Game Theory

Li Dazhu

Dept. of Philosophy
PKU
1301211434
lidazhu91@163.com

Apr. 28th, 2015 / Seminar

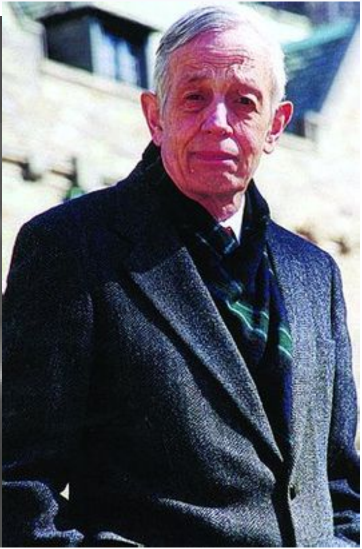

What is a Game

- A game is a formal representation of a situation in which a number of individuals interact in a setting of strategic interdependence.
- Each individual's welfare depends on the actions of herself and the other individuals.
- Four elements of a game
  - The *players*: who is involved?
  - The *rules*: who moves when? what do they know when they move? what can they do?
  - The *outcomes*: for each possible set of actions by the players,what is the outcome of the game?
  - The *payoffs*: what are the player's preferences (ie. utility functions) over the possible outcomes?

Matching Pennies:

- *Players*: There are two players, denoted A and B.
- *Rules*: Each player simultaneously puts a penny down, either heads up or tails up.
- *Outcomes*: Two pennies match (either both head up or both tail up) or not.
- *Payoffs*: If the two pennies match, player A pays 1 dollar to player B; otherwise, player B pays 1 dollar to player A.

Tick-Tack-Toe:

- *Players*: There are two players, X and O.
- *Rules*: The players are faced with a board that consists of nine squares arrayed with three rows of three squares each stacked on one another. The players take turns putting their marks (X or O) into an as-yet-unmarkedsquare. Player X moves first. Both players observe all choices previously made.
- *Outcomes*: The first player to have three of her marks in a row wins and receives 1 dollar from the other player. If no one succeeds in doing so after all nine boxes are marked, the game is a tie and no payments are made or received by either player.

The extensive form representation of a game

The extensive form captures who moves when, what actions each player can take, what players know when they move,what the outcomes ia as a function of the actions taken by the players, and the players' payoffs from each possible outcome.

The extensive form relies on the conceptual apparatus known as a game tree. The key elements of a game tree include decision nodes (initial decision node, decision nodes, and terminal nodes), branches and information sets.

Tick-Tack-Toe

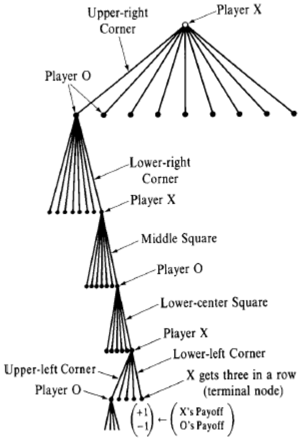

Matching Pennies Version C: Two players move sequentially, player B cannot see player A's choice until after player B has moved.

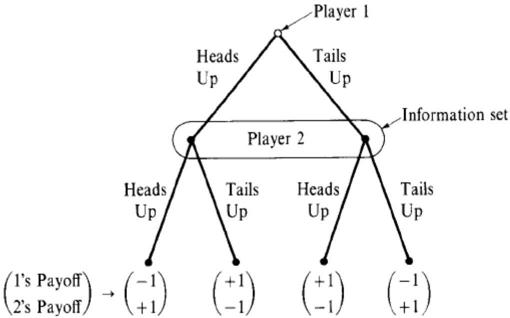

The meaning of the information set is that when it is player B's turn to move, she cannot tell which of these two nodes she is at because she has not observed player A's previous move.

Up to now, the outcome of a game has been a deterministic function of the players' choice. In many games, however, there is an element of chance. This, too, can be captured in the extensive form representation by including random moves of *nature*. We illustrate this point with still another variation, Matching Pennies Version D.

Matching Pennies Version D: Suppose that prior to playing Version B, the two players flip a coin to see who will move first.

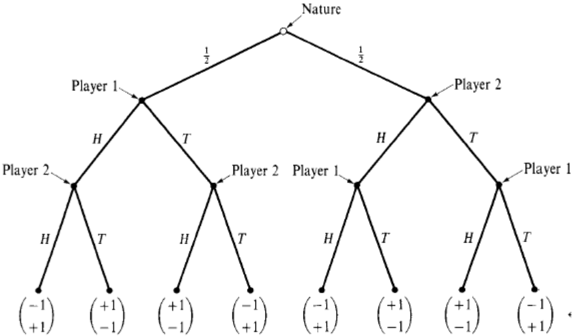

In addition to being depicted graphically, the extensive from can be described mathematically. Formally, a game represented in extensive form consists of the following items.

- A finite set of nodes N, a finite set of actions A, and a finite set of players $\{1, ...I\}$.
- A function p : $N \rightarrow N \cup \varnothing$specifying a single immediate predecessor of each node x; p(x) is nonempty for all $x \in N$ but one, designated as the initial node x_0. The immediate successor nodes of x are then $s(x) = p^{-1}(x)$, and the set of all predecessors and all successors of node x can be found by iterating p(x) and s(x). To have a tree structure, we require that these sets be disjoint (a predecessor of node x cannot also be a successor of it). The set of terminal nodes is $T = \{x \in N : s(x) = \varnothing\}$. All other nodes N$\backslash$T are known as *decisionnodes*.

- A function a : $N/x_0 \rightarrow A$ giving the action that leads to any noninitial node x from its immediate predecessor p(x) and satisfying the property that if $x', x'' \in s(x)$ and $x' \neq x''$ then $a(x') \neq a(x'')$. The set of choices available at decision node x is $c(x) = \{a \in A : a = a(x')$ for some $x' \in s(x)$.
- A collection of information sets $\breve{H}$, and a function $H : N \rightarrow \breve{H}$ assigning each decision node x to an information set $H(x) \in \breve{H}$.Thus,the information sets in $\breve{H}$ form a partition of N. We require that all decision nodes assigned to a single information set have the same choices available; formally $c(x) = c(x')$ if $H(x) = H(x')$. We can therefore write the choices available at information set H as $C(H) = \{a \in A : a \in c(x)$ for $x \in H\}$.

- A function $i : \overset{\frown}{H} \to \{0, 1, ..., I\}$ assigning each information set in $\overset{\frown}{H}$ to the player (or to nature:formally, player 0) who moves at the decision nodes in that set. We can denote the collection of player $i's$ information sets by $\overset{\frown}{H}_i = \{H \in \overset{\frown}{H} : i = i(H)\}$
- A function $\rho : \overset{\frown}{H} \times A \to [0, 1]$ assigning probabilities to actions at information sets where nature moves and satisfying $\rho(H, a) = 0$ if $a \notin c(H)$ and $\Sigma \rho(H, a) = 1$ for all $H \in \overset{\frown}{H}_0$.

- A collection of payoff functions $u = \{u_1(.), ..., u_I(.)\}$ assigning utilities to the players for each terminal node that can be reached, $u_i : T \to R$.
- Thus, formally, a game in extensive form is specified by the collection $\Gamma_E = \{N, A, I, p(.), a(.), \breve{H}, H(.), i(.), \rho(.), u\}$

Strategies and the normal form of a Game

A central concept of game theory is the notion of a player's *strategy*.

A *strategy* is a complete contingent plan,or decision rule,that specifies how the player will act in every possible distinguishable circumstance in which she might be called upon to move.

A *strategy* is a complete contingent plan that says what a player will do at each of her information sets if she is called on to play there.

DEFINITION(Strategy).

Let H_i denote the collection of player i's information sets, A is the set of possible actions in the game,and $C(H) \subset A$ the set of actions possible at information set H. A strategy for player i is a function $s_i : H_i \rightarrow A$ such that $s_i(H) \in C(H)$ for all $H \in H_i$.

DEFINITION(The normal form).

For a game with I players, the *normal form* representation Γ specifies for each player i a set of strategies S_i with $s_i \in S_i$ and a payoff function $u_i(s_i, ..., s_I)$ associated with the (possibly random) outcome arising from strategies $(s_1, ..., s_I)$. Formally, we write $\Gamma = [I, \{S_i\}, \{u_i\}]$.

The normal form of Matching Pennies Version B:

	s_1	s_2	s_3	s_4
H	−1, +1	−1, +1	+1, −1	+1, −1
T	+1, −1	−1, +1	+1, −1	−1, +1

Randomized Choices

Up to this point, we have assumed that players make their choices with certainty. However, there is no priori reason to exclude the possibility that a player could randomize when faced with a choice.

- Deterministic strategy: pure strategy.
- Randomized choice out of a set of pure strategies: mixed strategy.

DEFINITION(Mixed strategy).

Given player i's (finite) pure strategy set S_i , a mixed strategy for player i, $\sigma_i : S_i \to [0, 1]$, assigns to each pure strategy $s_i \in S_i$ a probability $\sigma_i(s_i) \geqslant 0$ that it will be played, where $\Sigma_{s_i \in S_i} \sigma_i(s_i) = 1$.

Player i's set of possible mixed strategies:

$$\Delta(S_i) = \{(\sigma_{1i}, ..., \sigma_{Mi}) \in R^M : \sigma_{mi} \geqslant 0 \text{ for all } m = 1, ..., M \text{ and } \Sigma_{m=1}^{M} \sigma_{mi} = 1\}$$

$$\Gamma_N = [I, | \Delta(S_i) |, | u_i(.) |]$$

Expected utility of a mixed strategy σ_i is given by $u_i(\sigma) = \Sigma_{s \in S}[\sigma_1(s_1)\sigma_2(s_2)...\sigma_I(s_I)]u_i(s)$, where $S = S_1 \times ... \times S_I$.

DEFINITION(Behavior Strategy)

Given an extensive form game Γ , a behavior strategy for play i specifies, for every information set $H \in \overset{\smile}{H}$ and action $a \in C(H)$, a probability $\lambda_i(a, H) \geqslant 0$ with $\Sigma_{a \in C(H)} \lambda_i(a, H) = 1$ for all $H \in \overset{\smile}{H}$.

Dominance

Prisoner's Dilemma

Let (N, S, u) be a game, $i \in N$. Strategy $s_i \in S_i$ *(weakly) dominates* $s_i' \in S_i$ if

(i) $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$ for all $s_{-i} \in S_{-i}$ and

(ii)$u_i(s_i, s_{-i}) > u_i(s_i', s_{-i})$ for at least one $s_{-i} \in S_{-i}$.

Strategy $s_i \in S_i$ *strictly dominates* $s_i' \in S_i$ if

$$u_i(s_i, s_{-i}) > u_i(s_i', s_{-i}) \text{ for all } s_{-i} \in S_{-i}.$$

If s_i (weakly/strictly) dominates all $s_i' \in S_i$, s_i is (weakly/strictly) dominant.

A strategy $s_i \in S_i$ is *strictly dominated* for player i in game $\Gamma = [I, \{S_i\}, \{u_i(.)\}]$ if there exists another strategy $s_i' \in S_i$ such that for all $s_{-i} \in S_{-i}$,

$$u_i(s_i', s_{-i}) > u_i(s_i, s_{-i})$$

In this case, we say that strategy s_i' *strictly dominates* strategy s_i.

A strategy $s_i \in S_i$ is *weakly dominated* for player i in game $\Gamma = [I, \{S_i\}, \{u_i(.)\}]$ if there exists another strategy $s_i' \in S_i$ such that for all $s_{-i} \in S_{-i}$,

$$u_i(s_i', s_{-i}) \geqslant u_i(s_i, s_{-i})$$

In this case, we say that strategy s_i' *weakly dominates* strategy s_i.

Unlike a strictly dominated strategy, a strategy that is only weakly dominated cannot be ruled out. For any alternative strategy that play i might pick, there is at least one profile of strategy for his rivals for which the weakly dominated strategy does as well.More generally, weakly dominated strategies could be dismissed if players always believed that there was at least some positive probability that any strategies of their rivals could be chosen.

Player 1 \ Player 2	L	R
U	1, −1	−1, 1
M	−1, 1	1, −1
D	−2, 5	−3, 2

Player 1 \ Player 2	L	R
U	5, 1	4, 0
M	6, 0	3, 1
D	6, 4	4, 4

Iterated Deletion of Strictly Dominated Strategies.

One feature of the process of iteratively eliminating strictly dominated strategies is that the order of deletion does not affect the set of strategies that remain in the end.This is fortunate, since we would worry if our prediction depended on the arbitrarily chosen order of deletion.

The DA's Brother.

Rationalizable Strategies

In general, player's common knowledge of each others' rationality and the game's structure allows us to eliminate more than just those strategies that are iteratively strictly dominated.

Here, we develop this point, leading to the concept of a *rationalizable strategy*.

The set of rationalizable strategies consists precisely of those strategies that may be played in a game where the structure of the game and the players' rationality are common knowledge among the players. Throughout this section, we focus on games of the form $\Gamma = [I, \{\Delta(S_i)\}], \{u_i(.)\}]$

DEFINITION(Best response).

In game $\Gamma = [I, \{\Delta(S_i)\}], \{u_i(.)\}]$, strategy σ_i is a *best response* for player i to his rival's strategies σ_{-i} if

$$u_i(\sigma_i, \sigma_{-i}) \geqslant u_i(\sigma_i', \sigma_{-i})$$

for all $\sigma_i' \in \Delta(S_i)$. Strategy σ_i is *never a best response* if there is no σ_{-i} for which σ_i is a best response.

DEFINITION.In game $\Gamma = [I, \{\Delta(S_i)\}], \{u_i(.)\}]$, the strategies in $\Delta(S_i)$ that survive the iterated removal of strategies that are never a best response are known as player i's rationalizable strategies.

Player 2

Player 1	b_1	b_2	b_3	b_4
a_1	0, 7	2, 5	7, 0	0, 1
a_2	5, 2	3, 3	5, 2	0, 1
a_3	7, 0	2, 5	0, 7	0, 1
a_4	0, 0	0, −2	0, 0	10, −1

- $\{a_1, a_2, a_3\}$ are rationalizable strategies for player 1.
- $\{b_1, b_2, b_3\}$ are rationalizable strategies for player 2.

Unfortunately, players may have many rationalizable strategies. If we want to narrow our predictions further, we need to make additional assumptions beyond common knowledge of rationality.

Nash Equilibrium

For ease of exposition, we initially ignore the possibility that players might randomize over their pure strategies, restricting our attention to game $\Gamma = [I, \{S_i\}, \{u_i(.)\}]$.

DEFINITION(Nash equilibrium).

A strategy profile $s = (s_1, ..., s_I)$ constitutes a *Nash equilibrium* of game $\Gamma = [I, \{S_i\}, \{u_i(.)\}]$ if for every i = 1,...,I,

$$u_i(s_i, s_{-i}) \geqslant u_i(s_i', s_{-i})$$

for all $s_i' \in S_i$.

In a Nash equilibrium, each player's strategy choice is a best response to the strategies *actually played* by his rival's. The italicized words distinguish the concept of Nash equilibrium from the concept of rationalizability. Rationalizability, which captures the implication of the players' common knowledge of each others' rationality and the structure of the game, requires only that a player's strategy be a best response to some reasonable conjecture about what his rivals will be playing, where *reasonable* means that the conjectured play of his rivals can also be so justified. Nash equilibrium adds to this the requirement that players be *correct* in their conjectures.

DEFINITION(NE of mixed strategy).

A mixed strategy profile $\sigma = (\sigma_1, ..., \sigma_I)$ constitutes a Nash equilibrium of game $\Gamma = [I, \{\Delta(S_i)\}, \{u_i(.)\}]$ if for every i = 1,...,I

$$u_i(\sigma_i, \sigma_{-1}) \geq u_i(\sigma_i', \sigma_{-i})$$

for all $\sigma_i' \in \Delta(S_i)$.

PROPOSITION.

Let $S_i^+ \subset$ denote the set of pure strategies that player i plays with positive probability in mixed strategy profile $\sigma = (\sigma_1, ..., \sigma_I)$. Strategy profile σ is a *Nash equilibrium* in game $\Gamma = [I, \{\Delta(S_i)\}, \{u_i(.)\}]$ if and only if for all i=1,...I,

(i) $u_i(s_i, \sigma_{-i}) = u_i(s_i', \sigma_{-1})$ for all $s_i, s_i' \in S_i^+$;

(ii) $u_i(s_i, \sigma_{-i}) \geqslant u_i(s_i', \sigma_{-i})$ for all $s_i \in S_i^+$ and all $s_i' \notin S_i^+$.

An implication of the proposition is that to text whether a strategy profile σ is a NE it suffices to consider only pure strategy deviation(i.e., changes in a player's strategy σ_i to some pure strategy s_i). As long as no player can improve his payoff by switching to any pure strategy, σ is a NE, so:

Pure strategy profile $s = (s_1, ..., s_I)$ is a NE of game $\Gamma = [I, \{S_i\}, \{u_i(.)\}]$ iff it is a mixed strategy NE of game $\Gamma = [I, \{\Delta(S_i)\}, \{u_i(.)\}]$.

Existence of Nash Equilibria.

(1)Every game $\Gamma = [I, \{\Delta(S_i)\}, \{u_i(.)\}]$ in which the sets $S_1, ..., S_I$ have a finite number of elements has a mixed strategy NE.

(2)A NE exists in game $\Gamma = [I, \{S_i\}, \{u_i(.)\}]$ if for all i=1,...,I,

- S_i is a nonempty, convex, and compact subset of some Euclidean space ;
- $u_i(s_i, ..., s_I)$ is continuous in $(s_i, ...s_I)$ and quasiconcave is s_i.

Of course, these results do not mean that we cannot have an equilibrium if the conditions of these existence results do not hold. Rather, we just cannot be assured that there is one.

Bayesian Nash Equilibrium

- Complete information
- Incomplete information

The presence of incomplete information raises the possibility that we may need to consider a player's beliefs about other players' preferences, his beliefs about their beliefs about his preferences, and so on.

One imagines that each player's preference are determined by the realization of a random variable. Although the random variable's actual realization is observed only by the player, its ex ante probability distribution is assumed to be common knowledge among all the players. Through this formulation, the situation of incomplete information is reinterpreted as a game of imperfect information: Nature makes the first move, choosing realizations of the random variables that determine each player's preference *type*, and each player observes the realization of only his own random variable. A game of this sort is known as a *Bayesian game*.

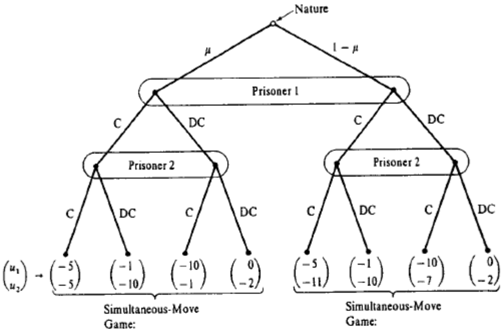

Prisoner 2

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

Prisoner 2

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

DEFINITION(Bayesian Nash equilibrium).

A (pure strategy) Bayesian Nash equilibrium for the Bayesian game $[I, \{S_i\}, \{u_i(.)\}, \Theta, F(.)]$ is a profile of decision rules $(s_1(.), ..., s_I(.))$ that constitutes a Nash equilibrium of game $\Gamma = [I, \{\overset{\smile}{S}_i\}, \{\overset{\smile}{u}_i (.)\}]$. That is, for every i=1,...,I,

$$\overset{\smile}{u}_i (s_i(.), s_{-i}(.)) \geqslant \overset{\smile}{u}_i (s_i'(.), s_{-i}(.))$$

for all $s_i'(.) \in \overset{\smile}{S}_i$,where $\overset{\smile}{u}_i (s_i(.), s_{-i}(.))$ is defined by

$$\overset{\smile}{u}_i (s_i(.), ..., s_I(.)) = E_\theta[u_i(s_1(\theta_1), ..., s_I(\theta_I), \theta_i)].$$

Proposition.

A profile of decision rules $(s_1(.), ..., s_I(.))$ is a Bayesian Nash equilibrium in Bayesian game $[I, \{S_i\}, \{u_i(.)\}, \Theta, F(.)]$ if and only if, for all i and all $\theta_i' \in \Theta_i$ occurring with positive probability

$$E_{\theta_{-i}}[u_(s_i(\theta_i'), s_{-i}(\theta_{-i}), \theta_i') \mid \theta_i'] \geqslant E_{\theta_{-i}}[u_(s_i'(\theta_i'), s_{-i}(\theta_{-i}), \theta_i') \mid \theta_i']$$

for all $s_i' \in S_i$ where the expectation is taken over realizations of the players' random variables conditional on player i's realization of his signal θ_i'.

To solve for the (pure strategy) Bayesian Nash equilibrium of this game, note first that type I of prison 2 must play confess with probability 1 because this is that type's dominant strategy. Likewise, Type II of prison 2 also has a dominant strategy: don't confess. Given this behavior by prison 2, prison 1's best response is to play don't confess if

$$[-10\mu + 0(1-\mu)] > [-5\mu - 1(1-\mu)],$$

or equivalently, if $\mu < \frac{1}{6}$, and is to play confess if $\mu > \frac{1}{6}$(He is indifferent if $\mu = \frac{1}{6}$)

Dynamic Game

- the credibility of a player's strategy
- subgame perfect Nash equilibrium
- a perfect Bayesian equilibrium(sequential equilibrium)
- forward induction

Subgame Perfection,Back Induction

We begin with an example to illustrate that in dynamic games the NE concept may not give sensible predictions.

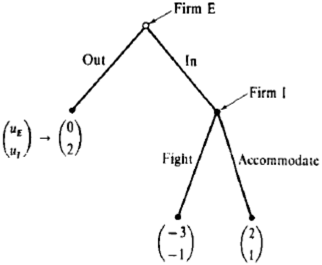

		Firm I	
		Fight if Firm E Plays "In"	Accommodate if Firm E Plays "In"
Firm E	Out	0, 2	0, 2
	In	−3, −1	2, 1

The NE (out, fight if firm E plays in) is not credible.

DEFINITION(Subgame).

A subgame of an extensive form game Γ is a subset of the game having the following properties:

(i)It begins with an information set containing a single decision node, contains all the decision nodes that are successors(both immediate and later) of this node, and contain only there nodes.

(ii) If decision node x is in the subgame, then every $x' \in H(x)$ is also, where H(x) is the information set that contains decision node x.

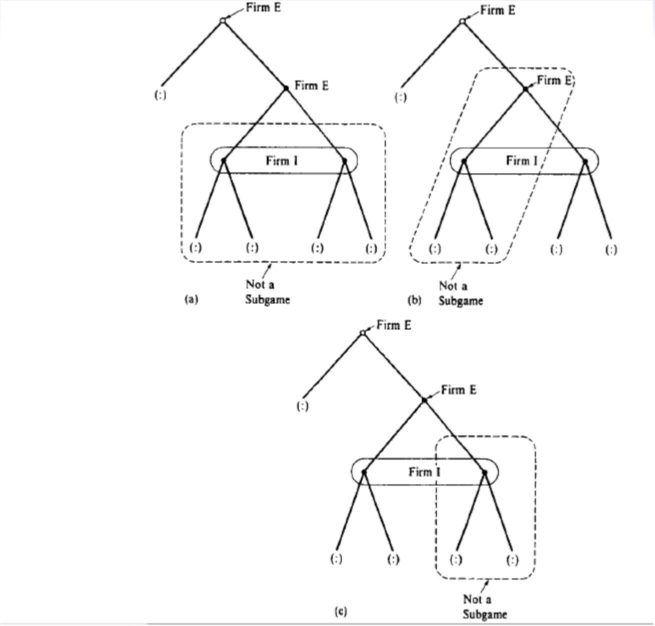
Firm E
Firm E
Firm I
(:) (:) (:) (:) (:)
Not a Subgame
(a)
Firm E
Firm E
Firm I
(:) (:) (:) (:) (:)
Not a Subgame
(b)
Firm E
Firm E
Firm I
(:) (:) (:) (:) (:)
Not a Subgame
(c)

DEFINITION.(Subgame perfect Nash equilibrium)

A profile of strategies $\sigma = (\sigma_1, ..., \sigma_I)$ is an I-player extensive form game Γ is a SPNE if it induce a NE in every subgame of Γ.

Any SPNE is a NE, but that not every NE is subgame perfect.

In fact, to identify the set of subgame perfect NE in a general(finite) dynamic gameΓ, we can use the backward induction procedure.

There is, however, an interesting tension present in the SPNE. In particular, the SPNE concept insists that players should play an SPNE wherever they find themselves in the game tree, even after a sequence of events that is contrary to the predictions of the theory(principle of sequential).

The Centipede game.

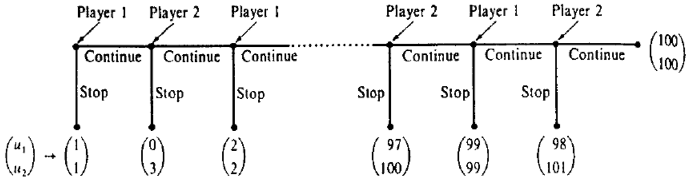

Beliefs and Sequential Rationality

Although subgame perfection is often very useful in capturing the principle of sequential rationality, sometimes it is not enough.

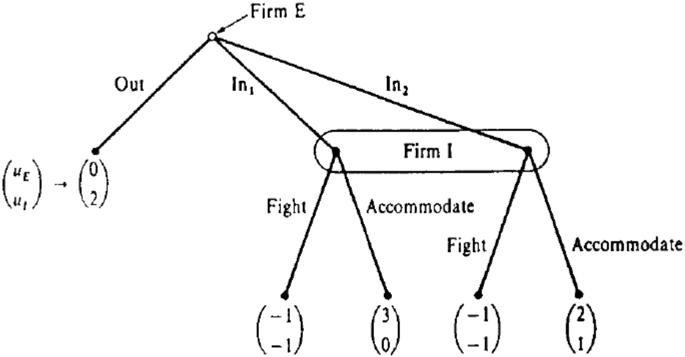

The criterion of subgame perfection is of absolutely if of absolutely no use: the only subgame is the game as a whole.

DEFINITION(System of belief).

A *system of beliefs* μ in extensive form game Γ is a specification of a probability $\mu(x) \in [0, 1]$ for each decision node x in Γ such that:

$$\Sigma_{x \in H} \mu(x) = 1$$

for all information sets H.

DEFINITION.(Sequential Rationality)

A strategy profile $\sigma = (\sigma_1, ..., \sigma_I)$ in extensive form game Γ is *sequentially rational at information set H given a system of beliefs* μ if, denoting by $\iota(H)$ the player who moves at information set H, we have

$$E[u_{\iota(H)} \mid H, \mu, \sigma_{\iota(H)}, \sigma_{-\iota(H)}] \geq E[u_{\iota(H)} \mid H, \mu, \sigma'_{\iota(H)}, \sigma_{-\iota(H)}]$$

for all $\sigma'_{\iota(H)} \in \Delta(S_{\iota(H)})$.If strategy profile σ satisfies this condition for all information sets H, then we say that σ is *sequentially rational given belief system* μ.

DEFINITION.(weak PBE)

A profile of strategies and system of belief (σ, μ) is a *weak perfect Bayesian equilibrium (weak PBE)* in extensive form game Γ if it has the following properties:

(i) The strategy profile σ is sequentially rational given belief system μ.

(ii)The system of beliefs μ is derived from strategy profile σ though Bayes' rule whenever possible. That is,for any information set H such that $Prob(H \mid \sigma) > 0$, we must have $\mu(x) = \frac{Prob(x|\sigma)}{Prob(H|\sigma)}$ for all $x \in H$.

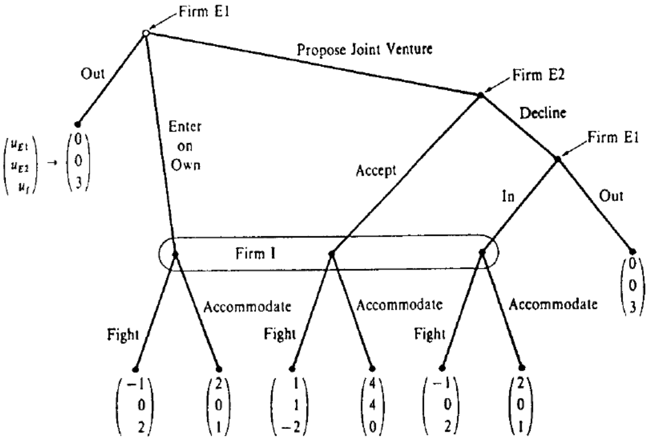
Firm E1
Propose Joint Venture
Out
Firm E2
Decline
Firm E1
Enter on Own
Accept
In
Out
Firm I
Accommodate
Fight
Accommodate
Fight
Accommodate
Fight
u_{E1}, u_{E2}, u_I → (0, 0, 3)
(−1, 0, 2)
(2, 0, 1)
(1, 1, −2)
(4, 4, 0)
(−1, 0, 2)
(2, 0, 1)
(0, 0, 3)

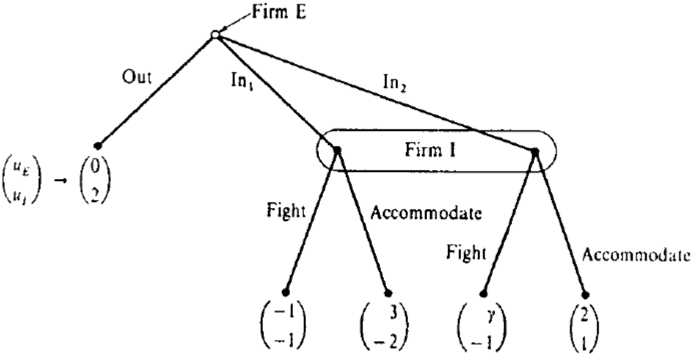
Firm E
Out
In₁
In₂
Firm I
(u_E, u_I) → (0, 2)
Fight
Accommodate
Fight
Accommodate
(−1, −1)
(3, −2)
(γ, −1)
(2, 1)

$\gamma > -1$

Reasonable Beliefs and Forward Induction

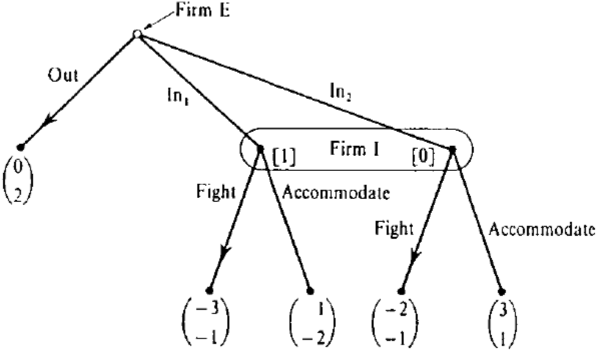

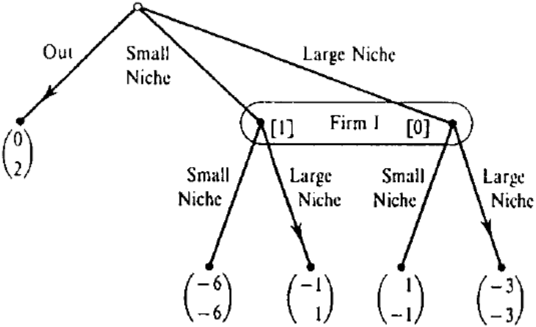

Out
Small Niche
Large Niche
0
2
[1]
Firm I
[0]
Small Niche
Large Niche
Small Niche
Large Niche
−6
−6
−1
1
1
−1
−3
−3

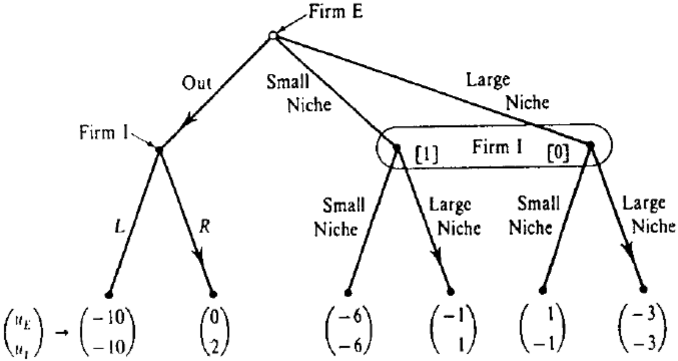
Firm E
Out
Small Niche
Large Niche
Firm I
[1]
Firm I
[0]
L
R
Small Niche
Large Niche
Small Niche
Large Niche
$\begin{pmatrix} u_E \\ u_I \end{pmatrix} \rightarrow \begin{pmatrix} -10 \\ -10 \end{pmatrix}$
$\begin{pmatrix} 0 \\ 2 \end{pmatrix}$
$\begin{pmatrix} -6 \\ -6 \end{pmatrix}$
$\begin{pmatrix} -1 \\ 1 \end{pmatrix}$
$\begin{pmatrix} 1 \\ -1 \end{pmatrix}$
$\begin{pmatrix} -3 \\ -3 \end{pmatrix}$

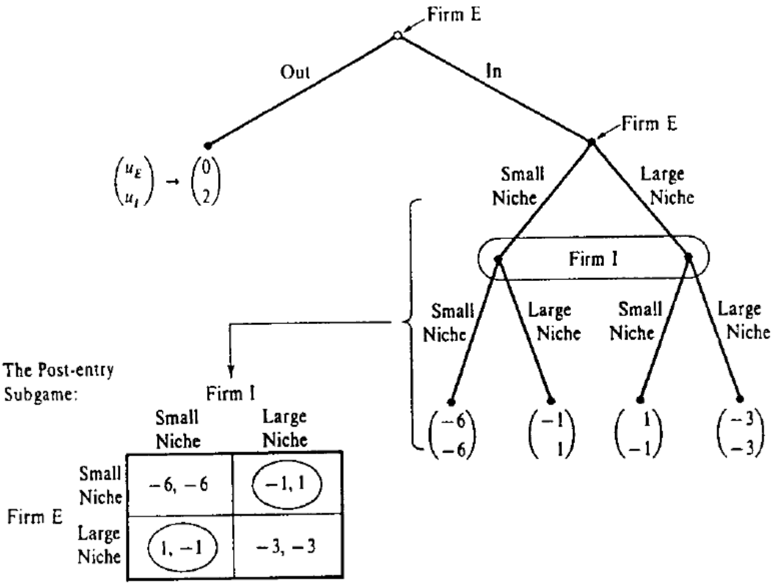
Firm E
Out
In
Firm E
Small Niche
Large Niche
Firm I
Small Niche
Large Niche
Small Niche
Large Niche
The Post-entry Subgame:
Firm I
Small Niche
Large Niche
Firm E
Small Niche
Large Niche
−6, −6
−1, 1
1, −1
−3, −3