
T	1, 1	1, 1	0, 0
B	0, 0	1, 2	1, 2

$\xrightarrow{M \succ L}$

	M	R
T	1, 1	0, 0
B	1, 2	1, 2

$\xrightarrow{B \succ T}$

	M	R
B	1, 2	1, 2

property: the order of elimination matters
remark: the Nash Equilibria can be lost during IEWDS

Rationalizability

x common knowledge of rationality (stricter)
you cannot believe your opponents' having probability of playing some strictly dominated strategies.
— belief in everyone's being rational

$p_i(a_{-i})$: probability player i puts on a_{-i}, with $\sum_{a_{-i} \in \prod_{j \neq i} A_j} p_i(a_{-i}) = 1$ "belief"

We call player i rational if he maximizes his expected payoff given $p_i(\cdot)$

Definition: Player's action a_i is called rationalizable if it is a best response to some mix of the opponent's rationalizable actions.
(player's i's belief of a_{-i})

Definition

The action a_i^* of player i is rationalizable if for each player j there exists a set $Z_j \subseteq A_j$ such that

1. $a_i^* \in Z_i$
2. *for every player j every action $a_j^* \in Z_j$ maximizes the expected payoff of player j given their belief p_j that assigns a positive probability only to action profiles in Z_{-j}.*

Independent vs Correlated Beliefs

W:

	L	R
U	8, 8, 8	0, 0, 0
D	0, 0, 0	0, 0, 0

X:

	L	R
U	4, 4, 4	0, 0, 0
D	0, 0, 0	4, 4, 4

Y:

	L	R
U	0, 0, 0	0, 0, 0
D	0, 0, 0,	8, 8, 8

Z:

	L	R
U	3, 3, 3	3, 3, 3
D	3, 3, 3	3, 3, 3

X is not rationalizable if we allow for product beliefs:

- Let $\Pr(U) = p$, $\Pr(L) = q$, the expected payoff from X is $4pq + 4(1-p)(1-q)$.
- X is not dominated by W and Y if and only if

$$4pq + 4(1-p)(1-q) \geq \max\{8pq, 8(1-p)(1-q)\} \quad \Rightarrow \quad p = q = 1/2.$$

- Then given $p = q = 1/2$ we get $\underbrace{v_3(X)}_{\text{expected payoff from X}} = 2 < 3 = v_3(Z)$, thus X is dominated by Z.

X is rationalizable if we allow for correltated beliefs:

- Let $\Pr(U, L) = p'$, $\Pr(D, R) = q'$, the expected payoff from X is $4p' + 4q'$.
- X is not dominated by W and Y if and only if $4p' + 4q' \geq \max\{8p', 8q'\} \quad \Rightarrow \quad p' = q' = 1/2$.
- Then given $p' = q' = 1/2$ we get $v_3(X) = 4 > 3 = v_3(Z)$, thus X is not dominated.

In finite games the set of correlatedly rationalizable actions coincides with the set of actions surviving the iterated elimination of strictly dominated strategies.

In two player finite games there is no need to think about correlations: all actions surviving IESDS are rationalizable and the other way around.

How reasonable it is to allow for correlated beliefs? Aumann (Econometrica, 1987) argues: “In games with more than two players, correlation may express the fact that what 3, say, thinks that 1 will do may depend on what he thinks 2 will do. This has no connection with any overt or even covert collusion between 1 and 2; they may be acting entirely independently. Thus it may be common knowledge that both 1 and 2 went to business school, or perhaps to the same business school; but 3 may not know what is taught there. In that case 3 would think it quite likely that they would take similar actions, without being able to guess what those actions might be.”

Nash Equilibrium

Definition

A NE (in pure strategies) is an action profile $a^* = \{a_1^*, \dots, a_n^*\}$ such that for any action a_i of any player i

$$u_i(a^*) \geq u_i(a_i, a_{-i}^*)$$

{ NE: nobody wants to deviate given others play the same strategies
Dominance: nobody wants to deviate regardless of others' strategies

Definition

The best response correspondence B_i of player i is defined as

$$B_i(a_{-i}) = \{a_i \in A_i : u_i(a_i, a_{-i}) \geq u_i(a_i', a_{-i}) \; \forall a_i' \in A_i\}$$

A set of 'interesections' of BR corespondances gives a full net of NE.
proof: a^* in a NE iff. $\forall i$, $a_i^* \in B_i(a_{-i}^*)$

Example Revisted: Cournot, Bertrand

the intersection of best responses ⇒ NE

Definition

A strategy profile α^* is a NE (in mixed strategies) if for any strategy α_i of player i

$$v_i(\alpha_i^*, \alpha_{-i}^*) \geq v_i(\alpha_i, \alpha_{-i}^*)$$

Mixed NE and Indifference

The mixed strategy profile α^* is NE iff. for each player i
{ α_i assigns zero probability to all pure actions $a_i \mid V_i(a_i, \alpha_{-i}^*) < V_i(\alpha_i^*, \alpha_{-i}^*)$
there are no actions a_i for which $V_i(a_i, \alpha_{-i}^*) > V_i(\alpha_i^*, \alpha_{-i}^*)$

Problem (O'Neill's Card Game)

	Ace	King	Queen	Jack
Ace	1, 0	0, 1	0, 1	0, 1
King	0, 1	0, 1	1, 0	1, 0
Queen	0, 1	1, 0	0, 1	1, 0
Jack	0, 1	1, 0	1, 0	0, 1

Label the probability that player 1 plays his pure actions p_1, p_2, p_3 and p_4. Similarly, for player 2, label these q_1, q_2, q_3 and q_4.

Note, that any playable action should lead to the same expected payoff!

$$\begin{cases} \mathbb{E}U_1(A) = q_1 \\ \mathbb{E}U_1(K) = q_3 + q_4 \\ \mathbb{E}U_1(Q) = q_2 + q_4 \\ \mathbb{E}U_1(J) = q_2 + q_3 \\ \sum q_i = 1 \end{cases} \Rightarrow \begin{cases} q_3 + q_4 = q_2 + q_4 \Rightarrow q_3 = q_2 \\ q_2 + q_3 = q_3 + q_4 \Rightarrow q_2 = q_4 \end{cases} \Rightarrow \begin{cases} q_1 = 2/5 \\ q_2 = 1/5 \\ q_3 = 1/5 \\ q_4 = 1/5 \end{cases}$$

Nash Equilibrium: Discussion

- General considerations:
  - unlike rationalizability, Nash equilibrium implies that player's beliefs are correct in equilibrium
  - common knowledge of rationality does not imply Nash equlibrium
  - one of possible solutions is to view Nash equilibrium as a stability concept: none of the players wants to deviate, but we do not ask a question why a specific action profile is chosen in the fist place
  - no predictions on which equilibrium will be played in a game with multiple equilibria (recall Battle of Sexes)
  - see link to evolutionary play later in this course.
- Mixed strategy equilibrium:
  - especially problematic, as players do not have any incentives to randomize in a specific way over the actions in the support of mixed strategies, since all actions deliver exactly the same expected payoff
  - stability interpretation: stable social configuration of actions
  - Aumann and A. Brandenburger (1995) interpret a mixed strategy of player i as payer j's belief about player i play
  - we will discuss purification later in this course.

Mixed NE with continuous Actions

price dispersion

{ heterogeneous information → discontinuity in demand
firms' mixed strategy equilibrium

Two firms produce a homogeneous good at marginal production costs c.
There is a unit mass of consumers. Each consumer has valuation v > c for the good, which is the same for all consumers.
Fraction λ of consumers are informed. These consumers know both the prices and buy at the lowest price. If the prices are equal, they split evenly.
Fraction 1 – λ of consumers are uninformed, they just buy the good at the closest store. We assume that they split equally across the firms.

Mixed Strategy of firm i defined by distribution $F_i(p) = Pr(\text{price} \le p)$
symmetric candidate equilibrium: $F(p)$

Claim 1. Support of $F(p)$ is a convex set. If there is a gap $[a, b]$ in the support, then $\pi(a) < \pi(b - \varepsilon)$ for sufficiently small ε.

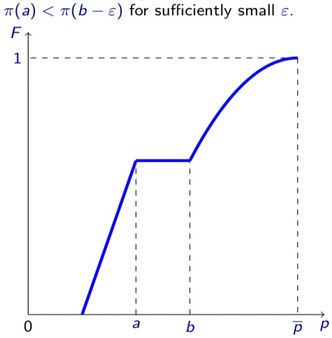

Claim 2. $F(p)$ is continuous. If there is an atom at p_0, then $\pi(p_0 - \varepsilon) > \pi(p_0 + \varepsilon) \Rightarrow$ there must be a gap above p_0.

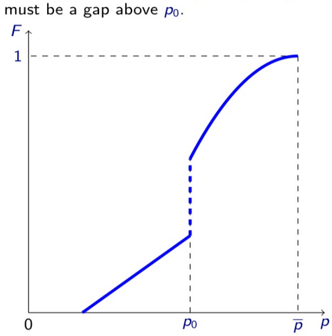

the profit of a firm charging p comes from:

uninformed customers: $\frac{1-\lambda}{2}(p-c)$

Informed customers, if it chargest the lowest price: $\lambda[1-F(p)](p-c)$

$\Rightarrow \pi(p) = \lambda[1-F(p)](p-c) + \frac{1-\lambda}{2}(p-c)$

upper bond: $\bar{p} = v$, $\pi = \frac{1-\lambda}{2}(v-c)$

solution (constant expected profits for any p): $\pi(p) = \pi$

$\therefore F(p) = 1 - \left(\frac{1-\lambda}{2\lambda} \cdot \frac{v-p}{p-c}\right)$, $\underline{p} = \frac{(1-\lambda)v + 2\lambda c}{1+\lambda}$ (set $F(\underline{p}) = 0$)

Correlated Equilibrium

Definition

A correlated equilibrium of a strategic game consists of:

- *a finite state Ω and a probability measure μ over this space*
- *for each player i a partition $\mathcal{P}_i(\omega)$*
- *for each player i a strategy $\sigma_i : \Omega \to A_i$ with $\sigma_i(\omega) = \sigma_i(\omega')$ whenever $\omega, \omega' \in P_i$ with $P_i \in \mathcal{P}_i$*

such that for all i and ω

$$\sum_{\omega' \in \Omega} \mu(\omega'|P_i(\omega)) u_i(\sigma_i(\omega'), \sigma_{-i}(\omega')) \geq \sum_{\omega' \in \Omega} \mu(\omega'|P_i(\omega)) u_i(s_i(\omega'), \sigma_{-i}(\omega'))$$

Suppose that the state space is generated by a through of the dice Ω = {1, 2, 3, 4, 5, 6}.
Partitions: P1 = ({1, 2, 3}, {4, 5, 6}), P2 = ({1, 3, 5}, {2, 4, 6})
Players condition their action on their partition:

whether ω is smaller than 4 or not for player 1

whether ω is odd or even for player 2

		L $\frac{2}{3}$	R $\frac{1}{3}$
$\frac{2}{3}$	U	6, 6	2, 7
$\frac{1}{3}$	D	7, 2	0, 0

	L	R
U	y	x
D	z	0

This game has two pure strategy Nash Equilibria: (D, L) and (U, R).

Alternating between these two equilibria delivers the average payoff of $4\frac{1}{2}$ and a mixed strategy equilibrium with expected payoff $4\frac{2}{3}$. Is there a correlated equilibrium with higher payoffs?

Let $\Omega = \{x, y, z\}$, let $\mathcal{P}_1 = (\{x, y\}, \{z\})$ and $\mathcal{P}_2 = (\{x\}, \{y, z\})$ with $\mu(x) = \mu(y) = \mu(z) = 1/3$.
U D R L

Let

- $\sigma_1(\{x, y\}) = U$ and $\sigma_1(\{z\}) = D$;
- $\sigma_2(\{x\}) = R$ and $\sigma_2(\{y, z\}) = L$.

Deviation incentives:

- If player 1 observes $\{x, y\}$ then $v_1(U) = \frac{1}{2}6 + \frac{1}{2}2 \geq \frac{1}{2}7 = v_1(D)$.
- If player 1 observes $\{z\}$ then $v_1(D) = 7 \geq 6 = v_1(U)$.
- Expected payoff is $\frac{2}{3}4 + \frac{1}{3}7 = 5$. — revelation principle in mechanism design

Supplement

don't survive IESDS / IDSDS. not a best response to <u>any</u> belief
Rationalizable (RAT): best response to <u>some</u> belief
support by rationalizable strategies
Nash Equilibrium (NE): best response to the <u>correct</u> belief

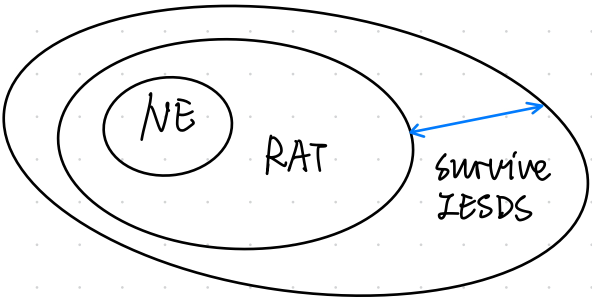

in 2-player game or $n \geq 2$ correlated strategy games:
RAT ≡ survive IESDS

"correlate our decisions"
⇒ (n-1) opponents merge into a super opponents against you

process of finding NE:

1. IESDS (consider mixed strategies)
   S_i is strictly dominated if there $\exists S_i'$ for every S_{-i} s.t. $u_i(S_i', S_{-i}) > u_i(S_i, S_{-i})$
   S_i is Not Best Response if for every S_{-i}, there $\exists S_i'$ s.t. $u_i(S_i', S_{-i}) > u_i(S_i, S_{-i})$
   —— only the first leads to deletion

   notice the "mixed" possibility doesn't matter for NE latter

	$L^{\frac{1}{4}}$	C	$R^{\frac{3}{4}}$
T	,3	,1	,0
B	,0	,1	,3

	L	C	R
T	,1.1	,1	,0
B	,0	,1	,1.1

2. pure NE
3. mixed NE
   —— if strategies 3x3, then a total of $(2^3-1) \times (2^3-1)$ combinations

if players have strict preferences over the outcomes then the number of Nash equilibria is odd!

$IB(p,q) = (BR_1(q), BR_2(p))$ if "==", Nash Equilibrium

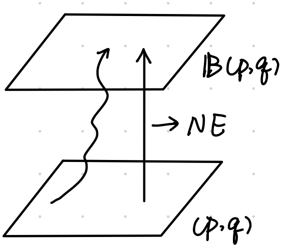

⇒

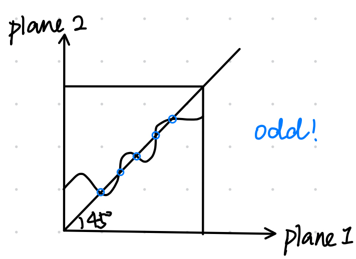

Supplement

Problem 2. Two 'oligarchs', A and B have a dispute over an asset of value X. Although both oligarchs now reside in city L, they decided to settle their dispute in court in city M, where they originally made their money. The reason is that the legal system in city M is simple and transparent: both contestants provide the judge with the 'evidence', which takes form of a briefcase with cash. The judge keeps the evidence from both contestants and decides the case in favour of the one who provided more evidence. If both contestants provided equal amount of evidence the judge flips a fair coin to make a decision. Assume that oligarchs submit their briefcases simultaneously and any contribution $y \in$ R+ can be put in the envelope.

1. Derive the best response of oligarch A to oligarch's B contribution.
2. Using your answer to part (1) show that there is no pure strategy equilibrium.
3. Prove that there is no contribution value played with a positive probability (Hint: consider small upward deviations).
4. Using your answer to (3), or otherwise, show if $F(y)$ is a symmetric mixed strategy equilibrium distribution then it does not have gaps in its support.
5. Argue that the lowest contribution in the support of $F(y)$ must be equal to zero.
6. Derive the equilibrium strategy of the oligarchs.

1. Assume A puts y_A. B puts y_B in the envelope.

$$u_A(y_A; y_B) = \begin{cases} X - y_A, & \text{if } y_A > y_B \\ \frac{X}{2} - y_A, & \text{if } y_A = y_B \\ -y_A, & \text{if } y_A < y_B \end{cases}$$

Best Response $B_A(y_B) = \begin{cases} \phi, & y_B < X \\ 0, & y_B \geq X \end{cases}$

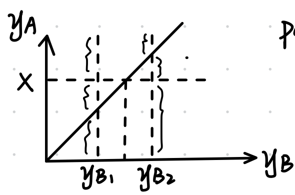

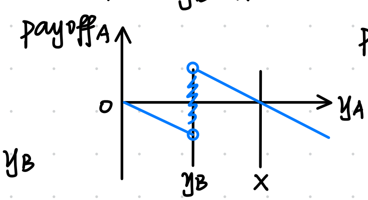

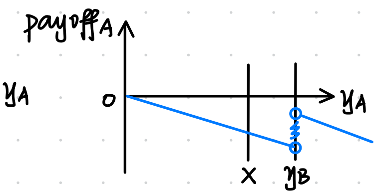

2.

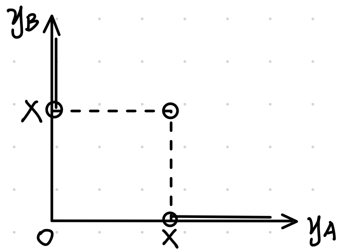

No pure Nash Equilibrium.

3. Proof: (F,F) cannot be an equilibrium when F looks like this:

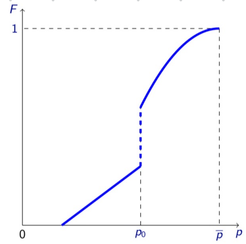

$F(p)$ is continuous.
There is no atom p_0.

Say F assigns mass $m(y_0)>0$.

$$\begin{aligned}
\text{Payoff}_A(y_0,F) &= m(y_0)(\tfrac{1}{2}X-y_0) && \text{Tie}\\
&\quad + \Pr(B \text{ plays} < y_0)\cdot(X-y_0) && \text{A wins}\\
&\quad + \Pr(B \text{ plays} > y_0)\cdot(-y_0) && \text{A loses}
\end{aligned}$$

$$\begin{aligned}
\text{Payoff}_A(y_0+\varepsilon,F) &= m(y_0)(X-y_0-\varepsilon) + \Pr(B \text{ plays} < y_0)\cdot(X-y_0-\varepsilon)\\
&\quad + \Pr(B \text{ plays} > y_0)\cdot(?)\\
&\geq m(y_0)(X-y_0-\varepsilon) + \Pr(B<y_0)(X-y_0-\varepsilon) + \Pr(B>y_0)(-y_0-\varepsilon)\\
&= m(y_0)\cdot\tfrac{1}{2}X + m(y_0)(\tfrac{1}{2}X-y_0) + \Pr(B<y_0)(X-y_0) + \Pr(B>y_0)(-y_0) - \varepsilon\\
&= m(y_0)\cdot\tfrac{1}{2}X + \text{Payoff}_A(y_0,F) - \varepsilon
\end{aligned}$$

$\forall\, m(y_0)>0,\ X>0.\ \exists\ 0<\varepsilon<m(y_0\cdot\tfrac{1}{2}X)$, s.t. $\text{Payoff}_A(y_0+\varepsilon,F) > \text{Payoff}_A(y_0,\varepsilon)$

$\Rightarrow$ not stable at y_0.

4. Proof: (F,F) cannot be an equilibrium when F looks like:

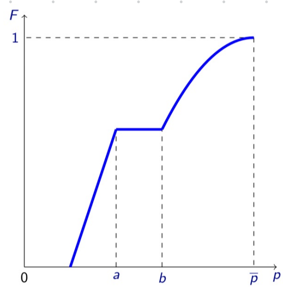

There's a gap $[y_0, y_1]$ in the support

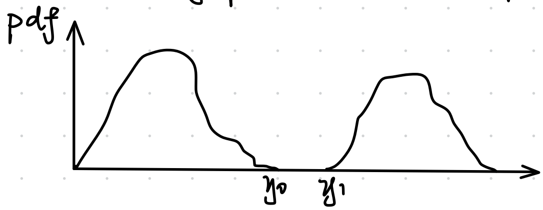

$$\text{Payoff}_A(y,F) = \Pr(B \text{ plays} < y)\, X - y = F(y)\cdot X - y$$

$$\text{Payoff}_A(y_0,F) = F(y_0)\cdot X - y_0 > \text{Payoff}_A(y_1,F) = F(y_1)X - y_1$$

5. Say $\text{supp}(F) = [\underline{y}, \bar{y}]$

From (3), $F(\underline{y})=0$, thus $\text{Payoff}_A(\underline{y},F) = -\underline{y} = 0 \Rightarrow \underline{y}=0$

(if $\underline{y}>0$, the payoff <0, but we can always guarantee 0 by setting $\underline{y}=0$)

6. $\text{Payoff}_A(y, F) = 0$ for $\forall y$:

$$F(y)(X-y) + [1-F(y)](-y) = 0$$

$\Rightarrow F(y) = \frac{y}{X}$ (uniform distribution)

Correlated Equilibrium

	L	R
T	x	y
B	z	w

"four traffic lights"

s.t. $x+y+z+w=1$

$\Rightarrow df = 3$

← told "T" $\Rightarrow$ x or y

Mixed Equilibrium

	L	R
T	pq	$p(1-q)$
B	$(1-p)q$	$(1-p)(1-q)$

$df = 2$

Correlated Equilibrium provide a reasonable generalization of mixed-strategy NE.

But Mixed NE could be weird:

	A	B
A	X, X	$0, 0$
B	$0, 0$	$1, 1$

$(X >> 1)$

mixed NE:

$$\left(\frac{1}{X+1}A + \frac{X}{X+1}B\right)$$

Remarks: NE omits / doesn't provide

1. selection among NEs
2. reasoning process / dynamic thing

Weeks 2:
Static Games of Imperfect Info

Bayes-Nash Equilibrium
Purification
Global games

Terminology

Strategic Games of Incomplete Information (Bayesian Games):
players have incomplete information about other players' preferences and characteristics.

Osborn: imperfect information
Harsanyi: games of incomplete information = games of imperfect information that have the following timing:

- nature draws a type for each player
- nature reveals to each player his or her type only
- players simoultaneously choose strategies
- (type-dependent) payoffs are realized

—— Thus games of incomplete information can be represented as games of complete but imperfect information about the players' type.

Definition

A game of incomplete information consists of

- *a set of players N;*
- *a set of* **states** *or types Ω;* state $\omega \in \Omega$
- *a set of actions for each player A_i;*
- *a set of* **signals** *and a signal function for each player;* space of signals: S; $\Omega \to S$
- *a* **belief** *associated to each signal received by each player;* signals + prior distribution → posterior belief
- *a Bernoulli payoff function over pairs of actions and states (a, ω).*

Consider a version BoS game, where one of the players can be either cooperative or not, another one is always cooperative:

	b	s
B	$\underline{2}, \underline{1}$	$0, 0$
S	$0, 0$	$\underline{1}, \underline{2}$

Cooperative Pl.1

	b	s
B	$0, \underline{1}$	$\underline{2}, 0$
S	$\underline{1}, 0$	$0, \underline{2}$

Non−Cooperative Pl.1

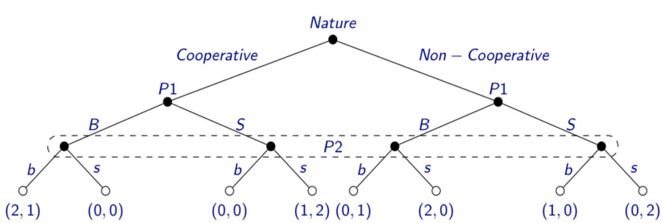

strategy space for
player 1: $\{BB, BS, SB, SS\}$
signal m^1_c: cooperative
signal m^1_{NC}: non-cooperative
player 2: $\{b, s\}$
(only pure strategies mentioned)

BoS: More Complicated Signals

Suppose: $P(m_c^2 \mid \text{Cooperative}) = q_1$, $q_1, q_2 > \frac{1}{2}$
$P(m_{NC}^2 \mid \text{Non-Cooperative}) = q_2$
prior belief $P(\text{Cooperative}) = p$

$\Rightarrow$ Player 2 needs to update his beliefs about the distribution of player 1 types conditional on the signal:

$$Pr(m_c^1 \mid m_c^2) = \frac{Pr(m_c^1 \cap m_c^2)}{Pr(m_c^2)} = \frac{Pr(m_c^2 \mid m_c^1) Pr(m_c^1)}{Pr(m_c^2 \mid m_c^1) Pr(m_c^1) + Pr(m_c^2 \mid m_{NC}^1) Pr(m_{NC}^1)} = \frac{pq_1}{pq_1 + (1-p)(1-q_2)}$$

, and similarly for $Pr(m_c^1 \mid m_{NC}^2)$

Player 1 also needs to form beliefs about the type of player 2:

$Pr(m_c^2 \mid m_c^1) = q_1$, $Pr(m_c^2 \mid m_{NC}^1) = 1 - q_2$

Bayes - Nash Equilibrium

Definition

A Nash equilibrium of a Bayesian game (Bayes-Nash equilibrium) is a Nash equilibrium of the strategic game defined as follows:

- **Players:** *Set of all pairs (i, s_i), where i is a player of Bayesian game, and s_i is a signal that player i can receive.*
- **Actions:** *Set of actions of player (i, s_i) is a set of actions of player i in the Bayesian game.*
- **Payoffs:** *Bernoully playoff function of player (i, s_i):*

$$u_{i,s_i} = \int_{\omega \in \Omega} u_i(a_i, a_{-i}(\omega), \omega) d\mathbb{P}(\omega \mid s_i)$$

where $a_{-i}(\omega)$ is an action profile of all other players given the signals they might get in state ω: $a_j(j, s_j(\omega))$.

Example: BoS with $Pr(\text{Cooperative}) = p = \frac{1}{2}$

(1) pure strategy BNE

	b	s
BB	1, 1	1, 0
BS	1.5, 0.5	0, 1
SB	0, 0.5	1.5, 1
SS	0.5, 0	0.5, 2

(SB, s) is unique BNE

(2) mixed strategy BNE

Cooperative Pl.1

	b	s
B	2, 1	0, 0
S	0, 0	1, 2

Non-Cooperative Pl.1

	b (β)	s
B	0, 1	2, 0
S	1, 0	0, 2

Suppose $\sigma'(m_c^1) = B$, $\sigma'(m_{NC}^1) = \alpha B + (1-\alpha) S$

player 2 indifferent: $\left.\begin{array}{l} V_2(b) = \frac{1}{2} \times 1 + \frac{1}{2}\alpha \\ V_2(S) = \frac{1}{2}(1-\alpha) \times 2 \end{array}\right\} \Rightarrow \alpha = \frac{1}{3}$

player 1 indifferent: $\left.\begin{array}{l} V_{NC}(B) = 2(1-\beta) \\ V_{NC}(S) = \beta \end{array}\right. \Rightarrow \beta = \frac{2}{3}$

(check the cooperative type doesn't want to deviate:
$V_C(B) = \frac{4}{3} \geq V_C(S) = \frac{1}{3}$)

$\Rightarrow \left\{ \begin{pmatrix} B \\ \frac{1}{3}B + \frac{2}{3}S \end{pmatrix}, \frac{2}{3}b + \frac{1}{3}s \right\}$

Examples

1. Swing Voter's Curse

- 2 states of the economy: A, B $\Rightarrow$ in state A, 1 better than 2, in state B, 2 better than 1.
- 2 candidates: 1, 2
- 2 voters: $u=1$ (best candidate wins), $u=0$ (worst), $u=\frac{1}{2}$ (tie)
  - citizen 1: informed of the state of the nature
  - citizen 2: $p(A)=0.9$, $p(B)=0.1$

Formalization:
- players: citizen 1 and 2
- states: $\{A, B\}$
- actions: vote for 1/2, do not vote: $\{1,2,0\}$
- signals: different in A or B / always the same
- beliefs: consistent with signal / $p(A)=0.9$, $p(B)=0.1$

payoffs:

	State A			State B		
	0	1	2	0	1	2
0	$\frac{1}{2},\frac{1}{2}$	1, 1	0, 0	$\frac{1}{2},\frac{1}{2}$	0, 0	1, 1
1	1, 1	1, 1	$\frac{1}{2},\frac{1}{2}$	0, 0	0, 0	$\frac{1}{2},\frac{1}{2}$
2	0, 0	$\frac{1}{2},\frac{1}{2}$	0, 0	1, 1	$\frac{1}{2},\frac{1}{2}$	1, 1

Best Response:

Player 1:

Player 2 plays	0	1	2
Player's 1 best response	(1, 2)	(0, 2)or (1, 2)	(1, 0)or (1, 2)

Player 2:

Player 1	(0, 2)	(1, 0)	(1.2)
0	0.55	0.95	1
1	0.95	0.9	0.95
2	0.1	0.55	0.55

Equilibria: ((1,2), 0), ((0,2), 1)

2. BoS 2x2 problem

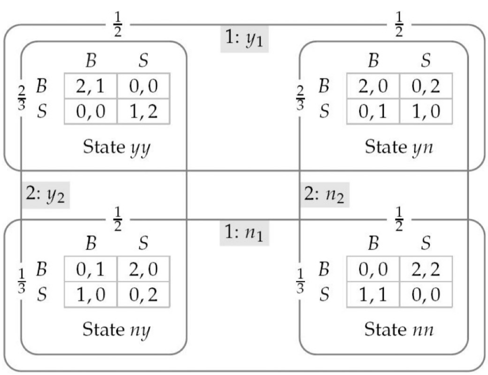

For each player denote WX a strategy "play W if I am y type and play X if I am n type".

Type y_1:

y_1/Pl.2	BB	BS	SB	SS
B	2	1	1	0
S	0	1/2	1/2	1

Type n_1:

n_1/Pl.2	BB	BS	SB	SS
B	0	1	1	2
S	1	1/2	1/2	0

BR^1:

Player 2 plays	BB	BS	SB	SS
Player's 1 BR	BS	BB	BB	SB

Similarly, BR^2:

Player 1 plays	BB	BS	SB	SS
Player's 2 BR	BS	BS, SS	SS	SB

$\Rightarrow$ two equilibria: (BB, BS), (SB, SS)

3. Cournot

Consider a Cournout Oligopoly model with the demand function p = 1 – q1 – q2. Suppose firm one has zero production costs, while firm two has costs either 0 (with probability α) or c ∈ (0, 2/(4 – α)). Firm two knows its type, firm one does not know it. Find BNE of the game.

Firm 1: $\pi_1 = [\alpha(1-q_1-q_{2L}) + (1-\alpha)(1-q_1-q_{2H})]\, q_1 = (1-q_1-q_2^e)\, q_1 \Rightarrow q_1 = \frac{1-q_2^e}{2}$

Firm 2: $\begin{cases} \text{Low cost: } \pi_{2L} = (1-q_1-q_{2L})\, q_{2L} \Rightarrow q_{2L} = \frac{1-q_1}{2} \\ \text{High cost: } \pi_{2H} = (1-q_1-q_{2H})\, q_{2H} - c q_{2H} \Rightarrow q_{2H} = \frac{1-q_1-c}{2} \end{cases}$

, where $q_2^e = \alpha q_{2L} + (1-\alpha) q_{2H}$

Solution: $q_1 = \frac{1+(1-\alpha)c}{3}$, $q_{2L} = \frac{2-(1-\alpha)c}{6}$, $q_{2H} = \frac{2-(4-\alpha)c}{6}$

Continuum of Types

1. Bertrand Competition

Setting and Strategies

Suppose that there is a single buyer who is willing to pay up to 1 pound for a product. There are two sellers, each seller has production cost ci , which is independently uniformly distributed on [0, 1]. Sellers simultaneously make their offers and the buyer buys at the lowest price or randomly chooses the seller if the prices are equal. Find the BNE of the game.

seller's linear increasing pricing strategy: $p_i = \alpha_i + \beta_i c_i$

$\pi_i = (p_i - c_i)\cdot Pr(p_i < p_{-i}) = (p_i - c_i)\cdot Pr(p_i < \alpha_{-i} + \beta_{-i} c_{-i})$

$= (p_i - c_i)\cdot Pr\left(\frac{p_i - \alpha_{-i}}{\beta_{-i}} < c_{-i}\right) = (p_i - c_i)\left(1 - \frac{p_i - \alpha_{-i}}{\beta_{-i}}\right)$

$\frac{\partial \pi_i}{\partial p_i} = 0 \Rightarrow p_i = \frac{1}{2}c_i + \frac{1}{2}(\alpha_{-i} + \beta_{-i}) \Rightarrow \beta_i = \frac{1}{2}, \alpha_i = \frac{1}{2}$

thus: $p_i = \frac{1}{2}c_i + \frac{1}{2}$

seller's arbitrary increasing pricing strategy $p(c)$:

Deviation: pretend to be a different type (call it $\hat{c}$)

$\pi(\hat{c}; c) = [p(\hat{c}) - c]\cdot Pr(p(\hat{c}) < p(c_{-i}))$

$= [p(\hat{c}) - c] \times Pr(\hat{c}_i < c_{-i}) = [p(\hat{c}) - c](1-\hat{c})$

$\frac{\partial \pi(\hat{c}; c)}{\partial \hat{c}} = p'(\hat{c})(1-\hat{c}) - p(\hat{c}) + c = 0$

the best possible deviation is to the firm's true type (i.e., $\hat{c} = c$):

$p'(c)(1-c) - p(c) + c = 0$

2. Double Auction

Valuation for the good:

$$\begin{cases} \text{buyer: } V_b \sim U[0,1] \\ \text{seller: } V_s \sim U[0,1] \end{cases}$$

Both buyer and seller simultaneously announce their prices P_b and P_s.
If $P_b \geq P_s$, the trade takes place at a price $(P_b+P_s)/2$.

payoffs (when $p_b < p_s$: 0, otherwise):

$$\begin{cases} \Pi_b(P_b, P_s, V_b) = V_b - \frac{P_b+P_s}{2} \\ \Pi_s(P_b, P_s, V_s) = \frac{P_b+P_s}{2} - V_s \end{cases}$$

Suppose that the seller plays a linear strategy: $p_s(V_s) = \alpha + \beta V_s$

$\Pi_b = \Pr(\text{trade}) \cdot (V_b - \text{expected price})$

(1) $\Pr(\text{trade}) = \Pr(P_b \geq P_s) = \Pr(P_b \geq \alpha + \beta V_s) = \Pr(V_s \leq \frac{P_b - \alpha}{\beta}) = \frac{P_b-\alpha}{\beta}$

(2) since $p_s \sim U[\alpha, \alpha+\beta]$:

expected price $= \frac{P_b + E(P_s \mid P_s \leq P_b)}{2} = \frac{1}{2}(p_b + \frac{1}{2}(p_b + \alpha))$

$\Rightarrow \Pi_b(P_b, P_s, V_b) = \frac{1}{4\beta}(p_b - \alpha)(4V_b - 3p_b - \alpha)$

FOC: $p_b = \frac{2}{3}V_b + \frac{1}{3}\alpha$

Suppose that the buyer plays a linear strategy: $p_b = \gamma + \delta V_b$

$\Pi_s = \Pr(\text{trade}) \cdot (\text{expected price} - V_s)$

$= \frac{1}{4\delta}(\gamma + \delta - p_s)(3p_s + \gamma + \delta - 4V_s)$

$\Rightarrow p_s = \frac{2}{3}V_s + \frac{1}{3}(\gamma + \delta)$

Equilibrium Strategies:

$$\begin{cases} P_b = \frac{1}{12} + \frac{2}{3}V_b \\ P_s = \frac{1}{4} + \frac{2}{3}V_s \end{cases}$$

trade occurs when $P_b \geq P_s \Rightarrow V_b \geq V_s + \frac{1}{4}$

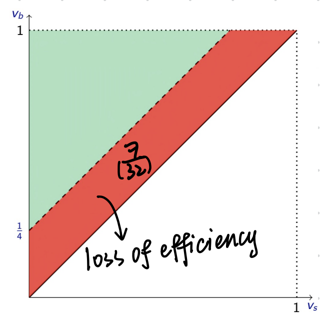

Actually, the linear equilibrium is the most efficient one.
(less efficient for other equilibria:)

choose some $x \in [0,1]$,

$$p_b = \begin{cases} x, & \text{if } V_b \geq x \\ 0, & \text{if } V_b < x \end{cases}, \quad p_s = \begin{cases} x, & \text{if } V_s \leq x \\ 1, & \text{if } V_s > x \end{cases}$$

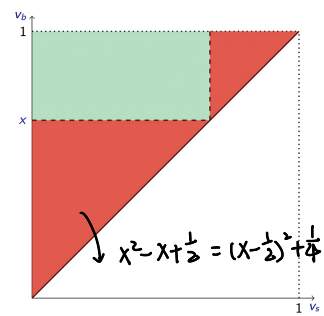

3. Adverse Selection

There is an asset (e.g. a firm) potentially available for sale.
The seller knows the value of the asset under its management.
The buyer does not know the value of the asset and thinks that its value is uniformly distributed on [0, 100].
The buyer believes that it has superior managing skills and under its management the value of the asset will go up by 50%.
The buyer proposes the price p for the asset, the seller can either accept or reject. If the offer is accepted, then the asset changes the ownership.

Formalization

- **Players:** the buyer and the seller.
- **States:** $v \in [0, 100]$.
- **Actions:** Buyer: $p \in \mathbb{R}_+$; seller: $\{a, r\}$.
- **Signals:** the seller gets a perfect signal $s = v$.
- **Beliefs:** Buyer has a uniform belief with support $[0, 100]$.
- **Payoffs:** 0 and v if the offer is rejected, $\frac{3}{2}v - p$ and p if the offer is accepted.

The seller accepts the offer whenever $p \geq v$.
The expected payoff of the buyer is:

$$\Pi_b = Pr(p \geq v) \cdot \left(\frac{3}{2} E(v \mid p \geq v) - p\right) = -\frac{p^2}{4}$$

$\Rightarrow$ buyer's best response: $p = 0$

Purification

Example: noisy BoS

	b	s
B	2,1	0,0
S	0,0	1,2

three Nash Equilibria:

$(B,b), (S,s), (\frac{2}{3}B+\frac{1}{3}S, \frac{1}{3}b+\frac{2}{3}s)$

perturbed game.

	b	s
B	$2+\varepsilon$,1	ε,δ
S	0,0	$1, 2+\delta$

ε and δ are players' type (extra payoff from their favourite action): private information

$\varepsilon, \delta \sim U[0,x]$ (independent)

suppose player 2 follows pure strategy

$$\begin{cases} \text{play } s \text{ if } \delta > \delta^* \\ \text{play } b \text{ if } \delta \le \delta^* \end{cases}$$

Best Response of player 1:

$$V_1(B) = (2+\varepsilon)\frac{\delta^*}{x} + \varepsilon\left(1-\frac{\delta^*}{x}\right) = 2\frac{\delta^*}{x} + \varepsilon$$

$$V_1(S) = 1\left(1-\frac{\delta^*}{x}\right)$$

$$V_1(B) \geq V_1(S) \Rightarrow \varepsilon \geq 1 - \frac{3\delta^*}{x} = \varepsilon^*$$

$$\Rightarrow \text{play } B \text{ if } \varepsilon \geq \varepsilon^*$$

$$\Rightarrow \left.\begin{aligned} \varepsilon^* &= 1-\frac{3\delta^*}{x} \\ \delta^* &= 1-\frac{3\varepsilon^*}{x} \end{aligned}\right\} \quad \varepsilon^* = \delta^* = \frac{x}{3+x}$$

$\begin{cases} \text{pure strategy: specific action without randomization} \\ \text{mixed strategy: involve randomization} \end{cases}$

Under pure strategy BNE, $Pr(B) = 1 - \frac{\varepsilon^*}{x} = \frac{2+x}{3+x}$

$$\lim_{x\to 0} \frac{3+x}{2+x} = \frac{2}{3}$$

Therefore, the limit point of outcomes of BNE equilibrium is mixed Nash Equilibrium of unperturbed game.

Proposition (Harsanyi purification)

The probability distributions over strategies induced by the pure (Bayesian Nash) equilibria of the perturbed game converge to the distribution of the (mixed Nash) equilibrium of the unperturbed game.

Global Games

Carlsson, van Damme (Econometica, 1993)

Definition: A game is randomly drawn from a class of games (global game), each player gets a noisy signal about which game is played.

—— As types of the players are randomly drawn conditional on the actual game played, types become correlated

Example:

Two investors simultaneously decide on whether they want to participate in a project or not:
1 participation costs for each investor is c;
2 the project succeeds only when both investors decide to invest;
3 if the project succeeds it brings v to each investor.

	i	n
I	v-c, v-c	-c, 0
N	0, -c	0, 0

(*) the participation costs are unknown ex ante and each players gets a conditionally independent signal about it.

Analysis:

1. Complete Information

$$\begin{cases} C<0: \text{ unique NE } (I,i) \\ C>V: \text{ unique NE } (N,n) \\ C\in[0,V]: \text{ three equilibria } (I,i),\ (N,n),\ \left(\frac{c}{v}I+\frac{v-c}{v}N,\ \frac{c}{v}i+\frac{v-c}{v}n\right). \end{cases}$$

2. Correlated Types

$v=10,\ c\in\{-2,-1,\cdots 11,12\}$, uniform distribution

signal: $\begin{cases} \text{one player learns true value } t=c. \\ \text{another player gets } t\in\{c-1,c+1\} \text{ equally likely} \end{cases}$

joint distribution (conditional on c):

t_1/t_2	c-1	c	c+1
c-1	0	$\frac{1}{4}$	0
c	$\frac{1}{4}$	0	$\frac{1}{4}$
c+1	0	$\frac{1}{4}$	0

suppose player 1 gets signal $t_1\in[1,11]$.

$$c|t_1=\begin{cases} t_1-1, & \text{w.p. } \frac{1}{4} \\ t_1, & \text{w.p. } \frac{1}{2} \\ t_1+1, & \text{w.p. } \frac{1}{4} \end{cases} \Rightarrow E(c|t_1)=t_1$$

$$t_2|t_1=\begin{cases} t_1-1, & \text{w.p. } \frac{1}{2} \\ t_1+1, & \text{w.p. } \frac{1}{2} \end{cases}$$ ($t_1\in\{-3,-2,12,13\}$ similarly)

See next page: correlated uncertainty delivers a unique prediction for almost all values of c.

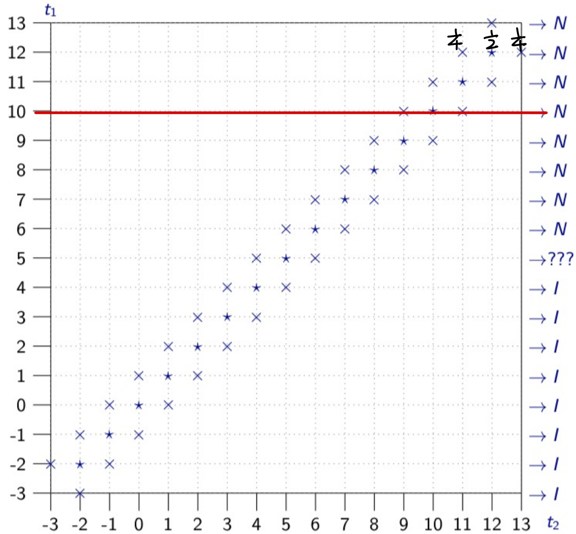

If $t_1 > 10$, $E(c|t_1) > 10$
⇒ not investing dominates

$V_1(I) = \Pr(a_2 = I | t_1 = 10) \times v - E(c|t_1)$
$\le [1 - \Pr(t_2 = 11 | t_1 = 10)] \times 10 - 10 < 0$
⇒ choose N

$V_i(I) = \Pr(a_{-i} = I | t_i = t) \times v - E(c|t_i)$
$\le [1 - \Pr(t_{-i} = t+1 | t_i = t)] \times v - t = \frac{v}{2} - c$

if both types invest for $t_i = t-1$
$V_i(I) = \Pr(I | t_i = t) \times v - E(c|t_i)$
$\ge \Pr(t_{-i} = t-1 | t_i = t) \times v - t = \frac{v}{2} - c$

$V_1(I) = \Pr(a_2 = I | t_1 = 0) \times v - E(c|t_1)$
$\ge \Pr(t_2 = -1 | t_1 = 0) \times v - 0 = \frac{1}{2} v > 0$
⇒ choose I.

If $t_1 < 0$, $E(c|t_1) < 0$ ⇒ investing

This conclusion is replicated in a game with a continuum of types and holds even if uncertainty is vanishingly small.
Thus, global games approach give a simple tool for equilibrium selection.

The outcome we obtained corresponds to a situation when a player has a “Laplacian prior” over opponent’s action: each action is taken with equal probability.
In this case a player should invest whenever $v/2 > c$ and abstain from investing if $v/2 < c$—same as in our BNE.

Risk Dominance

The risk-dominant equilibrium is the one with higher deviation-payoff product.

	L	R
T	a,b	c,d
B	e,f	g,h

(Suppose $a > e$, $b > d$, $g > c$, $h > f$)
deviation-payoff product: $\begin{cases} (T,L): (a-e)(b-d) \\ (B,R): (g-c)(h-f) \end{cases}$
(T,L) risk-dominant if $(a-e)(b-d) > (g-c)(h-f)$

Notice that in a symmetric game, risk dominance coincides with the equilibrium made up of strategies that are (strict) best responses to a 50:50 mix by the opponent.

	L	R
T	a, a	b, c
B	c, b	d, d

	L ($\frac{1}{2}$)	R ($\frac{1}{2}$)
T	5,5	3,2
B	2,3	4,4

If player 2's playing 50-50: player 1 should play T (the risk-dominant strat)

$$(a-c)^2 > (d-b)^2 \quad \Leftrightarrow \quad a - c > d - b \quad \Leftrightarrow \quad \frac{1}{2}a + \frac{1}{2}b > \frac{1}{2}c + \frac{1}{2}d.$$

Risk-dominance is a pairwise criterion.
See transitivity paradox:

	L	M	R
T	$\underline{6}, \underline{6}$	$0, 3$	$2, 0$
C	$3, 0$	$\underline{5}, \underline{5}$	$1, 4$
B	$0, 2$	$4, 1$	$\underline{4}, \underline{4}$

Check that (T, L) risk-dominates (B, R), which risk-dominates (C, M), which risk-dominates (T, L).
$6^2 > 2^2$ $\quad$ $3^2 > 1^2$ $\quad$ $5^2 > 3^2$

Global Random variable $\Rightarrow$ better (?) single (?) NE
(see PS 3-3)

Weeks 3 :
Knowledge in Games

Information and Knowledge

Information and Knowledge

Assume that there is a finite **state space Ω** and a **common prior μ**.

- A **state** (of the world) $\omega \in \Omega$ captures all relevant uncertainty: the state of nature (primitive uncertainty; e.g., payoffs), who knows what about it and about each others information, etc.

  Ω = {⚀, ⚁, ⚂, ⚃, ⚄, ⚅} with $\mu(\omega) = \frac{1}{6}$

- **Event** $E \subseteq \Omega$ is true (holds, obtains) in state ω if $\omega \in E$

  For example, E = {⚁, ⚃, ⚅}, i.e. E = "an even number appeared on the top of the dice".
  E holds in ω = ⚁ and does not hold in ω = ⚀.

Player Information

player i's information is represented by an information function which associates with each $\omega \in \Omega$ set $P_i(\omega)$ with properties:

- A collection P_i of sets $P_i(\omega)$ forms a **partition** of Ω.
- In state ω player i only knows that the state in $P_i(\omega)$

- Suppose that Ann knows whether the outcome of the dice was odd or even,

  Partition → $\mathcal{P}_A$ = ({⚀, $\overset{\omega_3}{⚂}$, $\overset{\omega_5}{⚄}$}, {⚁, ⚃, ⚅}).

  ↳ $\omega_1 \in P_A(\omega_1) = P_A(\omega_3) = P_A(\omega_5)$

- Suppose that Bob knows whether the number was 1-3, 4-5 or 6:

  $\mathcal{P}_B$ = ({⚀, ⚁, ⚂}, {⚃, ⚄}, {⚅}).

- In state ω = ⚁ Ann only knows that $\omega \in$ {⚁, ⚃, ⚅} and Bob only knows that $\omega \in$ {⚀, ⚁, ⚂}.

Knowledge

We say that player i **knows E in state ω** if $P(\omega) \subseteq E$

- the event that i knows E is $K_i(E) = \{\omega \in \Omega : P_i(\omega) \subseteq E\}$
- the mapping $K_i : 2^\Omega \to 2^\Omega$ is called player i's **knowledge operator**.

E = {⚀, ⚁, ⚂, ⚃} $K_B(E)$ = {⚀, ⚁, ⚂}

Event $E = \{\omega_1, \cdots \omega_k\} \overset{}{\underset{K_i}{\Rightarrow}} \{\omega_j\}$ s.t.
$E \subseteq$ Space 2^Ω $P_i(\omega_j) \subseteq E$
能让player i认知事件E为真的状态集

Bob knows E in $\omega \in$ {⚀, ⚁, ⚂} (e.g. $\mathcal{P}_B$(⚁) $\subseteq E$), but does not know it in ω = ⚃, as Bob cannot distinguish this outcome from ⚄.

Properties of Knowledge Operator

Definition:

$$K_i(E) = \{\omega \in \Omega : \mathcal{P}_i(\omega) \subseteq E\}$$ $K_B(E) = \{$⚀, ⚁, ⚂$\}$

Example:

$\mathcal{P}_B = ($\{⚀, ⚁, ⚂\}, \{⚃, ⚄\}, \{⚅\}$)$ $E = \{$⚀, ⚁, ⚂, ⚃$\}$

1. Truth: $K_i(E) \subseteq E$ (a player can only know something that is true)
2. Necessitation: $K_i(\Omega) = \Omega$ (a player knows the state space)
3. Inference: $E \subseteq F \Rightarrow K_i(E) \subseteq K_i(F)$

   $F = \{$⚀, ⚁, ⚂, ⚃, ⚄$\} \Rightarrow K_B(F) = \{$⚀, ⚁, ⚂, ⚃, ⚄$\} \Rightarrow K_B(E) \subseteq K_B(F)$.
4. Negative Introspection.

   $\neg K_i(E) \subseteq K_i(\neg K_i(E))$ (no unknown unknowns)

   不能让player i 认知事件E真假的状态集 ⊆ 能让他认知到自己对E无知的状态集
5. Positive Introspection:

   $K_i(E) = K_i(K_i(E))$ (a player knowns own information structure)

   player i 知道自己能认知事件E的状态集

Common Knowledge

1. The event that everybody knows E is:

   $K(E) = \cap_i K_i(E) = \{\omega \in \Omega : P_i(\omega) \subseteq E \text{ for all } i\}$
2. The event that everybody knows everybody knows E is:

   $K^2(E) = K(K(E))$

   More generally, $K^n(E) = \{\omega \in \Omega : P_i(\omega) \subseteq K^{n-1}(E) \text{ for all } i\}$
3. The event E is commonly known is: $K^\infty(E) = \cap_n K^n(E)$

   Event E is commonly known in state ω if $\omega \in K^\infty(E)$

An alternative definition of common knowledge:

1. An event $F \subseteq \Omega$ is self-evident between players A and B if for all $\omega \in F$, we have $P_i(\omega) \subseteq F$, $i = A, B$

   (that is, when F occurs, both players know it)
2. An event $E \subseteq \Omega$ is common knowledge between A and B at ω if there is a self-evident event F such that:

   ① $\omega \in F$ ② $F \subseteq E$

Example

This example is from Osborne and Rubinstein, 5.2.

$$\mathcal{P}_A = \left(\left\{\text{⚀}, \text{⚁}\right\}, \left\{\text{⚂}, \text{⚃}, \text{⚄}\right\}, \left\{\text{⚅}\right\}\right).$$

$$\mathcal{P}_B = \left(\left\{\text{⚀}\right\}, \left\{\text{⚁}, \text{⚂}, \text{⚃}\right\}, \left\{\text{⚄}\right\}, \left\{\text{⚅}\right\}\right).$$

$E = \left\{\text{⚀}, \text{⚁}, \text{⚂}, \text{⚃}\right\}$ does not contain an event that is self-evident for A and B. Hence, it is not common knowledge (in any state) according to the second definition.

Also not common knowledge according to first definition:

$$K_A(E) = \left\{\text{⚀}, \text{⚁}\right\}, \quad K_B(E) = E$$

$$K(E) = \left\{\text{⚀}, \text{⚁}\right\}, \quad K(K(E)) = \left\{\text{⚀}\right\}, \quad K(K(K(E))) = \emptyset.$$

The event $E' = \left\{\text{⚅}\right\}$ is common knowledge in ⚅.

Example: Dirty Faces

Each of three players, A, B and C have their face either dirty (1) or clean (0).
The state is a three-digit number, where the first number corresponds to the state of A's face, second number corresponds to B's and the third to C's.

$$\Omega = \{000, 001, 010, 011, 100, 101, 110, 111\}.$$

Each player only knows state of other players' faces:

$$\mathcal{P}_A = \{\{000, 100\}, \{001, 101\}, \{010, 110\}, \{011, 111\}\}$$
$$\mathcal{P}_B = \{\{000, 010\}, \{001, 011\}, \{100, 110\}, \{101, 111\}\}$$
$$\mathcal{P}_C = \{\{000, 001\}, \{010, 011\}, \{100, 101\}, \{110, 111\}\}$$

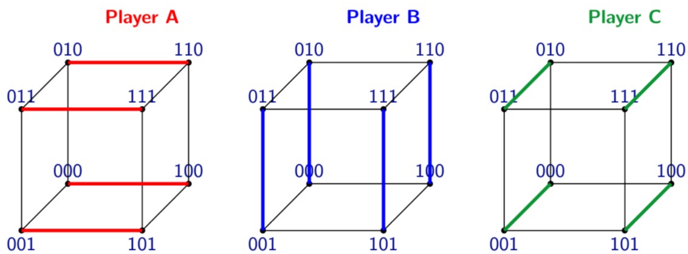

Let $E^* = \Omega \setminus \{000\}$ = "at least one face is dirty".
In which state (if any) E^* is common knowledge?

Answer:

$K(E^*) = \{011, 101, 110, 111\}$ $K(K(E^*)) = \{111\}$ $K(K(K(E^*))) = \emptyset$

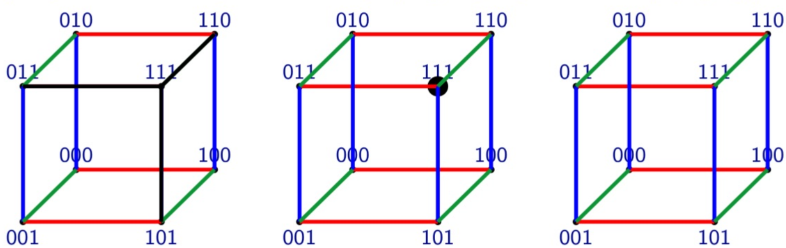

$$\mathcal{P}_A = \{\{000, 100\}, \{001, 101\}, \{010, 110\}, \{011, 111\}\}$$
$$\mathcal{P}_B = \{\{000, 010\}, \{001, 011\}, \{100, 110\}, \{101, 111\}\}$$
$$\mathcal{P}_C = \{\{000, 001\}, \{010, 011\}, \{100, 101\}, \{110, 111\}\}$$

Version II :

Sage enters the room, reports "all faces are clean!" if this is true (i.e., announces if 000 holds), says nothing otherwise.

All players know this, hence each player can distinguish 000 from the state where only his own face is dirty:

$$\mathcal{P}_A = \{\{000\}, \{100\}, \{001, 101\}, \{010, 110\}, \{011, 111\}\}$$
$$\mathcal{P}_B = \{\{000\}, \{010\}, \{001, 011\}, \{100, 110\}, \{101, 111\}\}$$
$$\mathcal{P}_C = \{\{000\}, \{001\}, \{010, 011\}, \{100, 101\}, \{110, 111\}\}$$

Now if E^* occurs, then everyone knows it is true. $K(E^*) = E^*$.
E^* is a public event, therefore $K^{\infty}(E^*) = E^*$, it is self evident.

Dynamic Game : Version II

raise one's hand as soon as sure his face is dirty (time preference)

Sage announces that there is at least one dirty face.

1. If there is exactly one dirty face, then the player with the dirty face knows it and immediately raises hand (due to costs of delay).
2. If no one raised hand, then it is common knowledge that all players know that there are at least two dirty faces. If there are exactly two dirty faces, then the a player with a dirty face observes exactly only one clean face and know that there is more than one clean face, so they infer that they must be the second player with a dirty face. Hence they raise hand in the second round.
3. If no player raised hand in the second round, this means that each player saw two faces, which implies that everyone's face is dirty.

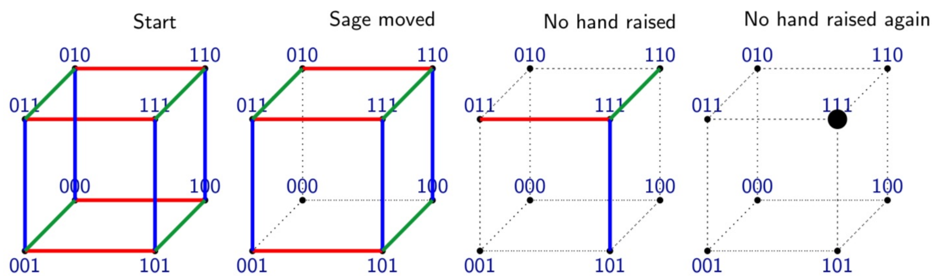

Almost Common Knowledge

Is high order mutual knowledge of an event "close" to common knowledge of the event?

No Agreeing to Disagree

for $i \in N$, $E \subseteq \Omega$ and $p \in [0,1]$, let $E_i^p = \{\omega \in \Omega : \mu(E \mid P_i(\omega)) = p\}$ be the event that player i assigns probability p to E.

1. Suppose it is common knowledge at ω that player 1 assigns probability p to E and that player 2 assigns probability p' to E.
   $\Rightarrow$ Then $p = p'$
2. Interpretation:
   common prior μ: players begin with identical beliefs about Ω
   Then any differences in beliefs is due to differences in information (i.e., $P_1(\omega) \neq P_2(\omega)$), thus they cannot 'agree to disagree'.
3. Belief and Common Knowledge

How can beliefs become common knowledge? Suppose that players can announce their beliefs...

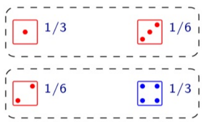

- $\Omega = \{$⚀, ⚁, ⚂, ⚃$\}$.
- Common prior $\mu($⚀$) = \mu($⚃$) = 1/3$, $\mu($⚁$) = \mu($⚂$) = 1/6$.
- Player A knows whether the number is odd or even: $\mathcal{P}_A = (\{$⚀, ⚂$\}, \{$⚁, ⚃$\})$
- Player B knows the color: $\mathcal{P}_B = (\{$⚀, ⚁, ⚂$\}, \{$⚃$\})$

- Take $\mathbb{E} = \{$⚁, ⚂$\}$ and suppose that the state is $\omega =$ ⚁.
- A announces that her posterior belief that E holds is $1/3$.
- B announces that his posterior belief that E holds is $1/2$.
- A concludes that E holds with probability 1 and announces it.
- B concludes that E holds.
- Check that if true state was ⚂ then the process leads to common knowledge that the posterior probability of E is $1/3$.

4. Posterior is not Common Knowledge

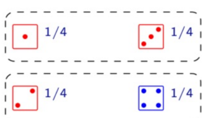

- $\Omega = \{$⚀, ⚁, ⚂, ⚃$\}$.
- Common prior $\mu($⚀$) = \mu($⚁$) = \mu($⚂$) = \mu($⚃$) = 1/4$.
- Player A knows whether the number is odd or even: $\mathcal{P}_A = (\{$⚀, ⚂$\}, \{$⚁, ⚃$\})$
- Player B knows the color: $\mathcal{P}_B = (\{$⚀, ⚁, ⚂$\}, \{$⚃$\})$

- Take $\mathbb{E} = \{$⚀, ⚁$\}$.
- A's belief that E holds is $1/2$ in every ω.
- In $\omega \in \{$⚀, ⚁, ⚂$\}$ player B assigns belief $2/3$ to E.
- In $\omega =$ ⚃ player B assigns belief 0 to E.
- In $\omega =$ ⚀ posteriors are mutually known.
- But there is no state where player 2's posterior is common knowledge. Indeed, 1 and 2's posteriors are not equal. Disagreement is common knowledge.

Epistemology

Epistemic conditions for solution concepts:
(what players need to know (e.g. about each others' rationality) in order to play)

1. Fix a game: $(N, (A_i)_{i\in N}, (u_i)_{i\in N})$
2. State $\omega \in \Omega$ specifies (describes), for all i:
   - the action taken by player i, $a_i(\omega)$
   - player i's beliefs, probability measure $\mu_i(\omega)$ over A_{-i}
   - what i knows, information set $P_i(\omega)$
3. Say that player i is rational in state ω if he plays a best response given his belief, i.e., $a_i(\omega)$ is best response against belief induced by $\mu_i(\omega)$
4. Assume that beliefs are consistent with knowledge. That is, the support of $\mu_i(\omega)$ is a subset of $\{a_{-i}(\omega') : \omega' \in P(\omega)\}$

Common knowledge of rationality implies rationalizability.
Suppose that in state ω it is common knowledge that players are rational. Then, for all i, $a_i(\omega)$ is rationalizable.

If beliefs are derived from a common prior and players are rational in every state, then we obtain a correlated equilibrium.
If μ is a (full support) common prior, players are rational in every state, $a_i(\omega) = a_i(\omega')$ for all i and $\omega, \omega' \in P_i(\omega)$, then $(\Omega, \mu, (P_i)_i, (a_i)_i)$ is a correlated equilibrium.

Nash equilibrium. Fix state ω, then an action profile $a(\omega)$ is a pure Nash equilibrium if in state ω for all i:
- i knows the other player's actions: $P_i(\omega) \subseteq \{\omega' : a_i(\omega') = a_{-i}(\omega)\}$
- i is rational: $a_i(\omega)$ is the best response to $a_{-i}(\omega)$

Weeks 4:
Dynamic Games

Extensive-Form Games
SPE / PBE / SE
Refinements
- forward induction
- intuitive criterion
- divinity
- Spencian signaling

Definition

Dynamic Games: multi-stage games, games played over time.

static game "strategic form" ⟶ dynamic game "extensive form" { commitment vs time consistency; when opponent's move unexpected

Models where players manipulate each others' information over time (signaling, jamming, reputation...) are prevalent in economics.

Example: the 'Entry Deterrence' game

A monopoly (called I) faces a potential entrant (E). If the entrant stays out then it receives 0, and the incumbent gets 2 units of utility.
If the entrant enters then the incumbent can either fight or acquiesce. The former yields (-1,-1), the latter (1,0) for E and I, respectively.

representation: game tree

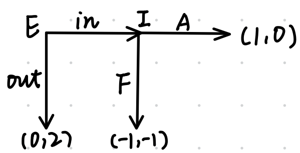

$\xrightarrow[\text{form}]{\text{strategic}}$

	F	A
In	-1, -1	<u>1</u>, <u>0</u>
Out	<u>0</u>, <u>2</u>	0, <u>2</u>

NE: (In, A), (Out, F), (Out, $pF+(1-p)A$) with $p \geq \frac{1}{2}$

1. Backward Induction
2. Subgame Perfect Equilibrium (SPE): only (In, A)
   SPE rules out Nash equilibria sustained by non-credible threats

Extensive-Form Games

- **Players, set N** Nature can be one of the players
- **Histories, set H** A history h is a sequence of movies to a given point in time
- **Available actions, set $A(h)$** terminal histories $Z \subseteq H$ are such that $\forall z$, $A(z) = \phi$
- **Player assignment function, $P(h) \in N$** moves at $h \in H \setminus Z$; if $P(h) =$ Nature, there is a probability distribution over $A(h)$
- **Information sets I** a collection of disjoint subsets of H whose union is H (i.e., a partition of H); for $h, h' \in I$, $P(h) = P(h')$ and $A(h) = A(h')$
- **Payoffs $\{u_i(z)\}_{i \in N}$** at every terminal history $z \in Z$

* Beliefs: probability distribution over nodes in an information set

Representation and Conversion to normal form

1. Lossless graphical representation: "Game tree +"

- Directed graph with a single initial node; edges represent moves
- Probabilities on edges representing Nature (chance) moves
- Nodes that the deciding player cannot distinguish (which are in the same information set) are connected by a dashed line.

2. Lossy conversion: Normal (or strategic form)

- A strategy is a player's complete plan of action, listing a move at every information set
- A strategy profile (one strategy from each player's set) determines an outcome (payoffs)
- Different extensive-form games may have same normal form.

(*) The number of a player's strategies = $\prod$ the numbers of actions available at each of his information sets.

Example: Simultaneous vs. Sequential moves

Alice (P1, row) and Bob (P2, column) play Battle of the Sexes.

	Cafe	*Pub*
Cafe	$\underline{5}$, $\underline{3}$	2, 2
Pub	0, 0	$\underline{3}$, $\underline{5}$

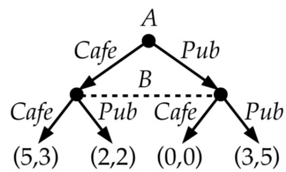

Alice first, observed by Bob (**singleton info sets** in extensive form!):

	CafeCafe	*CafePub*	*PubCafe*	*PubPub*
Cafe	$\underline{5}$, $\underline{3}$	$\underline{5}$, $\underline{3}$	$\underline{2}$, 2	2, 2
Pub	0, 0	3, $\underline{5}$	0, 0	$\underline{3}$, $\underline{5}$

Bob's strategy *XY* means, "*X if Alice plays Cafe, Y if she plays Pub*".

Example: Reduced normal form

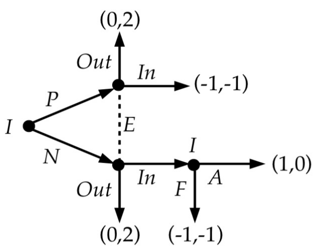

"Poison pill" in entry deterrence

P replaces PA, PF (equivalent)

	PA	*PF*	*NA*	*NF*
In	-1,-1	-1,-1	$\underline{1}$, $\underline{0}$	-1,-1
Out	$\underline{0}$, $\underline{2}$	$\underline{0}$, $\underline{2}$	0, $\underline{2}$	$\underline{0}$, $\underline{2}$

Solution Concepts

1. Every extensive-form game can be converted into its (reduced) normal form. Information about dynamics may be lost.
2. Every finite, normal-form game has a (mixed) Nash equilibrium.
3. A Nash equilibrium may fail time consistency: strategies may fail to best respond in some out-of-equilibrium continuation.

Refinements of Nash equilibrium in dynamic games:

1. Subgame-perfect equilibrium (SPE)
2. Perfect-Bayesian equilibrium (PBE)
3. Sequential equilibrium (SE)

SPE: { Definition: NE in every subgames (non-singleton nodes)
Method: Backward Induction

Subgame-Perfect Equilibrium (SPE)

finit games, complete info

Subgame: the continuation of the game (sub-tree) after a specific history (node), such that no information set is 'broken up'.

SPE: a strategy profile that is Nash equilibrium in every subgame.

Theorem (Kuhn, 1953): Every finite game in which players have perfect recall (do not forget own prior moves, not the same as perfect info!) has a subgame perfect equilibrium.

In simultaneous games, there are no subgames, thus all Nash are SPE

SPE + beliefs are specified by Bayesian Rules

Perfect Bayesian Equilibrium (PBE)

PBE: a SPE strategy-profile and beliefs for all players at their information sets, satisfying:

1. Sequential Rationality: Each player chooses optimally given his beliefs at each information set and the others' eqm strategies.
2. Bayesian beliefs: Beliefs are computed based on equilibrium strategies via Bayes' rule whenever possible. No restrictions on beliefs at 'unreached' information sets in two-player games.

Theorem (Fudenberg & Tirole, 1991): Every finite game has a PBE.

On-path beliefs generated by equilibrium strategies via Bayes' rule.
Off-path beliefs chosen by the modeler to support the equilibrium.

Sequential Equilibrium

Assessment: call (σ, μ) a potentially mixed strategy profile σ coupled with beliefs μ (for all players, at all information sets)

1. Assessment (σ, μ) is consistent if there is a sequence (σ^m, μ^m) that converges to (σ, μ) as $m \to \infty$, such that σ^m is a completely mixed strategy profile and μ^m is computed from σ^m via Bayes' rule.
2. An assessment (σ, μ) is a sequential equilibrium (SE) if it is sequentially rational (as in the definition of PBE) and consistent

> **Theorem (Kreps & Wilson, 1982):** Every finite game has a sequential equilibrium. All SE are PBE, and all PBE are SPE.

Off-path beliefs must be the limit of Bayesian beliefs generated by 'trembles', otherwise same as PBE

First-mover advantage: $\dfrac{d\left(\frac{d\pi_i}{dx_i}\right)}{dx_j} = \dfrac{\partial \pi_i}{\partial x_i \partial x_j}$

My marginal profit changes with your input efforts

$\begin{cases} \text{Bertrand: } BR_i = f(p_j) \ (+): \text{ second-mover} \\ \text{Stackelberg: } BR_i = f(q_j) \ (-): \text{ first-mover} \end{cases}$

Example: Reinhard Selten's 'horse'

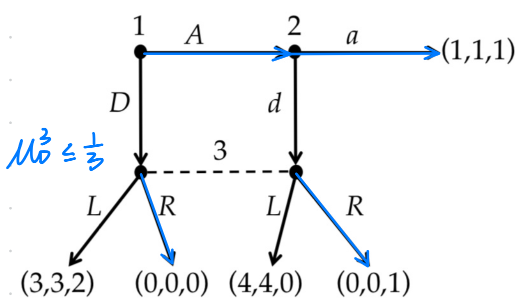

L:

	a	d
A	1,1,1	4,4,0
D	3,3,2	3,3,2

R:

	a	d
A	1,1,1	0,0,1
D	0,0,0	0,0,0

pure NE: (D, a, L), (A, a, R)

SPE: (D, a, L), (A, a, R)

PBE: (A, a, R),

player 3's belief: $\mu_D \le \frac{1}{3}$

$v_3(L) = 2\mu_D < (1-\mu_D) = v_3(R)$

$(A, a, pL + (1-p)R)$,

$\mu_D = \frac{1}{3}$, $p \le \frac{1}{4}$

"off-path mixing" by player 3.

(*) (D, a, L) not PBE because player 2 chooses d and induces outcome (A, d, L)

Example: PBE/SE Application (Centipede with doubt)

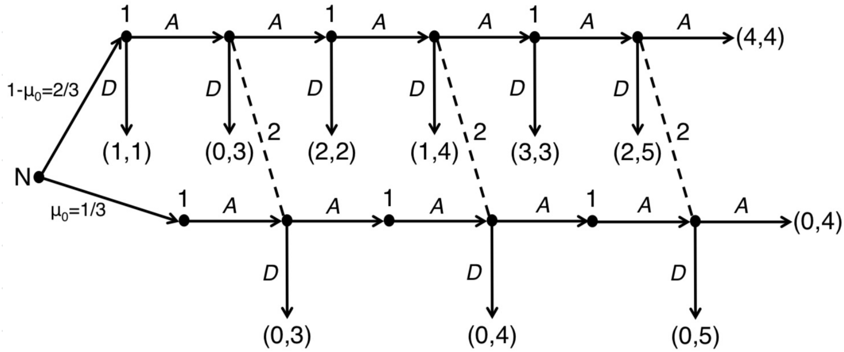

Suppose P2 (she) believes there is a $\mu_0 = 1/3$ chance that P1 (he) is "irrational," i.e., a *commitment type* playing A at every node.

P2's prior is commonly known; P1 knows if he is rational or not.

⇒ the unique PBE:

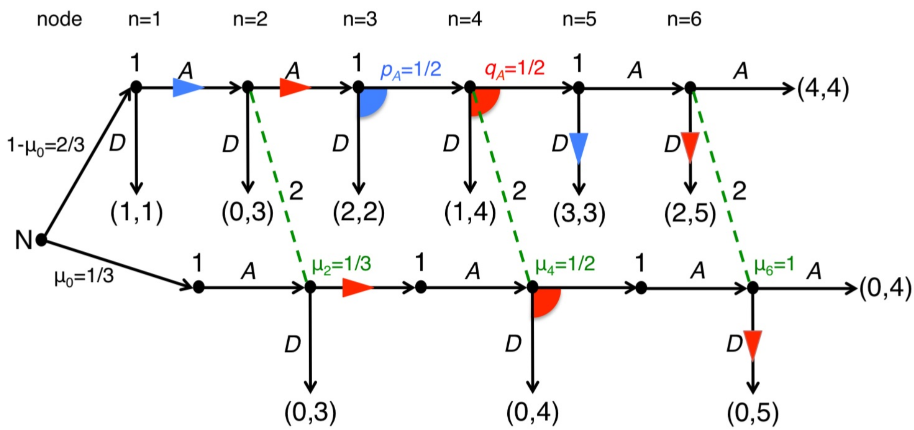

"Sane" P1 plays A to build a reputation of being "crazy". P2 goes along (rationally); players burst the bubble near the end.

In the unique PBE of the "Centipede game with doubt" both players play A in early periods, provided μ0 is not too small.
With 14 < μ0 < 12, we get A in all but the last four periods. The smaller μ0 the longer is the endgame where players mix.
The PBE is unique — the proof of this is a bit more involved.

Extensive-form Refinements

Forward Induction

thinking about the rational behinds the seemingly irrational point: "why we are here"

BoS with option X

Consider BoS where Alice has an attractive outside option:

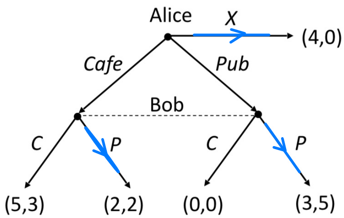

	C	P
Cafe	5, 3	2, 2
Pub	0, 0	3, 5
X	4, 0	4, 0

(X, P): Nash, SPE, PBE ($\mu_{pub} \geq \frac{1}{6}$), SE
If Bob believes $\mu_{pub} \geq \frac{1}{6}$, then he plays P, so Alice plays X. (PBE)

(Cafe, C). Nash, PBE ($\mu_{Cafe} = 1$)

(*) SE supports (X, P)

for $m = 1, 2, \cdots$ let $\sigma_A^m = \frac{1}{m^2}\text{Cafe} + \frac{1}{m}\text{Pub} + (1 - \frac{1}{m^2} - \frac{1}{m})X$, $\sigma_B^m = \frac{1}{m}\text{Cafe} + (1 - \frac{1}{m})\text{Pub}$

Using Bayes' rule, Bob's belief that Alice has played Pub conditional on having played either Cafe and Pub "by mistake" is:

$$\mu_{pub}^m = \Pr[\text{Pub} \mid \text{not } X] = \frac{m}{m+1}$$

As $m \to \infty$, $\sigma_A^m \to X$, $\sigma_B^m \to P$, $\mu_{pub}^m \to 1$. Therefore, (X, P) with $\mu_{pub} = 1$ is SE.

But: Alice guarantees 4 by playing X. By not playing X, she must be expecting more. Only Cafe can yield more for her. Hence $\mu_{pub} = 0$?

Intuitive Criterion

Beer-Quiche game (Cho & Kreps, 1987)

A pub-goer (P1) orders either Beer or Quiche. A bully (P2) sees this and either fights him or not. P2 wants to fight if P1 is weak but not when P1 is strong. Only P1 knows his type; the prior probs are 20% weak, 80% strong.
Weak P1 likes Quiche, strong P1 likes Beer; neither type of P1 wants a fight.

players: Nature, P1, P2

Histories & Actions:

0. Nature picks P1's type, weak (P1w, 20%) or strong (P1s, 80%)
1. P1 observes his type, chooses either Beer or Quiche.
2. P2 observes P1's action but not his type, picks Fight or Not.

Payoffs:

$\begin{cases} \text{P1 gets 1 for consuming favourite item, 2 for avoiding a fight} \\ \text{P2 gets 1 for fightening weak or not fighting strong P1, 0 otherwise} \end{cases}$

Solution:

Beer-Quiche is extensive form.

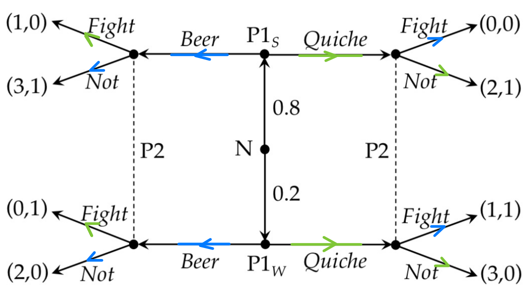

"Pooling" PBE where both types of P1 have beer.

"Pooling" PBE where both types of P1 have quiche

	FF	FN	NF	NN
BB	1, 0, 0.2	0, 1, 0.2	3, 2, 0.8	3, 2, 0.8
BQ	1, 1, 0.2	1, 3, 0.0	3, 1, 1.0	3, 3, 0.8
QB	0, 0, 0.2	2, 0, 1.0	0, 2, 0.0	2, 2, 0.8
QQ	0, 1, 0.2	2, 3, 0.8	0, 1, 0.2	2, 3, 0.8

→ Player 2's choice WZ: "W if P1 chose Beer, Z if P1 chose Quiche"

↓ player 1's choice XY: "X if strong, Y if weak"

Payoffs written in each cell in the order P1s, P1w, P2.
pure NE: (BB, NF), (QQ, FN)

PBE beliefs for (BB, NF):
P2: Pr[P1s | Beer] = 0.8 ⇒ P2 doesn't fight after Beer
free to choose Pr[P1s | Quiche] $\le \frac{1}{2}$ ⇒ P2 fights after Quiche

PBE beliefs for (QQ, FN):
P2: Pr[P1s | Quiche] = 0.8 ⇒ P2 doesn't fight after Quiche
free to choose Pr[P1s | Beer] $\le \frac{1}{2}$ ⇒ P2 fights after Beer

P1 playing BQ or QB: incentive to deviate

Sequential Equilibrium construction:

for pooling on Beer

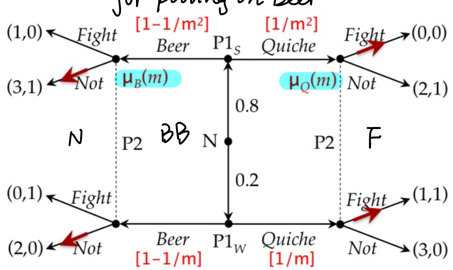

for pooling on Quiche

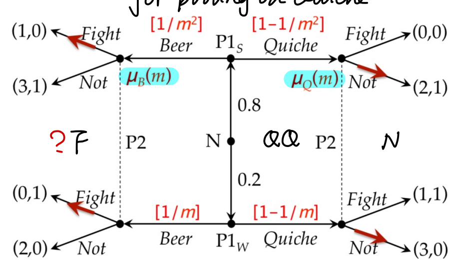

As $m \to \infty$, $\mu_B(m) \to 0.8$, $\mu_Q(m) \to 0$ (in both case)
Counterintuitive!

P1W cannot possibly gain by deviating to Beer (his equilibrium payoff is his maximal feasible one). Then why does P2 put at least 50% belief on this type upon observing a deviation to Beer?

Conclusion (Intuitive Criterion)

upon observing a deviation, put zero weight on types whose equilibrium payoff exceeds all possible payoffs from deviating (assuming opponent's action is rational for some beliefs)

Result by Cho & Kreps: In the Beer-Quiche game, the PBE/SE outcome selected by the Intuitive Criterion is Pooling on *Beer*.

Divinity and Stategic Stability

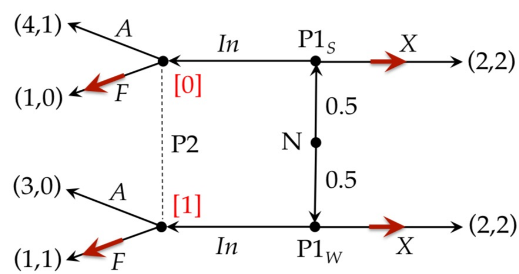

Pooling PBE/SE with (X, F) and $\mu_S = 0$ satisfies the Intuitive Criterion because both types of P1 could gain from deviation to *In*.

This equilibrium "feels wrong" as $P1_S$ gains from deviating to *In* for more mixed replies of P2 than $P1_W$ does. Formalized as *Divinity*.

Divinity D1 (Cho & Sobel, 1991): Upon observing the Sender's out-of-equilibrium action, the Receiver must assign 0 beliefs to a Sender type if there is *another type* that gains compared to its eqm payoff for a *larger set* of sequentially rational Receiver responses.

D1, D2, Universal...

Strategic Stability (Kohlberg & Mertens, 1986): A strategy profile σ^* is stable if, for any sequence $\varepsilon^m = \{\varepsilon^m(s_i) : s_i \in S_i, i \in N\} > 0$, with $\varepsilon^m \to 0$ as $m \to \infty$, there exists a sequence $\sigma^m \to \sigma^*$ such that for all i, σ_i^m is a best response to σ_{-i}^m subject to $\sigma_i^m(s_i) \geq \varepsilon^m(s_i)$.

Trembling-Hand Perfection → Strategic Stability

That is, σ^* is stable if it is the limit of Nash equilibria subject to any type of trembles as the trembles vanish. Robust, may fail to exist.

Spencian Signal

Thorstein Veblen (1899): conspicuous consumption, Veblen goods
(waste to display status) (demand increases with prizes)

Potlatch: feast to display wealth

Michael Spence (1973, 2001 Nobel): formal model

Spence's (1973) signaling game

1. Nature picks worker's type $\theta \in \{\theta_L, \theta_H\}$; $Pr(\theta = \theta_H) = \lambda$
2. Worker observes θ and choose education level $e \geq 0$
3. (At least two) firms observe e but not θ and set wage w.
4. Worker accepts at most one offer.

$\begin{cases} \text{Worker's payoff: } w - e/\theta, \text{ zero outside option} \\ \text{Firm's payoff: } \theta - w \text{ if it employs the worker, } 0 \text{ if it does not.} \end{cases}$

PBE: $(e(\theta)_{\theta=\theta_L,\theta_H}, w(e)_{e\geq 0}, \mu(e)_{e\geq 0})$, where $\mu(e) = Pr[\theta = \theta_H | e]$

Find all PBEs $\Rightarrow$ apply the Intuitive Criterion $\Rightarrow$ select a unique outcome

Example of a pooling equilibrium
(both types set $e(\theta) = e^* \Rightarrow \mu(e^*) = \lambda$)

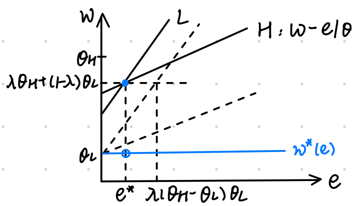

$w(e^*) = \lambda\theta_H + (1-\lambda)\theta_L$

for other $e' \neq e^*$, let $\mu(e') = 0$, therefore $w(e') = \theta_L$

check $0 \leq e^* \leq \lambda(\theta_H - \theta_L)\theta_L$, otherwise θ_L deviates to $e' = 0$

Example of a separating equilibrium

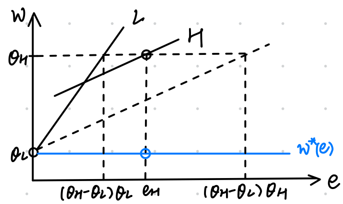

Weeks 5 :
Bargaining, Evolutionary Games

Bargaining : Nash's Axioms & Risk Aversion
Noncooperative Bargaining Model
Evolutionary state strategies (ESS)
- the Replicator Dynamic
- (Stochastic) Best Reply Dynamics

Bargaining

Cooperative Game Theory: specify reasonable conditions solutions should satisfy without writing down full details of bargaining process

Non-cooperative Game Theory: write a specific bargaining game and analyze using tools studied earlier in term

X = set of possible agreements

D = disagreement

player i has utility function $u_i: X \cup D \to \mathbb{R}$

Disagreement Point: $d = (d_1, d_2)$, $d_i = u_i(D)$

$U = \{(u_1(x), u_2(x)) : x \in X\} \cup d$

Assumption: $\begin{cases} d \in U \\ \exists (v_1, v_2) \in U \text{ s.t. } v_1 > d_1,\ v_2 > d_2 \\ U \text{ is convex, closed and bounded} \end{cases}$

A bargaining problem is a pair (U, d)

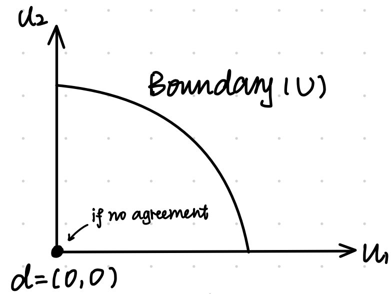

the Bargaining Set

Nash's Axioms

Let F be a function which assigns a unique outcome $F(U,d) \in U$ to every bargaining problem (U, d)

1. WP (Weak Pareto Efficiency)

   if $u = F(U,d)$ there does not exist $(v_1, v_2) \in U$ such that $v_1 \geq u_1$ and $v_2 \geq u_2$ (≥ 1 strict)

2. Sym (Symmetry)

   (U,d) is a symmetry problem if $d_1 = d_2$ and $(v_1, v_2) \in U \Leftrightarrow (v_2, v_1) \in U$.

   If (U,d) is a symmetric problem and $u = F(U,d)$ then $u_1 = u_2$.

3. INV (Invariance to Equivalent Payoff Representations)

   Given $\alpha_i > 0$ and β_i let:

   $\begin{cases} u_i' = \alpha_i u_i + \beta_i \\ U' = \{(\alpha_1 u_1 + \beta_1, \alpha_2 u_2 + \beta_2) : (u_1, u_2) \in U\} \\ d' = (\alpha_1 d_1 + \beta_1, \alpha_2 d_2 + \beta_2) \end{cases}$

   Then $u = F(U,d) \Leftrightarrow (\alpha_1 u_1 + \beta_1, \alpha_2 u_2 + \beta_2) = F(U', d')$

4. IIA (Independence of Irrelevant Alternatives)

   If $U' \subseteq U$, $d' = d$ and $F(U,d) \in U'$ then $F(U', d) = F(U, d)$

Nash's Theorem (Nash, 1950)

There is a unique function that satisfies WP, SYM, INV, IIA, namely $v = F(U,d)$ maximizes $(v_1 - d_1)(v_2 - d_2)$ s.t. $(v_1, v_2) \in U$, $v_1 \geq d_1$, $v_2 \geq d_2$

Proof: 1. Nash's solution $\Rightarrow$ Nash's axioms (not necessarily equilibrium)

2. "unique": build up from special cases

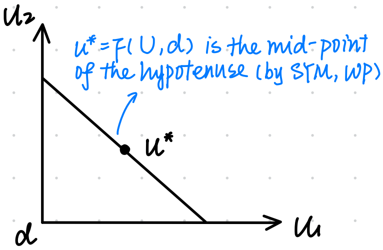

U is an isosceles triangle and $d = (0,0)$

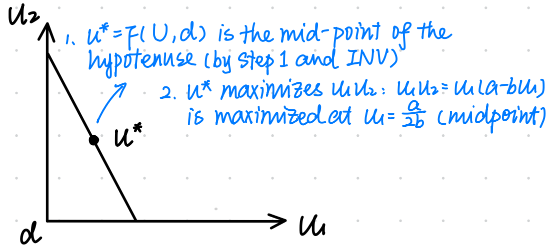

U is a right-angle triangle and $d = (0,0)$

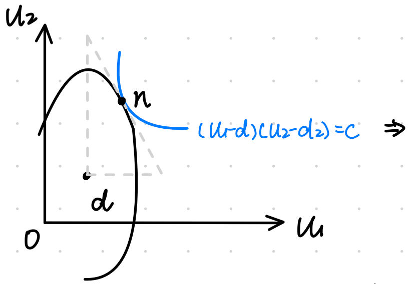

General Convex bargaining problem (U, d)

$\Rightarrow$

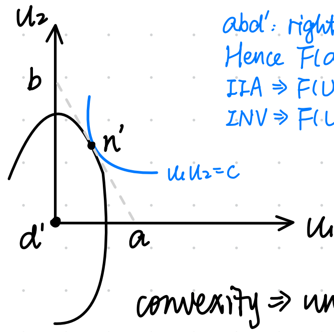

abd': right-angle triangle
Hence $F(abd')$ = Nash solution
IIA $\Rightarrow F(U', d') = n' = N(U', d')$
INV $\Rightarrow F(U, d) = n = N(U, d)$

convexity $\Rightarrow$ uniqueness

Effects of Risk Aversion

Utility function $u(x)$ displays constant relative risk aversion if $u(x) = x^e$, $0 < e < 1$, $e = 1$ corresponds to risk neutrality.

Example: division of pie of size 1 (agent 1 is risk-neutral and agent 2 risk averse with $e < 1$)

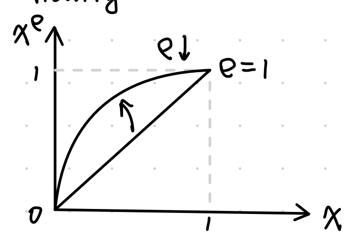

Nash solution: $\max u_1(x_1)u_2(x) = x(1-x)^e$
Equivalently: $\max \ln x + e\ln(1-x)$
FOC: $\frac{1}{x} - e\frac{1}{1-x} = 0 \Rightarrow x = \frac{1}{1+e} > \frac{1}{2}$
—— Nash solution gives more of the pie to the risk neutral player (less risk averse)

Non-cooperative Models of Bargaining

Discuss a sequence of models of bargaining process:

1. the Nash Demand Game
2. the Ultimatum Game
3. Offer-counteroffer
4. Offer-counteroffer with discounting
5. Offer-counteroffer with breakdown
6. Infinitely-repeated versions

1. The Nash Demand Game

players $i = 1, 2$
utility functions $u_i : [0,1] \to \mathbb{R}$
strategies: $x_1, x_2 \in [0,1]$
outcome: (x_1, x_2) if $x_1 + x_2 \leq 1$; $(0,0)$ otherwise

Equilibrium: $(x, 1-x)$, $(1,1)$

2. The Ultimatum Game

player 1 proposes a division (x_1, x_2) s.t. $x_1 + x_2 \leq 1$, $x_1, x_2 \geq 0$
player 2 accepts (Y) or rejects (N)
outcome: (x_1, x_2) if Y; $(0,0)$ if N

SPE: player 2 accept any offer, player 1 offers $(1,0)$

3. Offer-Counteroffer

Round 1: player 1 makes offer, player 2 says Yes or No
Round 2: if player 2 said No, she makes offer and player 1 responds

SPE: player 1 offers $(0,1)$ and player 2 accepts

4. Offer-Counteroffer with Discounting

player i discounts future payoffs by $\delta_i < 1$ per period
offer-counteroffer played once per period

SPE: player 1 offers $(1-\delta_2, \delta_2)$ and player 2 accepts

5. Offer-Counteroffer with Breakdown

No discounting: $\delta_1 = \delta_2 = 1$

If rejection occurs in period 1, bargaining breaks down with prob α

SPE: player 1 offers $(\alpha, 1-\alpha)$ and player 2 accepts

6. Infinitely Repeated Offer-Counteroffer with Breakdown

No discounting: $\delta_1 = \delta_2 = 1$

If rejection occurs in period 1, bargaining breaks down with prob α

players alternate in making offers until acceptance or breakdown

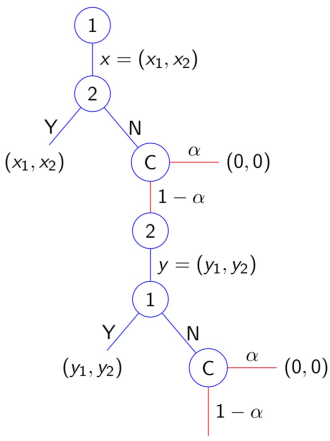

Define $\hat{x}(\alpha) = (\hat{x}_1(\alpha), \hat{x}_2(\alpha))$ and

$\hat{y}(\alpha) = (\hat{y}_1(\alpha), \hat{y}_2(\alpha))$ by:

$$\begin{cases} u_2(\hat{x}_2(\alpha)) = (1-\alpha)\, u_2(\hat{y}_2(\alpha)) \\ u_1(\hat{y}_1(\alpha)) = (1-\alpha)\, u_1(\hat{x}_1(\alpha)) \end{cases}$$

Note that: $\hat{x}_1(\alpha) > \hat{y}_1(\alpha)$ and $\hat{x}_2(\alpha) < \hat{y}_2(\alpha)$

Stationary Subgame Perfect Equilibrium:

$$\begin{cases} \text{player 2 accepts } x = (x_1, x_2) \text{ iff: } \\ \quad u_2(x_2) \geq (1-\alpha)\, u_2(\hat{y}_2(\alpha)) \\ \text{player 1 accepts } y = (y_1, y_2) \text{ iff: } \\ \quad u_1(y_1) \geq (1-\alpha)\, u_1(\hat{x}_1(\alpha)) \end{cases}$$

$\Rightarrow$ player 1 always offer $\hat{x}(\alpha)$, player 2 $\hat{y}(\alpha)$

$\Rightarrow$ Outcome: $\hat{x}(\alpha)$ in period 1.

proposition 1.

when $\alpha \to 0$, and breakdown outcome is the disagreement point (d_1, d_2), the SPE offer $\hat{x}(\alpha)$ is very close to the Nash bargaining solution:

$$\lim_{\alpha \to 0} \hat{x}(\alpha) = (n_1, n_2)$$

where (n_1, n_2) maximizes $(u_1 - d_1)(u_2 - d_2)$ over all $u \in U$

proof: cross multiply: $(1-\alpha)\, u_1(\hat{x}_1(\alpha))\, u_2(\hat{x}_2(\alpha)) = (1-\alpha)\, u_2(\hat{y}_2(\alpha))\, u_1(\hat{y}_1(\alpha))$

$\Rightarrow (u_1(\hat{x}_1(\alpha)), u_2(\hat{x}_2(\alpha)))$ and $(u_1(\hat{y}_1(\alpha)), u_2(\hat{y}_2(\alpha)))$ lie on a level curve (a hyperbola) of the function $u_1 u_2$

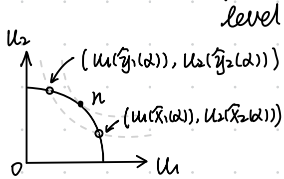

$|\hat{x}(\alpha) - \hat{y}(\alpha)| \to 0$ as $\alpha \to 0$

$\Rightarrow$ both $\hat{x}(\alpha)$ and $y(\alpha)$ converge to n

7. Infinitely Repeated Offer-Counteroffer with Discounting

SPE: $\begin{cases} u_2(\hat{x}_2) = \delta_2 u_2(\hat{y}_2) \\ u_1(\hat{y}_1) = \delta_1 u_1(\hat{x}_1) \end{cases}$

proposition 2:

If δ_1, δ_2 very close to 1 the SPE offer $\hat{x}(\delta)$ is very close to the generalized Nash solution:

$$\max \; u_1(x_1)^{\frac{1}{p_1}} u_2(x_2)^{\frac{1}{p_2}}, \text{ where } p_i = -\ln \delta_i$$
$$\text{s.t. } x_1 + x_2 \le 1, \; x_1, x_2 \ge 0$$

proof: If $p_i = -\ln \delta_i$, then $\delta_i = e^{-p^i}$ and so

$$u_2(\hat{x}_2) = e^{-p_2} u_2(\hat{y}_2)$$
$$u_1(\hat{y}_1) = e^{-p_1} u_1(\hat{x}_1)$$

$$\Rightarrow \quad \left[\frac{u_1(\hat{x}_1)}{u_1(\hat{y}_1)}\right]^{1/p_1} = e = \left[\frac{u_2(\hat{y}_2)}{u_2(\hat{x}_2)}\right]^{1/p_2}$$

$$\Rightarrow \quad u_1(\hat{x}_1)^{1/p_1} u_2(\hat{x}_2)^{1/p_2} = u_1(\hat{y}_1)^{1/p_1} u_2(\hat{y}_2)^{1/p_2}$$

$\Rightarrow$ $(u_1(\hat{x}_1), u_2(\hat{x}_2))$ *and* $(u_1(\hat{y}_1), u_2(\hat{y}_2))$ lie on the same level curve of the function $u_1^{1/p_1} u_2^{1/p_2}$

(*) Kalai-Smorodizky (KS) solution

an alternative to Nash bargaining solutions:

1. U: two-player bargaining set
   $d = (0,0)$: disagreement point
2. $u_i^+ = \max u_i \cdot (u_1, u_2) \in U, \; i = 1,2$
3. Among all pairs $(u_1, u_2) \in U$ s.t. $\frac{u_1}{u_1^+} = \frac{u_2}{u_2^+}$, let (u_1^*, u_2^*) be such that $\frac{u_1^*}{u_1^+} = \frac{u_2^*}{u_2^+}$ is a maximum

Properties: 1. not satisfy Nash's independence axiom
2. linear growth of each player's gains as pie grows

Evolutionary Game Theory

Evolutionarily Stable Equilibrium: static idea intended to capture robustness to invasion by mutants (shocks, pertubations)

Explicit dynamics:

1. Replicator Dynamic (deterministic)
2. Stochastic Shocks

Another Definition for ESS:
α^* is ESS iff $\forall \alpha \neq \alpha^*$, there is $\bar{\varepsilon} > 0$ s.t. $\forall \varepsilon \in (0, \bar{\varepsilon})$, we have:
$U(\alpha^*, (1-\varepsilon)\alpha^* + \alpha\varepsilon) > U(\alpha, (1-\varepsilon)\alpha^* + \alpha\varepsilon)$

Biological Framework

- a large population of individuals or organisms
- each endowed with a behavioral strategy (inherited)
- play a given two-person game in randomly assigned pairs
- payoff is an individual's rate of reproduction (Darwinian Fitness)

A two-person symmetric game G:

A = action space

$u(a, a')$ = payoff to an a-player when the opponent plays a'

{ pure strategy: an action $a \in A$
  mixed strategy: a probability distribution over actions, α

{ Monomorphic population: everyone has the same strategy
  Polymorphic population: different individuals, different strategies

Evolutionarily Stable Strategies (ESS):

static concept

"a strategy is ESS"
"NE about strategy portfolio"

A strategy α^* is an ESS if the following two conditions hold:

One Definition {
1. (α^*, α^*) is a symmetric NE of G   also: $U(\alpha^*, \alpha^*) \geq U(\alpha, \alpha^*)$
2. if $\beta \neq \alpha^*$ is a best response to α^* ($U(\beta, \alpha^*) = U(\alpha^*, \alpha^*)$), then $U(\beta, \beta) < U(\alpha^*, \beta)$

$\Rightarrow$ (α^*, α^*) is a strict NE if $\forall \beta \neq \alpha^*$, $U(\beta, \alpha^*) < U(\alpha^*, \alpha^*)$ $\Rightarrow \alpha^*$ is ESS

$\Rightarrow$ mixed NE cannot be strict

Example 1: Dynamics with population growth

	a	b
a	2,2	0,0
b	0,0	1,1

p: proportion playing a $\Rightarrow$ (a,a), (b,b) are ESS

$n(t)$ = size of population at time t (time discrete)

$n_a(t), n_b(t)$ = number of a-players / b-players at t

$\Rightarrow n_a(t+1) = \lambda[2p(t)]n_a(t) + n_a(t)$ $2p(t)$: fitness (expected payoff)

$n_b(t+1) = \lambda[1-p(t)]n_b(t) + n_b(t)$

λ = proportionality factor (depends on time scale)

$\Rightarrow p < \frac{1}{3}$ is the basin of attraction of $p=0$; $p > \frac{1}{3}$ for $p = 1$.
(set of points from which converges to $p=0$)

It could be that A is ESS, but couldn't resist two mutants' invading strategy simultaneously.

Example 2: Hawk-Dove Game

		A	P
p	A	$\frac{v-c}{2}, \frac{v-c}{2}$	$v, 0$
$1-p$	P	$0, v$	$\frac{v}{2}, \frac{v}{2}$

A: aggressive, P: passive

$\begin{cases} v: \text{value of resources} \\ c: \text{cost of fighting} \end{cases}$

$v > c$: (A, A) is a strict NE and therefore an ESS

$v < c$: $(v/c, 1-v/c)$ is a mixed NE and an ESS.

$$U(\alpha^*, \beta) = \tfrac{v}{c}p\tfrac{(v-c)}{2} + \tfrac{v}{c}(1-p)v + (1-\tfrac{v}{c})p.0 + (1-\tfrac{v}{c})(1-p)\tfrac{v}{2}$$

$$U(\beta, \beta) = p^2\tfrac{(v-c)}{2} + p(1-p)v + (1-p)p.0 + (1-p)^2\tfrac{v}{2}$$

$$\Rightarrow U(\alpha^*, \beta) - U(\beta, \beta) = (c/2)(v/c - p)^2 > 0 \text{ whenever } p \neq v/c$$

for any value p, β is best reply to α^*

$U(\alpha^*, \alpha^*) = U(\beta, \alpha^*)$

— 'indifference' condition

Example 3: some games have no ESS

	A	B	C
A	$\frac{1}{2}, \frac{1}{2}$	$-1, 1$	$1, -1$
B	$1, -1$	$\frac{1}{2}, \frac{1}{2}$	$-1, 1$
C	$-1, 1$	$1, -1$	$\frac{1}{2}, \frac{1}{2}$

claim $\alpha^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ is the unique symmetric NE $\Rightarrow$ only possible ESS

Any pure strategy a. $U(a, a) = \frac{1}{2} > U(\alpha^*, a) = \frac{1}{6}$ and is best response to α^*

Conditional Strategies

One approach to defining ESS for asymmetric games is to 'symmetrize' the game

Consider a two-player game C (not necessarily symmetric):

$U_i(\alpha, \beta)$ = payoff to role i when role 1 plays α and role 2 plays β

suppose the ex ante probability of being in each role is $\frac{1}{2}$

$\sigma = (\sigma_1, \sigma_2)$: conditional strategy portfolio

Define the two-person symmetric game with payoff function:

		Intruder A	Intruder P
Owner	A	$\underline{\frac{v-c}{2}}, \underline{\frac{v-c}{2}}$	$\underline{v}, \underline{0}$
	P	$\underline{0}, \underline{v}$	$\underline{\frac{v}{2}}, \underline{\frac{v}{2}}$

$U((\sigma_1, \sigma_2), (\tau_1, \tau_2)) = \frac{1}{2}U_1(\sigma_1, \tau_2) + \frac{1}{2}U_2(\tau_1, \sigma_2)$

$\Rightarrow (A, P)$ is an ESS.

if $c \geq v$.

Proposition: A conditional strategy (σ, σ) is an ESS if and only if σ is a pure strict equilibrium of game C

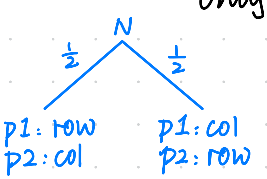

β is an ESS of a symmetrized game iff β is a strict NE in the original game

The Replicator Dynamic

Let G be a symmetric two-person game with m actions

Payoff matrix $A = (a_{ij})$ is $m \times m$

a_{ij} = payoff to row when row plays i and column plays j

$n_i(t)$ = number of individuals programmed to play action a_i at time t

- $n(t) = \sum_{i=1}^{m} n_i(t)$ = total number of individuals at t
- $p_i(t) = n_i(t) / n(t)$ = proportion of population playing action i at time t

$\pi_i(t)$ = average payoff to an i player $= \sum_{j=1}^{m} a_{ij} p_j$

Evolutionary Process:

$n_i(t+1) = n_i(t) + \lambda n_i(t)\pi_i(t) = n_i(t) + \lambda n_i(t)\sum_{j=1}^{m} a_{ij}p_j(t) = n_i(t)(1+\lambda[Ap]_i)$

, where λ = scale factor (small) that depends on time length of periods

$\Rightarrow n(t+1) = n(t) + \lambda \sum_{i=1}^{m}[p_i(t)n(t)\sum_{j=1}^{m} a_{ij}p_j(t)]$

$= n(t) + \lambda n(t)\sum_{i,j} p_i(t) a_{ij} p_j(t) = n(t)(1+\lambda pAp^T)$

, where pAp is the growth rate of the whole population = payoff of p against itself

$p_i(t+1) = \frac{n_i(t+1)}{n(t+1)} = \frac{n_i(t)(1+\lambda[Ap]_i)}{n(t)(1+\lambda pAp)} \approx p_i(t) + \lambda p_i(t)([Ap]_i - pAp^T)$

$\dot{p}_i = \frac{dp_i(t)}{dt} \approx p_i(t+1) - p_i(t)$ → average payoff among all strategies essentially

$\Rightarrow \dot{p}_i = \lambda p_i([Ap]_i - pAp^T)$ — Replicator Equation

proportion playing i grows if it achieves a payoff greater than the averse against the population strategy.

(*) If there are only 2 strategies 1 and 2, played with p and $1-p$ respectively, then the replicator equation for p can be written as:

$\dot{p} = p(1-p)([Ap]_1 - [Ap]_2)$

Intuition: strategy 1 does better than the population average if and only if it does better than strategy 2

Example:

	a	b
a	2,2	0,0
b	0,0	1,1

. Picture of dynamics:

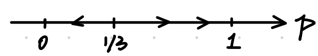

replicator equation: $\dot{p} = p(1-p)(2p-(1-p)) = p(1-p)(3p-1)$

— an ESS is asymptotically stable under the replicate dynamic

Example: Prisoner's Dilemma

	C	D
C	3,3	0,5
D	5,0	1,1

, where C = cooperate, D = defect

infinitely repeated Prisoner's Dilemma ⇒ infinite strategies

Focus on three strategies:

- C: cooperate unconditionally
- D: defect unconditionally
- T: start by cooperating, defect if opponent defected in previous round, otherwise cooperate

⇓ payoff matrix for infinitely repeated PD with stopping probability s ($\leq \frac{1}{2}$)

	C	D	T
C	$3/s$	0	$3/s$
D	$5/s$	$1/s$	$4+1/s$
T	$3/s$	$1/s-1$	$3/s$

In each encounter between the two players:

- First round occurs for sure
- Each subsequent round occurs with probability $1-s$.

Replicator Dynamic in Infinitely Repeated PD

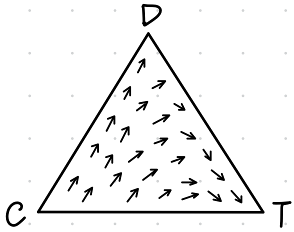

If there is a sufficiently high proportion of T-players initially, the evolutionary process converges to a mixture of T- and C-players

Otherwise: converges to all D-players (likely)

Randomness in the Infinitely Repeated PD

1. Payoffs from a given interaction are variable
2. Matching is not perfectly uniform
3. Number of children is variable
4. Mutations occur

—— Resulting process is a stochastic dynamical system

(*) When the population is large these sources of noise are small in the aggregate, and are well-approximated by a Normal random variable

Best Reply Dynamics stochastic stability

Framework: 1. players from large population randomly matched in pairs to play symmetric game.

2. each period a randomly chosen player is allowed to revise their strategy, who chooses a best response against the current distribution of opponents.

Definition (absorbing states): never leave once reach 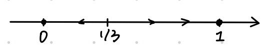

Stochastic Best Reply Dynamics:

$\begin{cases} \text{choose a best reply to the distribution of opponents, } pr = 1-\varepsilon \\ \text{choose a non-best reply, } pr = \varepsilon \end{cases}$

1. Long-run behavior doesn't depend on initial conditions
   probability that i people are playing strategy a in any period converges to a constant π_i, π_i is also the long-run proportion of periods i people play a (follows from theory of Markov Chains)
2. ε small $\Rightarrow$ most time near pure equilibria (all play a or b), mistakes cause to move between occasionally, a harder to escape.
   —— (a, a) is the unique stochastically stable equilibrium

Example: the Stag Hunt (Risk dominance and Equilibrium)

		R	S
(p)	R	3,3	4,0
	S	0,4	5,5

(S,S): Pareto-dominant

(R,R): Risk Dominant

stochastically stable

Main idea:

Cost for all S to all R : e^1

Cost for all R to all S . e^2

compare e^1 and e^2

OR: compare size of basin of attraction

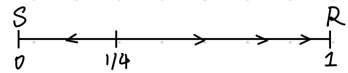

R is best reply if $p \geq \frac{1}{4}$. S is best reply if $p \leq \frac{1}{4}$

(All R has the larger basin of attraction)

—— the risk dominant equilibrium is the unique stochastically stable equilibrium (if the population is large) population same for row and col

---

Relationship between ESS and RD:

$\dot{p} = f(p)$

- p is a steady state if $f(p) = 0$ ($\dot{p} = 0$)
- p is a stable steady state whenever we're near p, we stay near p
- p is asymptotically stable iff whenever we're near p, we end up at p

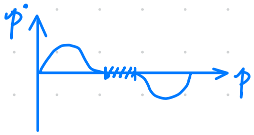

1. Every NE is a steady state of RD
2. Every stable steady state RD is NE
3. Every ESS is a asymptotically stable steady state of RD
   (ESS is robust with disturbance)

RD and Markov Chain:

Argue that the only absorbing states of the process correspond to the pure strategy Nash equilibria of the game.

$\Rightarrow$ $\forall$ state s. we can construct a path so that it ends up at an absorbing state (just wait long enough)

Markov Chain:

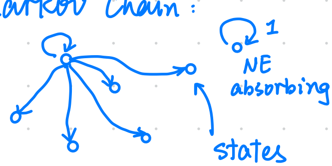

Stochastic stability:

nothing is stable (people make mistake)

"stable" = stay for a long time

Weeks 6:
Information Transmission

Reputation and Stackelberg payoff
Cheap Talk (Crawford-Sobel)

signalling games: proving who you are
reputation building: pretending someone you are not
(a long-run player with commitment types)
communication: how to convince a listener with a conflict of interest

Aristotle in <u>Rhetorics</u> (修辞学) (347BC) identified three means of persuasion:

逻辑 logos (appeal to evidence and deduction): persuasion games
品德 ethos (speaker's credibility): <u>cheap talk</u> games 虚谈博弈
情感 pathos (listener's emotion): not aware of games yet

Reputation

A sequence of short-run players (Et, t = 1,...,T) to play against a single long-run player (I). Each Et decides whether to enter (In) or stay out (Out). Player I chooses fight (F) or acquiesce (A) whenever Et enters.
All earlier actions are observable; the payoffs are below.

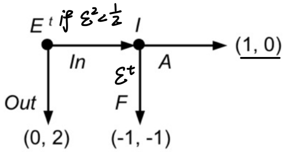

$\forall$ finit T, the unique SPE by backward induction is (In, A)

But: would you enter as E^3 if I fought E^1 and E^2 before you?

Suppose that I is either "sane" (with payoffs given in the game), or "crazy" (commitment type), always playing F against an entrant.
I knows whether he is sane or crazy, Et does not.

Denote the belief of E^t that I is crazy by ε^t.
- $\varepsilon^1 > 0$ is the prior probability of I being crazy.
- PBE (beliefs consistent with play), ε^t is the expected probability of I being crazy at the beginning of stage t.

Find the unique equilibrium (PBE) working backwards (assume T=2).

If E^2 enters then sane I plays A, crazy I plays F.
E^2 enters whenever $\varepsilon^2(-1) + (1-\varepsilon^2)\cdot 1 > 0$, that is, if $\varepsilon^2 < \frac{1}{2}$

Working backward, we now determine PBE at $t=1$.
Claim 1: if $\varepsilon^1 \geq \frac{1}{2}$, E^1 plays Out; however, if E^1 comes In, then player I replies with F. In either case, $\varepsilon^2 = \varepsilon^1$, and so E^2 stays Out.
Claim 2: if $\varepsilon^1 < \frac{1}{2}$ and E^1 plays In then the sane I mixes F and A at $t=1$.
Claim 3: suppose $\varepsilon^1 < \frac{1}{2}$ and E^1 plays In. If I plays A then, at $t=2$, E^2 plays In. If I plays F then E^2 mixes In with probability $\frac{1}{2}$.

Intuition of the claims:

1. **Sane I must not imitate the crazy type 'too perfectly'.**
   If sane I plays F with probability 1 (like the crazy type does), then playing F does not build a reputation for craziness.
2. **When sane I imitates the crazy type at $t=1$ (by playing F), this act would not deter E^2 with probability 1, because then sane I would be tempted to overuse the deterrent.** 威慑

Complete the derivation for $\varepsilon' < \frac{1}{2}$:

1. Sane I mixes at $t=1$ to make E^2 indifferent between In and Out.
   Let q be sane I's probability of playing F at $t=1$. By Bayes' rule:
   $$\varepsilon^2 = \Pr(\text{crazy} \mid F \text{ at } t=1) = \frac{\varepsilon'}{\varepsilon' + (1-\varepsilon')q}$$
   E^2 is indifferent between In and Out iff $\varepsilon^2 = \frac{1}{2}$.
   Hence I needs to set $q = \frac{\varepsilon'}{1-\varepsilon'}$
2. If E' enters then crazy I plays F at $t=1$ with probability 1, and sane I does the same with probability q.
   The total probability that I plays F at $t=1$ is: $\varepsilon' + (1-\varepsilon')q = 2\varepsilon'$
   $\Rightarrow$ E' plays Out iff this probability exceeds $\frac{1}{2}$, that is, iff $\varepsilon' > \frac{1}{4}$

PBE of Chain Store game with reputation, $T=2$

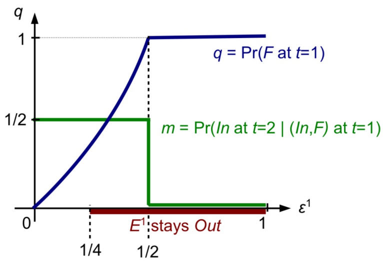

I imitates type crazy to build reputation

E^2 mixes strategy In and Out

Characteristics of the unique PBE for any $T \geq 2$.

1. E' stays Out iff $\varepsilon' > \frac{1}{2^T}$, which goes to 0 as $T \to \infty$
2. If E' comes in, I plays F with probability 1 iff $\varepsilon' > \frac{1}{2^{T-1}}$
3. For all $t > 1$, if I has fought all earlier entrants then E^t stays Out with positive probability. If I ever played A, then E^t comes In for sure.
   — If I ever plays A then he will play A from then on. Otherwise he keeps playing F with positive probability.

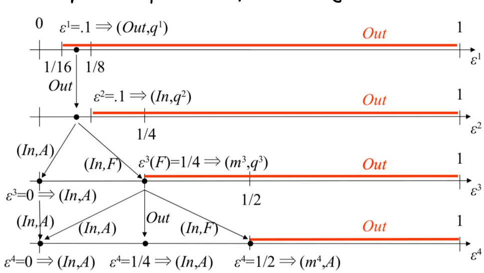

illustration: $T=4$, $\varepsilon' = 0.1$

As T → ∞, the average payoff of the sane, long-lived player I converges to 2, as if he could credibly threaten with F. Possibility of "crazy type" allows I to get his Stackelberg payoff. Same is true in any game between a long-run and a myopic player:

Theorem (Fudenberg and Levine, 1989):
If the long-lived player (P1) has each possible commitment type with positive prior probability then in the unique PBE of the T-times repeated game against short-lived opponents the average payoff of P1 converges to $\max_{a_1} u_1(a_1, BR_2(a_1))$, where $BR_2(a_1)$ is the short-lived player's best response function.

Crawford - Sobel cheap talk game

Vince Crawford and Joel Sobel (1982):

1. Nature picks random state θ ∈ [0, 1]; P1 learns θ, P2 does not.
2. P1 (he) sends a message m ∈ M ; set M is rich, e.g. contains [0, 1].
3. Having observed m, P2 (she) picks y ∈ [0, 1].

Conflict: P2 wants to pick y = θ; P1 wants her to pick a bigger y.
Formally: u2 = –(θ – y)2 and u1 = –(θ + b – y)2 with b > 0.

Application / interpretation:
P2 is prime minister, P1 is expert;
P1 knows optimal policy, θ. P1 has known bias, b > 0. (Ideology, self-interest, etc.)
Advice has no direct payoff implication unlike in signaling.

Equilibrium (PBE):

1. P1's strategy is $S_1 : [0,1] \to \Delta(M)$ = mix over M
2. In equilibrium P2 replies by $S_2(m) = E[\theta \mid S_1(\theta) = m]$
3. $S_1(\theta)$ is a mix over m's maximizing $-(\theta + b - S_2(m))^2$

Messages and Induction.

We say that an action y is induced in state θ if some message m with s2(m) = y is sent in state θ with positive probability.
An action y is induced if it is induced in some state.
Messages have no meaning; only induced (re)actions matter.

Theorem (V. Crawford and J. Sobel, 1982):
In any PBE the total number of induced y actions is finite.

Example: uniform θ, two induced actions

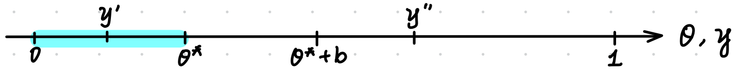

Equilibrium y' is induced in states $\theta < \theta^*$ and action y'' at $\theta > \theta^*$.

P2's best response: $y' = \frac{\theta^*}{2}$, $y'' = \frac{1+\theta^*}{2}$

P1's rationality: $\theta^* + b = \frac{y' + y''}{2}$

$\Rightarrow \theta^* = \frac{1}{2} - 2b$. Two-action PBE iff $b \leq \frac{1}{4}$

Proof that the number of induced action is finite:

In state θ, player P1's ideal point (reaction y) is $\theta + b$ and P2's is θ.

Let $y' < y''$ be two induced actions and let $\theta^* = (y' + y'')/2 - b$.

P1 prefers y'' over y' if $\theta > \theta^*$, hence y' is not induced in $\theta > \theta^*$.

P2 knows this, therefore in equilibrium $y' \leq \theta^*$.

We conclude that $y' \leq \theta^* < \theta^* + b = (y' + y'')/2 < y''$.

The distance between any two induced actions exceeds b, therefore there are fewer than $1/b$ induced actions.□

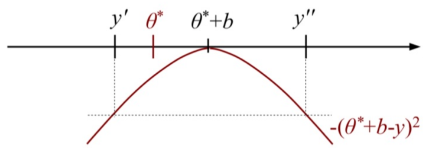

Discussion

- Talk is cheap, yet it can be informative
- Infinitely many states, but only finitely many induced actions
  (greater bias $\Rightarrow$ fewer actions can be induced in equilibrium)

Ex ante, both players prefer equilibria with more communication. (Expected quadratic loss = "variance" is smaller on a finer partition.

Is this a good explanation for polarized/ simplified/ brief communication?

Directions pursued in the literature

1. (partially) verifiable messages

   Messages with meaning (not all m can be sent)

   If θ is provable and P2 has an action that all P1 types dislike, then there is a PBE with full revelation: the highest-type P1 must prove himself, all lower types compelled to do the same. Falls apart if message is observed with noise.

2. Communication over multiple issues

Multi-dimensional state, $\theta = (\theta_1, \cdots \theta_n)$ with $\sum_{i=1}^{n} \theta_i \leq 1$

Comparative equilibria exist in which P1 reports a ranking of dimensions (e.g., “θ1 > θ3 > θ2”) but not the levels. As n → ∞ this amounts to full revelation (Chakraborty & Harbaugh, 2007).

Game:
1. Nature picks $\theta_1, \cdots \theta_k \in [0,1]$, distributed i.i.d.
2. P1 sends a message $m \in [0,1]^k$
3. Having observed m, P2 (she) picks $y \in [0,1]^k$

Payoffs $u_1 = -\sum_{k=1}^{K} (\theta_k + b - y_k)^2$, $u_2 = -\sum_{k=1}^{K} (\theta_k - y_k)^2$; $b > 0$

⇒ if b is small then P1 and P2 can communicate via finitely many messages along each dimension via Crawford & Sobel

⇒ No informative 'Crawford-Sobel partition equilibrium' if b is large

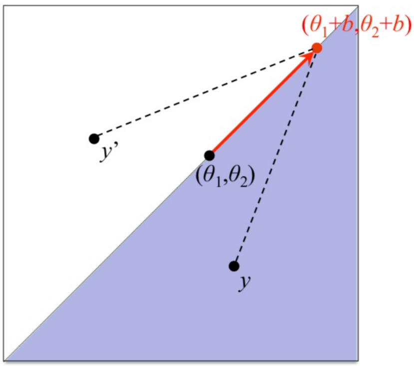

Proposition (Chakraborty & Harbaugh, 2007): For all $b > 0$, there is an equilibrium in which P1 induces $y = E[(\theta_1, \theta_2)|\theta_1 > \theta_2]$ if $\theta_1 > \theta_2$ and $y' = E[(\theta_1, \theta_2)|\theta_1 < \theta_2]$ otherwise.

3. Communication with multiple experts

⇒ harder for experts to bias decisions (Battaglini 2002, Ambrus & Takahashi 2007)

Game:
1. Two experts (P1, P2) know the multi-dimensional state, $\theta \in \Theta \subseteq \mathbb{R}^k$
2. P1 and P2 simultaneously send messages $m_1, m_2 \in \mathbb{R}^k$
3. Receiver (decision maker P0) sets $y \in \mathbb{R}^k$

Ideal reaction of P0: $y = \theta$; the experts have known biases

—— How much does P0 gains by comparing the experts' reports?

Let $B_i(\theta) = \{y \in \mathbb{R}^k : \text{Expert } i \text{ prefers } y \text{ to } \theta \text{ in state } \theta\}$

Theorem (Battaglini, 2002; Ambrus & Takahashi, 2007):
The cheap talk game with multiple experts has a perfect Bayesian equilibrium such that $y = \theta$ is induced in every state θ if, and only if, for all $\theta', \theta'' \in \Theta$ there exists $\hat{y} \in \Theta$ such that $\hat{y} \notin B_1(\theta'') \cup B_2(\theta')$.

Example: Three-way allocation of a budget

(Expert 1 biased towards Politics, expert 2 biased towards Philosophy. $B_i(\theta)$ is the set of policies that i prefers to the policy $y=\theta$ in state θ.)

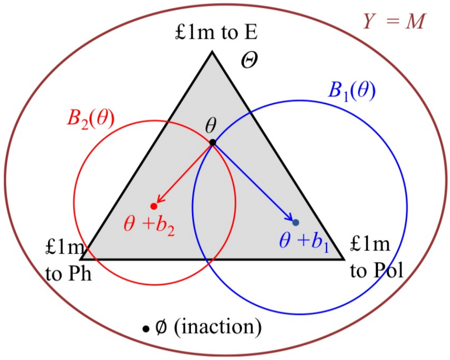

To get the truth from the experts is possible if $\forall \theta', \theta'' \in \Theta$, $\exists \hat{y} \in \Theta$ s.t. $\hat{y} \notin B_1(\theta'') \cup B_2(\theta')$.

If $\theta' = \theta''$, one of them is lying. P0 can punish both without knowing who's lying by setting $\hat{y} \notin B_1(\theta'') \cup B_2(\theta')$

Alternative: ask P1 the correct split between Econ and Phil and P2 about the split between Econ and Pol – neither will lie. you infer θ.

4. **Communication over time** ask again and again; compare reports with earlier ones
5. **Commitment and delegation** P1 chooses y directly? P2 commit to y as function of m?

Weeks 7:
Auction

Auction Formats & Equilibrium:
- Coins Games: Winner's Curse
- Wallet Games

Auction Design: Case Study
Efficient Mechanism Desi

Common Auction Formats

1. First-price sealed-bid Auction
   - Bidders simultaneously write down their 'best-and-final' bids
   - Higher bidder wins the object and pays the bid he wrote
   - e.g., oil, mineral rights, real estate, construction, procurement contracts
2. Dutch Auction
   - Price starts high and gradually falls until one bidder agrees to buy the object at the price.
   - e.g., flowers in Holland, some fish and agricultural products.

Above are the same game.

3. Japanese Auction (Ascending Auction)
   - price starts low and gradually increases until only one bidder left.
   - Highest bidder then wins and pays own bid

Coins Game: the Winner's Curse

The coins game is 'Common Values':

the actual value of the prize is the same for all bidders, but different bidders have different estimates of this value.

Suppose bidders act as if the value of the coins is their estimate of the value before any information is revealed by other bidder's behaviour.

—— Then the winner is likely to have paid too much.

If you win, you probably had the highest estimate.

You should drastically shade your bid below your initial estimate in this auction.

Wallet Game (phone number)

(Pure Common Values: Ascending Auction)

Symmetric Equilibrium.

Type z_1 quits at same price as opponent type $z_2 = z_1$

If type z_1 wins at her quitting price, $b(z_1)$, the type she will have beaten z_2 close to z_1.

$\Rightarrow$ Type z_1 of player 1 quits at price $b(z_1) = 2z_1$.

—— independent of distributions of z_1 and z_2.

$$v_1 = v_2 = v = z_1 + z_2$$

Type z_1 of player 1 quits at price $z_1 + z_1 = 2z_1$
- e.g. $z_1 = 900$ bids up to, and then quits at, price 1800p (=£18)
- e.g. $z_1 = 030$ bids up to, and then quits at, price 60p

Question: At price 60p, z_2 is at least 030
- So why doesn't $z_1 = 030$ bid further?
- v_1 is at least 60p ⇒ On average $v_1 > 60$p

Winner's Curse!
- If $z_1 = 030$ wins at 61p, then $v_1 \leq 30 + \frac{61}{2} = 60\frac{1}{2}$
- So z_1 should have quit earlier!

(Pure Common Values: Sealed-bid Auction)

Those without information: should bid zero

Those with information: depends on the distributions

e.g. if z_1 & z_2 are independently and uniformly distributed on $[0, K]$, for any $K>0$, then, in equilibrium, bidder i bids z_i.

(Asymmetric Case: Ascending Auction)

$\begin{cases} \text{Prize is } V_1 = z_1 + z_2 + £1 \text{ if Bidder 1 wins} \\ \text{Prize is } V_2 = z_1 + z_2 \text{ if Bidder 2 wins} \end{cases}$

⇒ Bidder 1 bid £2 more than before

If she wins at price p, she infers $z_2 = p/2$, $V_1 = z_1 + \frac{p}{2} + 1$

So she will bid until $p = V_1 \Rightarrow p = 2z_1 + 2$

⇒ Bidder 2 quit £2 earlier than before

Now if bidder 2 wins at p, she infers $z_1 = (p-2)/2$

So her value $V_2 = (p-2)/2 + z_2 = p \Rightarrow p = 2z_2 - 2$

⇒ Bidder 1 bid £4 more than before

⇒ Bidder 2 quit £4 earlier than before

... ⇒ Bidder 1 never quits. Bidder 2 quits bidding at price $\leq z_2$.

Polar Cases of Bidder Value

1. Common Values

– The eventual value of the object will be the same for all bidders, but different bidders may have different estimates of this value

– Each bidder's estimate of the value would be altered by knowledge of other bidders' estimates

2. Private Values

– Each bidder has some value for the object (i.e., maximum price she is willing to pay) which does not depend on the information or values of other bidders

Auction Design : Case Study

Key problems in Auction Design :

1. Collusion / Coordination (including auctioneer)
2. Entry Deterrence
3. Subject to 1 & 2, reveal information

—— An auction is just a market and standard economics applies

Revenue Equivalence Theorem.

All auctions with no reserve price yield the same expected revenue for the seller.

Formal statement (Revenue Equivalence Theorem)

"Assume each of n risk-neutral potential buyers has a privately-known value (or signal in the common-values case) independently drawn from a common distribution F(z) that is strictly increasing and atomless on $[\underline{z}, \bar{z}]$

Then any auction mechanism in which
(i) the object always goes to the buyer with the highest value (or signal), and
(ii) any bidder with value (or signal) z expects zero surplus
, yields the same expected revenue, and results in a buyer with value (or signal) z making the same expected payment."

1. All auctions with a suitable (public) reserve price yield the maximum possible expected revenue for the seller
   Moreover, absent collusion problems, using a reserve price typically only makes much difference if you have very few bidders
2. These results don't usually apply if more than one object is sold.
3. Other caveats :
   - asymmetric wallet game

UK 3G mobile phone auction

5 mobile-phone licenses on sale
Bidders each allowed to win at most 1 license
Ran (simultaneous) ascending auction (prices rise until only 5 left)
13 bidders entered, raised £22.5 billion = 2.5% GNP

Crucial details in UK and Netherland cases

- UK auctioned 5 licences and UK had 4 strong bidders (incumbent operators)
- Netherlands auctioned 5 licences but had 5 strong bidders (incumbent operators)

—— weaker companies have no incentive to bid in the Netherlands case.

- Denmark: sealed-bid auction successfully (4 licenses, 4 incumbent bidders)
- Austria: 6 bidders for 12 lots, ascending auction
  { Bidders permitted to win >1 lot each
  { Asking price rises until bidders ask for 12 units in total ⇒ Collusion, 2 lots / company

Why Ascending Auction rather than Sealed-bid Auction in UK?

Simultaneous ascending auction likely to be efficient if can attract entrants & prevent demand reduction

In UK 3G auction:

1. Bidders allowed to win 1 license only ⇒ no scope for dividing spoils
   ⇒ Demand reduction / collusion not a major worry
2. 5 licenses, 4 incumbent bidders. so at least one license goes to entrants
   ⇒ Entry not a major worry

Abilities firms require to collude:

1. agree division of market
2. detect defection from agreement
3. credibly punish defection
4. deter new entry

Even badly-run auctions are usually better than the alternatives

Auctions	Administrative Allocation ("Beauty Contests")
• Efficient: winners are bidders with highest values	• Often inefficient
• Transparent • Speedy • Fair	• Hard to specify criteria • Time-consuming • Outcome often contested
• Seller gets most of value (without deadweight losses) 3G auctions: €100 billion (8 EU countries, pop. 250m)	• Seller gets little or nothing 3G beauty contests: €2bn (7 EU countries, pop. 130m)

Northern Rock Bank Run (2007.9)

Bank of England wanted to sell multiple types of loans to commercial banks, building societies (Type = quality of collateral used by borrower)

① 6-month loans against 'poor' collateral, e.g., MBS
② 6-month loans against 'good' collateral, e.g., UK government bonds
Total allocation = £2,500 million

1. Not running a separate auction for each variety
   Market power: too little competition if separate auctions

2. Not running Simultaneous Multiple Round Auction (SMRA)
(as pioneered by Paul Milgrom and Bob Wilson)

But: { may take too long
may aid collusion / predation
hard to allow the mix of varieties sold to depend upon the bids

Product-Mix Auction

① Bidders for spectrum licenses are oligopoly
⇒ entry & collusion / coordination are 1st order issues
—— use game theory models / insights

② Many potential bidders for loans in UK
⇒ entry & collusion / coordination not 1st order issues
⇒ focus on making bidding easy and efficient to extract information and implement competitive outcome
—— use competitive models / insights

Product-Mix Auction: { 1. each participant simultaneously states preferences
—— sealed bid auctions for multiple units of multiple differentiated goods
2. implement competitive equilibrium allocation

1. Auctioneer expresses preferences as a supply function
   - e.g. Bank of England can "bid" this:

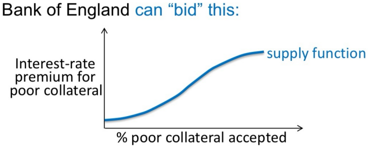

2. Bidders express preferences between goods
   - e.g. potential borrower can bid "(60bp,50bp; £200m)" to mean

"I would like to borrow £200m, and would pay up to 50 basis points using my good collateral, but up to 60 bp if I can use my poor collateral (and at lower interest rates I prefer to use poor collateral if rate difference < 10bp)"

3. Market power: low when all goods in same auction

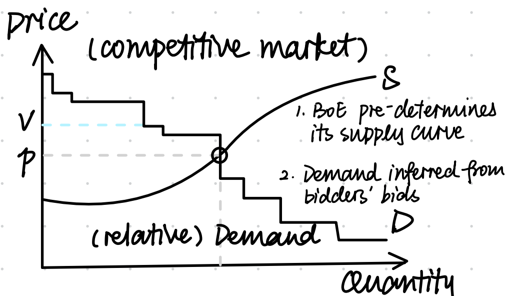

A consumer who reveals she has value v pays price p.

Theorem: under some conditions, the Product Mix-Auction achieves an 'efficient' allocation. All bidders, and the auctioneer, get exactly what they would have chosen at the final prices.

—— Bidding is efficient, informative and easy

(*) Another case mentioned in lecture: Ecosystem for Turtle Doves

Efficient Mechanism Design

1. VCG Mechanism: Efficient mechanisms with transfers
2. Gale-Shapley Algorithm: Stable matching without transfers

EMD { under private values (u_i only depends on θ_i but not θ_j)
with transferable utilities

Elements x of set X are various social decisions.

Agent i's utility measured in money: $u_i(x, \theta_i) + p_i$ for $i \in N$;
θ_i is its type (private information) and p_i her transfer received

Is there a way to carry out the socially efficient decision rule, $x^*(\theta) = \arg\max_{x \in X} \sum_{i \in N} u_i(x, \theta_i)$ for every state $\theta = (\theta_i)_{i \in N}$, without directly observing θ.

Applications: auctions ($x \in X$ describes who gets what)
public good provision (binary x shows if the bridge is built)

VCG Mechanism

Vickrey (1961: auction), Clarke (1971: PG), Groves (1973: general)

Vickrey-Clarke-Groves (VCG) mechanism:
(1) Each $i \in N$ reports a type, $\hat{\theta}_i$. Denote $\hat{\theta} = (\hat{\theta}_i)_{i \in N}$.
(2) Carry out decision $x^*(\hat{\theta})$, and set i's transfer equal to

$$p_i^*(\hat{\theta}) := \sum_{j \neq i} u_j(x^*(\hat{\theta}), \hat{\theta}_j) - \pi_i(\hat{\theta}_{-i}).$$

Here $\pi_i(\hat{\theta}_{-i})$ is any function with argument $\hat{\theta}_{-i} = (\hat{\theta}_j)_{j \neq i}$.

Each i receives the sum of the other agents' gross (pre-transfer) payoffs in the efficient allocation computed at the reported types, minus a transfer that may depend on the others' reported types.

VCG with $\pi_i(\hat{\theta}_{-i}) \equiv 0$ is called the team mechanism

Theorem: In any VCG mechanism it is **weakly dominant** to report $\hat{\theta}_i = \theta_i$, hence the outcome is **efficient** (ignoring transfers).

proof: Rewrite efficiency of x^* as follows.

For all θ_i, $\hat\theta_i$ and x,

$$u_i(x^*(\theta_i,\hat\theta_{-i}),\theta_i) + \sum_{j\neq i} u_j(x^*(\theta_i,\hat\theta_{-i}),\hat\theta_j) \geq u_i(x,\theta_i) + \sum_{j\neq i} u_j(x,\hat\theta_j)$$

In VCG, if i's own type is θ_i, and her and the others' reports are $\hat\theta_i$ and $\hat\theta_{-i}$, then i's payoff: $u_i(x^*(\hat\theta_i,\hat\theta_{-i}),\theta_i) + \sum_{j\neq i} u_j(x^*(\hat\theta_i,\hat\theta_{-i}),\hat\theta_j) - T_i(\hat\theta_{-i})$

—— which is highest at $\hat\theta_i = \theta_i$ by the efficiency of x^* as written above

Vickrey's Pivot Mechanism (Nobel prize '96).

Pivot Mechanism: VCG with $T_i(\hat\theta_{-i}) = \max_{x\in X} \sum_{j\neq i} u_j(x,\hat\theta_j)$

Vickrey payment: in allocation $x^*(\hat\theta)$, player i plays:

$$-p_i^{**}(\hat\theta) := \max_{x\in X} \sum_{j\neq i} u_j(x,\hat\theta_j) - \sum_{j\neq i} u_j(x^*(\hat\theta),\hat\theta_j)$$

(the effect of her presence on the welfare of others)

Hence, i's net payoff is her marginal contribution:

$$m_i(\hat\theta) = \sum_{j\in N} u_j(x^*(\hat\theta),\hat\theta_j) - \max_{x\in X} \sum_{j\neq i} u_j(x,\hat\theta_j)$$

Single unit for sale: Second-price auction (winner pays the second-highest bid), it's optimal to bid truthfully.

Vickrey's auction for K identical goods

Suppose I agents are interested in buying K identical goods.
Each agent has decreasing, privately-known marginal valuations.

Notation: ~~v1i > . . . > vKi~~ are i 's valuations for kth unit. $v_1^i > \cdots > v_K^i$

Ex-ante distribution of ~~vki~~ v_k^i won't matter, hence ignored.

1. Each agent $i \in I$ submits K bids, say $\hat v_1^i > \ldots > \hat v_K^i$.
2. Highest K bids (pooled, from all bidders) win.
3. Each i pays the sum of the bids **submitted by others** that would have won had i not participated in the auction.

Proposition: it is weakly optimal for i to submit $\hat v_k^i = v_k^i$, $\forall k \in K$.

Observation:
1. the price a winner pays is not the 'highest rejected' or 'lowest winning' bid; each i may pay different prices.
2. i's bids don't affect how much she pays, only whether she wins
3. Principle: pay your externality (the bids you crowded out)

Example: 3 people for 3 apples.

A: 8p, 2p, 0   A pays 5p for her apple.
B: 7p, 6p, 5p   B pays 4p + 2p for her two apples.
C: 4p, 0, -6p

Ausubel's auction (a real-time implementation)

open-bid, ascending-price counterpart of Vickrey's K-unit auction.
⇒ trust

Price clock p is set to 0. Each i indicates her demand by showing at most K fingers. As p rises agents may reduce their demand.

If, at price p, the total number of fingers held up by agents other than i falls one short of the number of still-available units, then we say agent i has clinched a unit at price p.

The unit clinched by i as well as i's demand for it are removed.

The price clock resumes unless all units have been allocated.

Some known issues and recent solutions to VCG

1. VCG requires private valuations (u_i not to depend on θ_j)
   Dasgupta & Maskin (QJE, 2000) extended the computation of Vickrey payments to interdependent valuations.
2. VCG can be manipulated by players colluding (merging) or a single player pretending to be two or more players
   Collusion-proof mechanisms: Che & Kim (Econometrica, 2006)
3. VCG itself does not accommodate dynamic problems with players arriving & leaving and surplus generated over time. But see:
   Athey & Segal (Econometrica, 2013: dynamic team mechanism)
   Bergemann & Välimäki (Econometrica, 2010: dynamic pivot)

Market Design for Matching

Matching and market design

Typical problems:
- match young children to primary schools
- medical residents to hospitals
- donated kidneys to patients.

Common characteristics:
- Participants on both sides of the market may have preferences over potential matches. They may lie to gain advantage. Monetary transfers are not allowed.

Motivating example:

Allocate K economics tutors (new APs) to K Oxford colleges.
Each tutor is able to rank the colleges in a strict, transitive order
Each college has a strict ranking of the tutors
(Monetary transfers not allowed)

Task: stable pairings establishment
(each tutor to be matched to a college in such a way that no tutor would prefer a different college and that would also prefer him over the tutor allocated to them)

Alice:	N > K > J
Bob:	N > K > J
Carol:	K > N > J

N:	Carol > Alice > Bob
K:	Bob > Carol > Alice
J:	Bob > Carol > Alice

Is there an algorithm in which { tutors / colleges state preferences truthfully
such that leads to a stable matching

Gale-Shapley deferred acceptance algorithm:

Start: Each tutor names his or her favourite college.

Loop: Colleges demanded by multiple tutors **provisionally** pick a tutor. Tutors not picked by a college choose again, with the restriction that they cannot choose a college that rejected them before.

The loop is repeated until each tutor is allocated to a college.

Round 1: A $\mapsto$ New, B $\mapsto$ New, C $\mapsto$ Keble. New wants A, so B is not picked.
Round 2: B $\mapsto$ Keble. Keble picks B, so now C is without a college.
Round 3: C $\mapsto$ New. New picks C, so now A is left without a college.
Round 4: A $\mapsto$ Keble. Keble picks B, so A has to choose again.
Round 5: A $\mapsto$ Jesus, B $\mapsto$ Keble, C $\mapsto$ New. Stable!

properties:
1. If every participant behaves truthfully, then the Gale-Shapley algorithm leads to the stable matching that is optimal for the side that makes the initial demands (here: tutor-optimal)
2. It is dominant strategy for tutors to tell the truth.
3. If there are multiple stable matchings, then the receiving side (colleges) can manipulate to obtain their favorite pairing
4. If there are multiple stable matchings, there's no stable matching algorithm that's impossible to manipulate.

Top Trading Cycle (TTC) algorithm (David Gale)

1. Each tutor points their 'top' most-preferred college (could be their own), and each college points to their default tutor.
2. Identify cycles: carry out the suggested trades and remove the participants. Repeat from step 1 till done.

Theorem: truth-telling is dominant, the outcome is stable

Weeks 8 :
Repeated Games

Folk Theorem & Perfect Folk Theorem
Renegotiation
Finit Repetition

Repeated Games in Practice:

1. social norms, customs, threats, punishments, renegotiation and revenge
2. model tacit collusion (默契勾结) among firms in an industry

Framework: A stage game played a finit / infinit number of times.
In period t, each $i = 1, \cdots n$ simultaneously picks stage-game strategy $a_i \in A_i$ for stage payoff $u_i(a) \in \mathbb{R}$, where $a = (a_i, a_{-i})$

Notation:

1. Denote player i's stage-game mixed strategy by α_i and i's expected payoff when playing a_i against the others' mixing by $u_i(a_i, \alpha_{-i})$
2. Denote the state-game action profile played at t by $a^t = (a_1^t, \cdots a_n^t)$
The history of play at time t is $h^t = (a^1, \cdots a^{t-1})$

payoff $= 1 + 1 + 1 + \cdots$ (discount: δ)
$\Rightarrow \frac{1}{1-\delta}$ $\boxed{\times (1-\delta)}$

Definition:

1. A repeated game strategy for player i specifies an action (pure or mixed) in each period t as a function of the history at t.
2. Payoff: δ-discounted present value of stage-game payoffs, where $\delta \in (0,1)$ is the probability of repetition times the discount factor.
3. Average Discounted Value (ADV):
   $ADV_t = (1-\delta) \cdot PV_t$ (definition)
   - ADV recursive: $ADV_t = (1-\delta) \cdot$ today's payoff $+ \delta \cdot ADV_{t+1}$
   - For infinitive repetition: $ADV = PV \cdot (1-\delta)$
   - ADV of "v for k periods, then v' forever": $(1-\delta^k)v + \delta^k v'$
   interpret δ as weight here

Folk Theorem

Payoff constraint 1: Feasibility

Feasibility. We say $(v_1, \ldots, v_n)$ is a feasible average discounted payoff of the repeated stage game iff

$$(v_1, \ldots, v_n) \in \text{co}\{(w_1, \ldots, w_n) : w_i = u_i(a) \text{ for some } a \in A\}.$$

"co" means "convex hull", i.e., smallest convex containing set.

1. Players generate such $(v_1, \cdots v_n)$ via Public randomization device
2. 'Coordinated Cycling' over pure outcomes of the stage game.
For any δ, if the number of repetitions is large enough, $(v_1, \cdots v_n)$ can be approximated arbitrarily closed in ADV.

Payoff constraint 2: Individual Rationality

Player i's minmax payoff is $\underline{v}_i = \min_{\alpha_{-i}} \max_{\alpha_i} u_i(\alpha_i, \alpha_{-i})$
(if all players other than i coordinate on punishing i, but i knows this and gives his best response, then i gets $\underline{v}_i$.)

Individual Rationality. If $(v_1, \ldots, v_n)$ are average discounted Nash equilibrium payoffs of a repeated game then for all i, $v_i \geq \underline{v}_i$.

Example: a contribution game

If Pr(P2 *gives*) $\geq 1/2$ then P1's best reply is *give* $\Rightarrow u_1 \geq 1$.

If Pr(P2 *gives*) $\leq 1/2$ then P1's best reply is *not* $\Rightarrow u_1 = 1$.

Therefore $\underline{v}_1 = 1$.

By symmetry, $\underline{v}_2 = 1$.

	give	*not*
give	<u>2</u>,<u>2</u>	0,1
not	1,0	<u>1</u>,<u>1</u>

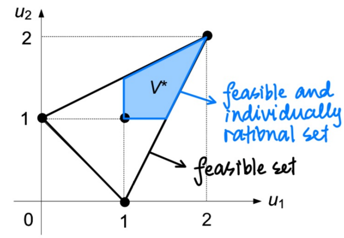

Three comments on minimax strategies and payoffs:

1. When i minmaxes j, player i's payoff may be less than $\underline{v}_i$ (punishing others can hurt more than being punished!)
2. Minimax(ing) strategies may be mixed. Example:

(*) player i plays strategy that minimizes player j's rational payoff

Matching pennies
$\underline{v}_1 = \underline{v}_2 = 0$.
Minmax by $(0.5 \cdot H + 0.5 \cdot T)$

	H	T
H	1,-1	-1,1
T	-1,1	1,-1

proof: (coincide lower and upper bounds)

$\begin{cases} \text{P1 can hold P2 to 0 payoff by P1 mixing 50-50\%, so } \underline{v}_2 \leq 0 \\ \text{P2 can guarantee 0 payoff by P2 mixing 50-50\%, so } \underline{v}_2 \geq 0 \end{cases}$

$\Rightarrow$ Hence $\underline{v}_2 = 0$

Folk Theorem for Nash Equilibria

Folk Theorem for Nash equilibria with ∞ repetition, discounting. If $v \in V^*$ then there is $\underline{\delta} < 1$ such that for all $\delta \geq \underline{\delta}$, the ∞-repeated game has a Nash equilibrium in which each player i's ADV equals v_i.

But not credible

Proof. For simplicity, suppose v is generated by stage-game action profile a, that is, $v_i = u_i(a)$ for all i.

<u>Proposal: Each i plays a_i if a has been played in all previous periods. Otherwise all play to minmax the player that deviated first.</u>

If i deviates she expects at most $\bar{v}_i$ (her maximal payoff) once, then $\underline{v}_i$ (her minmax payoff) forever. She won't deviate if

$$v_i \geq (1-\delta)\bar{v}_i + \delta \underline{v}_i \iff \delta \geq (\bar{v}_i - v_i)/(\bar{v}_i - \underline{v}_i).$$

Hence $\underline{\delta} = \max_i\{(\bar{v}_i - v_i)/(\bar{v}_i - \underline{v}_i)\}$. □ So that nobody wants to deviate

Perfect Folk Theorem

Subgame Perfection

NE is not the right concept in repeated games because it allows players to use threats that are not credible.

Example:

Example:

	L	R
U	2,2	1,3
D	-1,1	-1,1

$\underline{v}_1 = \underline{v}_2 = 1.$

[Note: P1 minmaxes P2 by playing D.]

Folk Theorem suggests how to sustain (U, L) in a NE of the infinitely repeated game if δ is close to 1: If P2 deviates, P1 "punishes" her by playing D forever —— not credible!

More reasonable: Subgame Perfect Equilibrium (SPE)
—— play NE in every subgame, even in the continuation after a deviation

{ Subgame: continuation after t
{ stage game: at t

But: playing stage-game Nash at every t is of course SPE.

Theorem (J. Friedman, 1971):
Suppose v is a feasible payoff vector such that for all i, v_i exceeds player i's payoff in some Nash equilibrium of the stage game.
For δ sufficiently close to 1 there is a subgame perfect equilibrium of the infinitely repeated game with average discounted payoffs v.

1. 'Grim Trigger' in the Prisoner's Dilemma

Prisoner's dilemma:

	C	D
C	1,1	-1,2
D	2,-1	0,0

$\underline{v}_1 = \underline{v}_2 = 0.$

unique NE of the stage game: (D, D)
Hence: (D, D) in every period is SPE for any repetition, any δ.

Infinite repetition: Any feasible $(v_1, v_2) \geq (0, 0)$ is in SPE for high δ.
– In particular, $(v_1, v_2) = 1$ can be sustained in SPE for $\delta \geq \frac{1}{2}$:
{ play C at $t=1$ and as long as both players have played C.
{ if someone plays D, play D from then on forever.
(doesn't deviate iff $1 \geq (1-\delta) \cdot 2$, i.e. $\delta \geq \frac{1}{2}$)

Clever Tool / notation: Automata

Automation:
(representation)

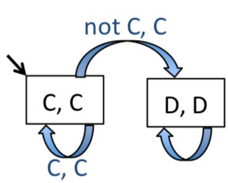

boxes: states
arrows: transitions

2. One-shot Deviation Principle

A proposed strategy profile is SPE if and only if, no player has an incentive to deviate at any state while obeying the transitions.

> **Theorem (One-shot deviation principle):**
> In a repeated game with discounting, a strategy profile is SPE iff the following holds for all histories, all t, all i:
>
> Provided all players other than i play their proposed strategies at and after t, player i cannot gain by deviating in period t and then reverting back to his proposed equilibrium strategy from $t+1$ on.

Proof: 1. Bottom line is that there's no need to check complex deviations
2. However, we must check unilateral, one-shot deviations in every continuation or subgame — every state of automation

- Suppose towards contradiction that player i's repeated-game strategy si cannot be improved in one step (at any t, after any history ht), yet i has a strictly better strategy si′ at a particular t and history ht.
- Due to discounting, payoffs far in the future make little difference. Any gain is less than Δ in present value from payoffs after t + TΔ.
- Therefore if si′ increases i's payoff (from that generated by si) by Δ, then it should do so in finitely many steps (within TΔ periods).
- Let's suppose si′ which improves on si by Δ in present value differs from si for some histories at t, . . . , t + TΔ, but it specifies the same action as si does for all histories from period t + TΔ + 1 onwards.

- Consider strategy s′′, which agrees with si′ for all histories in periods up to and including t + TΔ – 1, but agrees with si for all histories from t + TΔ onwards. (So s′′ reverts to si one period earlier than s′i.)
- In period t+TΔ, after any history, si′′ = si cannot be improved by a one-period deviation, hence s′′ is weakly better than s′ from then on. Since s′′ agrees with s′ before t + TΔ, it is weakly better than s′ at t.
- Therefore s′′, which deviates from si at t for only TΔ – 1 periods, is strictly better than si in period t at history ht .
- Repeat the argument TΔ times, each time shortening the duration of the deviation from si that improves it at t, history ht. Eventually we get a one-period improvement, a contradiction proving the claim.

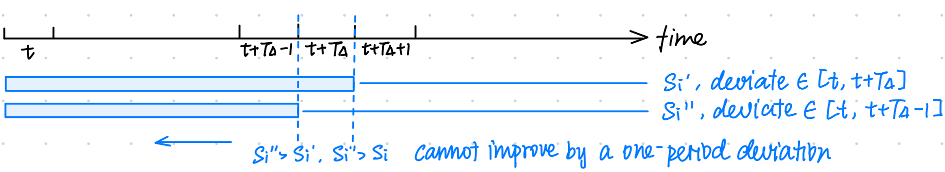

Checking for SPE:

	L	R
U	6, 6	3, 7
D	7, 3	-1,-1

Two pure, one mixed Nash in this stage game

How to sustain (U,L) forever in SPE of the infinitely repeated game with δ close to 1?

First, consider the proposal:

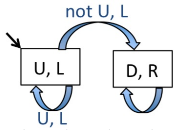

$\Rightarrow$ not SPE (in (D,R), each would deviate)
on-path payoff: $-1-\delta-\delta^2-\cdots$
one-shot deviation: $3-\delta-\delta^2-\cdots$ (more!)

SPE for δ close to 1:
($\delta \geq \frac{4}{7}$)

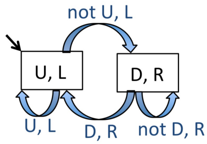

$\Rightarrow$ SPE in (U,L)

At (U,L): $\begin{cases} \text{on-path: } 6+6\delta+6\delta^2+\cdots \\ \text{off-path: } 7-\delta+6\delta^2+\cdots \end{cases} \Rightarrow \delta \geq \frac{1}{7}$

At (D,R): $\begin{cases} \text{on-path: } -1+6\delta+6\delta^2+\cdots \\ \text{off-path: } 3-\delta+6\delta^2+\cdots \end{cases} \Rightarrow \delta \geq \frac{4}{7}$

stick-carrot (opposed to grim trigger):
punish hard for a limited time, then return to cooperation

Perfect Folk Theorem

Fudenberg & Maskin:

> **Theorem (D. Fudenberg and E. Maskin, 1986):**
> Assume V^* (set of feasible and strictly IR payoffs) is n-dimensional.
> For any $(v_1, \ldots, v_n) \in V^*$ there is $\underline{\delta} \in (0,1)$ such that for all $\delta \geq \underline{\delta}$ the infinitely repeated game has a subgame perfect equilibrium with average discounted payoffs $(v_1, \ldots, v_n)$.

1. Dimensionality condition: We need to be able to punish deviators individually, so the stage game payoffs need to vary 'independently' across players.

2. Notation before proof:
   - Let $v = (v_1, \cdots v_n)$ be the payoff that we want to support in SPE in the infinitely repeated game.
   - Let the strategy profile that minmaxes i be $m^i = (m_1^i, \cdots m_n^i)$
     Normalize payoffs so $\underline{v}_i \equiv u_i(m^i) = 0$; however $u_j(m^i) \gtrless 0$
   - Denote i's maximal payoff in the stage game by $\bar{v}_i$.
   - Pick w that is feasible and $0 < w < v$
     Let T be such that $Tw_i > \bar{v}_i$ for all i.
   - Let $\varepsilon > 0$ be so that $w^i \equiv (w_i, w_{-i}+\varepsilon) < v$ is feasible for all i.
     (that is, $w_i^i = w_i$, but $w_j^i = w_j + \varepsilon$ for all i and $j \neq i$)

Illustration of the Notation:

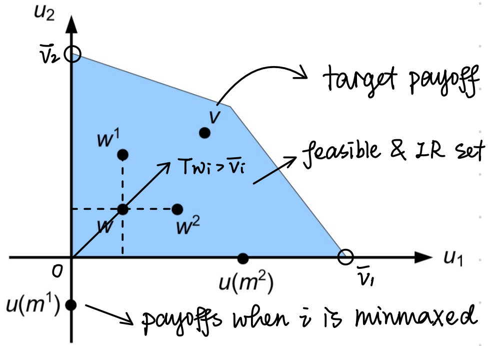

3. Construction of the Equilibrium

I. **Collaboration**: play stage game actions that generates v, repeat unless player j deviates, in that case go to phase II^j.

II^j. **Punishment**: play m^j (minmax j) for exactly T periods.

$\begin{cases} \text{if no one deviates, then go to phase } \text{III}^j \\ \text{if player } k \text{ deviates then start over phase } \text{II}^k \end{cases}$

III^j. **Reconcilation**: play stage game actions that generate w^j repeat unless player k deviates, in that case go to phase II^k.

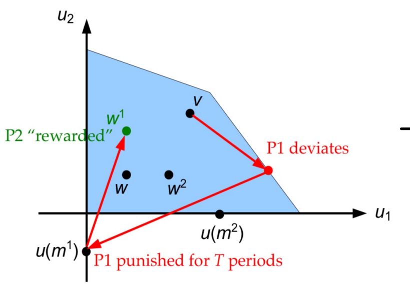

P1 deviates in Phase I

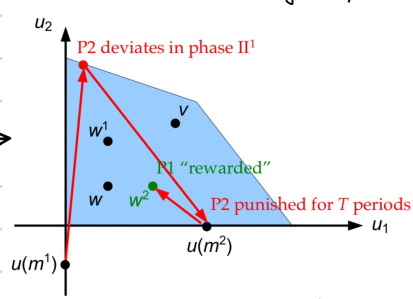

P2 deviates in Phase II'

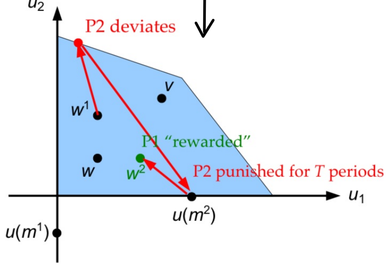

P2 deviates in Phase III'

4. Formal Proof by one-shot deviation principle

Phase I. Player i's payoff is $v_i \equiv (1-\delta^{T+1})v_i + \delta^{T+1}v_i$.
If i deviates then i gets at most $(1-\delta)\bar{v}_i + \delta^{T+1}w_i$.

Which one is bigger, for δ close to 1?

$$\begin{aligned}(1-\delta^{T+1})v_i &= (1-\delta)(1+\delta+\ldots+\delta^T)v_i \\ &> (1-\delta)Tv_i > (1-\delta)Tw_i > (1-\delta)\bar{v}_i,\end{aligned}$$

where the first inequality holds for δ near 1, the second inequality holds by $v > w$, and the final one by the definition of T.

Since $\delta^{T+1}v_i > \delta^{T+1}w_i$, we find that

$$v_i \equiv (1-\delta^{T+1})v_i + \delta^{T+1}v_i > (1-\delta)\bar{v}_i + \delta^{T+1}w_i,$$

so a deviation in Phase I is not profitable for i.

Phase II^i: While i is minmaxed, he cannot get more than 0.
Deviation only postpones the end of punishment, not worth it.

Phase II^j: At its start i expects $(1-\delta^T)u_i(m^j) + \delta^T(w_i+\varepsilon)$.
With $K \leq T$ periods left this is $(1-\delta^K)u_i(m^j) + \delta^K(w_i+\varepsilon)$.
By deviating i gets at most $(1-\delta)\bar{v}_i + \delta^{T+1}w_i$.

Player i does not deviate in Phase II^j iff

$$(1-\delta^K)u_i(m^j) - (1-\delta)\bar{v}_i + (\delta^K - \delta^{T+1})w_i + \delta^{T+1}\varepsilon \geq 0.$$

As $\delta \to 1$, the first three terms tend to zero, the last one to $\varepsilon > 0$.

Phase III: By conforming player i gets at least w_i.
By deviating player i gets at most $(1-\delta)\bar{v}_i + \delta^{T+1}w_i$.

As in the analysis of Phase I, for δ near 1,
$(1-\delta^{T+1})w_i = (1-\delta)(1+\delta+\ldots+\delta^T)w_i > (1-\delta)Tw_i > (1-\delta)\bar{v}_i$.

Hence $w_i \equiv (1-\delta^{T+1})w_i + \delta^{T+1}w_i > (1-\delta)\bar{v}_i + \delta^{T+1}w_i$, and so a deviation is not profitable in Phase III either. □

Take-away on ∞-repetition Folk Theorems

1. Any feasible and individually rational payoff can be sustained in SPE of the indefinitely-repeated game for δ sufficiently close to 1.
   —— Perfect Folk Theorem (Fudenberg & Maskin, 1986)
2. Three key observations:
   - credible punishment need not be 'Nash reversion'
   - Forgiveness: punishment need not last forever
   - Reconcilation: worse for the deviator, better for punisher(s)

Renegotiation

Criticism towards repeated game SPE with 'punishment phases':
Both players are worse off at the beginning of a punishment phase, so they may want to renegotiate it.

—— Collaborative repeated game that not only subgame perfect but also 'renegotiation proof'?

Farrell and Maskin (1989):

Weak Renegotiation Proofness: No continuation play in the 'book of plays' is Pareto dominated by any other.

book of play: A SPE consists of subgame-perfect continuation plays after every possible history. Call this the 'book of plays' according to the SPE.

proof: otherwise players would renegotiate to the mutually preferred continuation

(*) weak renegotiation proofness involves a comparison within the playbook of a given SPE, not across different SPEs.

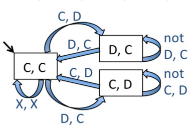

Check SPE (one-deviation principle):
@[C,C], need $1+\delta \geq 2-\delta, \Leftrightarrow \delta \geq 1/2$
@[D,C], need $-1+\delta \geq 0-\delta, \Leftrightarrow \delta \geq 1/2$
@[C,D], same as [D,C] by symmetry.

The automaton on the left, above, is "Tit-for-tat then forgive":

1. Play (C, C) unless someone deviates
2. If P1 defects then play (C, D) once (try again if not successful) and return to (C, C); if P2 defects then (D, C) once and return.

The computation on the right, above, shows this is SPE for $\delta \geq 1/2$.

Renegotiation proof: P1 prefers starting from [D,C] to [C,C] to [C,D]; P2 prefers continuation plays in the exact reverse order.

	L	R
U	6, 6	3, 7
D	7, 3	-1,-1

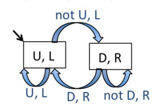

The stick-carrot strategies above, right, form a SPE for δ near 1 (specifically, for $\delta \geq 4/7$), but aren't renegotiation proof. (Why?)

The following construction is a renegotiation proof SPE:

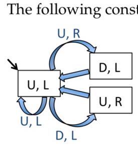

Check SPE (one-deviation principle):
@[U,L], need $6+6\delta \geq 7+3\delta$, i.e., $\delta \geq 1/3$.
@[D,L]: Nash, no gainful one-shot deviation.
@[U,R]: same.

P1 ranks [D,L] > [U,L] > [U,R], P2 in reverse.

Theorem (J. Farrell and E. Maskin, 1989):
Assume $n = 2$. Let $v = (v_1, v_2)$ be a feasible, strictly IR payoff-pair, and assume it is generated by action profile $a = (a_1, a_2)$.

If there exist stage game actions $a' = (a'_1, a'_2)$ and $a'' = (a''_1, a''_2)$ with

(1) $v_1 > c_1 \equiv \max_{\tilde{a}_1} u_1(\tilde{a}_1, a'_2)$ and $v_2 < u_2(a'_1, a'_2)$,

(2) $v_1 < u_1(a''_1, a''_2)$ and $v_2 > c_2 \equiv \max_{\tilde{a}_2} u_1(a''_1, \tilde{a}_2)$,

then for δ near 1 the infinitely repeated game has a renegotiation proof SPE with average discounted payoffs $v = (v_1, v_2)$.

Converse: If an infinitely repeated game has a renegotiation proof SPE with average discounted payoffs v then there are action profiles a' and a'' satisfying (1) and (2) with weak inequalities.

In PD, we support (C, C) using:
$$\begin{cases} a' = (C, D), \\ a'' = (D, C). \end{cases}$$

Moral:
Forgive but not Forget.
(punishment must reward the punisher before cooperation can be resumed)

Finite Repetition Folk Theorems

Theorem (J-P. Benoît and V. Krishna, 1985):
Assume $n = 2, \delta = 1$. Let $v' = (v'_1, v'_2)$ and $v'' = (v''_1, v''_2)$ be Nash equilibrium payoffs in the stage game with $v'_1 < v''_1$ and $v'_2 > v''_2$.

If $v = (v_1, v_2)$ is a feasible payoff vector such that $v_1 \geq v'_1$, $v_2 \geq v''_2$ and $v \geq \lambda v' + (1 - \lambda)v''$ for all $\lambda \in [0, 1]$, then, for T sufficiently large, it can be approximated arbitrarily closely by the average payoff of a subgame perfect equilibrium of the T-repeated game.

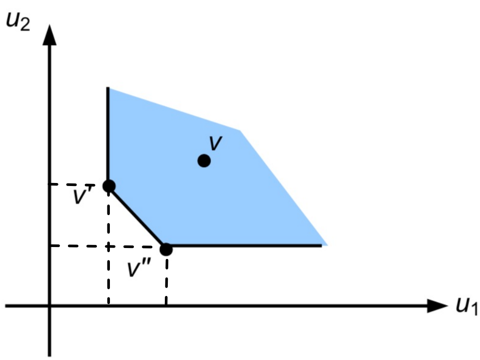

V', V'': Nash Equilibrium payoffs
V: feasible payoff
—— v approximated as the average payoff of a SPE provided T sufficiently large

Benoit & Krishna: Finite repetition expands the set of SPE if the stage game has multiple Nash equilibria ranked differently by the players.

proof of Benoît & Krishna's Theorem:

Set T^* s.t. $\begin{cases} T^* \dfrac{V_1'' - V_1'}{2} \geqslant \overline{V}_1 - V_1 \quad ① \\ T^* \dfrac{V_2' - V_2''}{2} \geqslant \overline{V}_2 - V_2 \end{cases}$

P1 conforms: $V_1 + T^* \frac{V_1' + V_2''}{2}$
P1 deviates: $\overline{V}_1 + T^* V_1'$
$\Rightarrow$ conforms > deviates: inequation ①

Assume that $T \gg T^*$. Proposed equlibrium strategies:

1. For $t \leq T - T^*$, play to generate (V_1, V_2) unless someone deviates
2. If no one has deviated at any $t \leq T - T^*$, then from then on alternate between playing (V_1', V_2') and (V_1'', V_2'').

$\begin{cases} \text{if Player 1 deviated at } t \leq T - T^*, \text{ then play } (V_1', V_2') \text{ to the end} \\ \text{if Player 2 deviated at } t \leq T - T^*, \text{ then play } (V_1'', V_2'') \text{ to the end} \end{cases}$

Remarks: PD has no NE, thus unique SPE = (D, D) on finite repetition

Construction in a familiar example:

	L	R
U	6, 6	<u>3</u>, <u>7</u>
D	<u>7</u>, <u>3</u>	-1,-1

Pure Nash: (U, R) and (D, L);
mixed Nash: $\left(\frac{4}{5}U + \frac{1}{5}D, \frac{4}{5}L + \frac{1}{5}R\right)$.

Nash payoffs: (3,7); (7,3); (5.4,5.4).

Suppose the stage game is played $T < \infty$ times, no discounting.

SPE to sustain (U, L) in all periods but the final one:

- Play (U, L) until the last period unless someone deviates; if no-one has deviated, play mixed Nash in period T.
- If P1 deviates then play (U, R) in all subsequent periods.
- If P2 deviates then play (D, L) in all subsequent periods.

Recent Developments

1. Imperfect Public Monitoring

past actions are unobservable, there are public signals imperfectly correlated with aggregate action

Example: the OPEC cartel

- OPEC's goal is to constrain total oil production and keep oil prices high ("stable").
- Explicit collusion: Each member country has a quota. Members' production is difficult to monitor, but world-wide oil prices are publicly observable.
- A drop in oil price can be the result of a member country exceeding its quota <u>or</u> a fall in oil demand.

Cooperative / collusive equilibria may still be constructed under imperfect public monitoring using trigger strategies.

—— price wars potentially triggered by demand shocks

2. Imperfect Private Monitoring

suppose at each t, each player observes a private signal that agrees with the opponent's action with probability $1-\varepsilon$ and differs from it with probability ε. (Can't observe own payoff till the end.)

What goes wrong in PD, with (say) grim trigger punishment?

- Suppose Player 1 observes that Player 2 played D. But in equilibrium Player 1 knows that Player 2 did not deviate.
- Player 1 knows that if he carries out the required punishment (playing D at least once) then Player 2 likely observes D.
- But that would trigger punishment from Player 2. Since Player 1 prefers cooperation, he won't play D as required.
- As a result, Player 2 can deviate without risking punishment. The equilibrium breaks down!

J. Ely, J. Välimäki (2002) and M. Piccione (2002):

sustain "(C,C) forever in PD" via belief-free equilibrium

- Construct an equilibrium where i is indifferent between choosing C and D if j is playing C in the same round.
- Also make i indifferent between C and D in case j is playing D. So i is indifferent between C and D at all t, after any history.
- If Player j defects, i can punish her by decreasing the probability of playing C .

In such an equilibrium, i does not care whether he observed D by mistake or he thinks j has really deviated.

—— Folk Theorem for all games: Hörner and Olszewski (2006).

Insights for competition policy

... makes collusion easier:

1. multimarket contact (cheat on one market, retaliate on all)
2. 'meet the competition' clause: customers help detect defection.
3. 'most favored customer' clause (= price commitment)
4. trade body that monitors and reports on the firm's actions

... makes collusion harder:

1. leniency towards whistleblowers
2. business cycles
3. imperfect information, differentiation of products