2.- Game Theory Overview

Game Theory studies situations in which the outcome of the decisions of an individual depends on the decisions taken by other individuals as well as on his own

- Strategic Interdependence -

Plan

2.1 Introduction

Any *(non cooperative)* game has the following elements

1. The Players Who is to take decisions

2 The Rules Who moves first, who moves next, what can be done and what can not, ...

3 The Strategies What actions are available to the players

4 The Outcomes What is the consequence of each combination of actions by the players

5 The Payoffs What each player obtains in each possible outcome

Example: Rock - Scissors - Paper

1. The Players *Player 1* and *Player 2*

2 The Rules Simultaneously, both players show either *rocks*, *scissors* or *paper*

3 The Strategies Each player can choose among showing *rocks*, *scissors* or *paper*

4 The Outcomes *Rock* beats *scissors*, which beats *paper*, which beats *rock*

5 The Payoffs The winning player receives 1 from the loser (thus getting -1)

Game Scenarios

	Complete Information	Incomplete Information
Static	Players have all relevant information and move simultaneously	Not all players have all the information and move simultaneously
Dynamic	Players have all relevant information and move in turns	Not all players have all the information and move in turns

Examples

	Complete Information	Incomplete Information
Static	Rock, Scissors, Paper	Sealed-bid Auction
Dynamic	Chess	English Auction

2.2 Static Games of Complete Information

Players take decisions without knowing the actions taken by other players

Basic Assumptions

1. Players are rational and self-interested (Rationality)

2. Players have full information about all the elements of the game (Complete Information)

3. Players know that players are rational and fully informed, and they know that others know that players are rational and fully informed, and they know that others know that they know that players are rational and fully informed, and ... (Common Knowledge)

Common Knowledge examples

1. The Byzantine Generals and the Coordinated Attack Problem
2. The Three Hats Problem
3. The Princes Bride
4. La vida es sueño (Pedro Calderón de la Barca - 1636)

Notation

- $\mathcal{I} = \{1, 2, \ldots, I\}$ is the *Set of Players*
- $S^i = \{s_1^i, s_2^i, \ldots\}$ is the *Set of Pure Strategies of player* $i \in \mathcal{I}$
- $S = S^1 \times S^2 \times \cdots \times S^I$ is the *Set of pure strategy profiles*
- $(s_i^*,\, s_{-i}) = (s_1, s_2, \ldots, s_{i-1}, s_i^*, s_{i+1}, \ldots, s_I)$
- $u^i : S \to \mathbb{R}$ is the *(von Neumann - Morgenstern) Payoff function of player* $i \in \mathcal{I}$

A *Game* is fully characterized

$$\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$$

This representation of a game is know a normal form game or strategic form game,.

Example 1: The Prisoners' Dilemma

- 2 players $\mathcal{I} = \{1, 2\}$
- 2 (identical) strategies for each player $S^i = \{C, D\}$ (Cooperate, Defect) $i \in \mathcal{I}$
- Strategy profiles in $S = \{C, D\} \times \{C, D\} = \{(C, C), (C, D), (D, C), (D, D)\}$
- Symmetric payoffs given by
  - $u^1(C, C) = 3 \qquad u^2(C, C) = 3$
  - $u^1(C, D) = 0 \qquad u^2(C, D) = 5$
  - $u^1(D, C) = 5 \qquad u^2(D, C) = 0$
  - $u^1(D, D) = 1 \qquad u^2(D, D) = 1$

This is usually represented by a *bi-matrix table*

Example 1: The Prisoners' Dilemma (continued)

$1\backslash 2$	C	D
C	$3,3$	$0,5$
D	$5,0$	$1,1$

Independently of the other player strategy, each player will strictly prefer to play $Defect\,(D)$

Any game satisfying the relationships below is "a prisoners' dilemma"

$1\backslash 2$	s_1^2	s_2^2
s_1^1	R,R	S,T
s_2^1	T,S	P,P

with $T > R > P > S$

$T \longrightarrow$ Temptation payoff $R \longrightarrow$ Reward payoff
$P \longrightarrow$ Punishment payoff $S \longrightarrow$ Sucker payoff

This is a game in which there are gains from cooperation, but a binding agreement is not possible (neither stable !). Situations such as price-fixing/collusion and free-riding can be understood using a prisoner's dilemma model.

It is a compelling example that individual self-interest might not lead to social efficiency

Example 2: The Battle of the Sexes

$\mathcal{I}$ $=\{1, 2\}$, $S^i = \{F, O\}$ (Football, Opera)

$1\backslash 2$	F	O
F	$2, 1$	$0, 0$
O	$0, 0$	$1, 2$

The players want to coordinate, but have conflicting interests. Any coordination is a possible result.

Example 3: Meeting in Barcelona

$\mathcal{I}$ $=\{1, 2\}$, $S^i = \{P, R\}$ (Plaça Catalunya, Rambles)

$1\backslash 2$	P	R
P	$2, 2$	$0, 0$
R	$0, 0$	$1, 1$

The players want to coordinate, and have common interests. Any coordination is a possible result, but one of them is "more possible" than the other (Schelling's Focal Point)

Example 4: Rock-Scissors-Paper

$\mathcal{I} = \{1, 2\}$, $S^i = \{R, S, P\}$ (Rock, Scisors, Paper)

$1\backslash 2$	R	S	P
R	$0, 0$	$1, -1$	$-1, 1$
S	$-1, 1$	$0, 0$	$1, -1$
P	$1, -1$	$-1, 1$	$0, 0$

In this case, what is possible outcome of the game ? (Zero-Sum Game)

Main Question

- What behavior should we expect form a rational, fully informed, and self-interested player ?
- If all players behave in such way, what will be the outcome of the game ?

- EQUILIBRIUM -

- People (players) don't do silly things (elimination)
- People's behavior will be *stable* if nobody has incentives to change

The Equilibrium

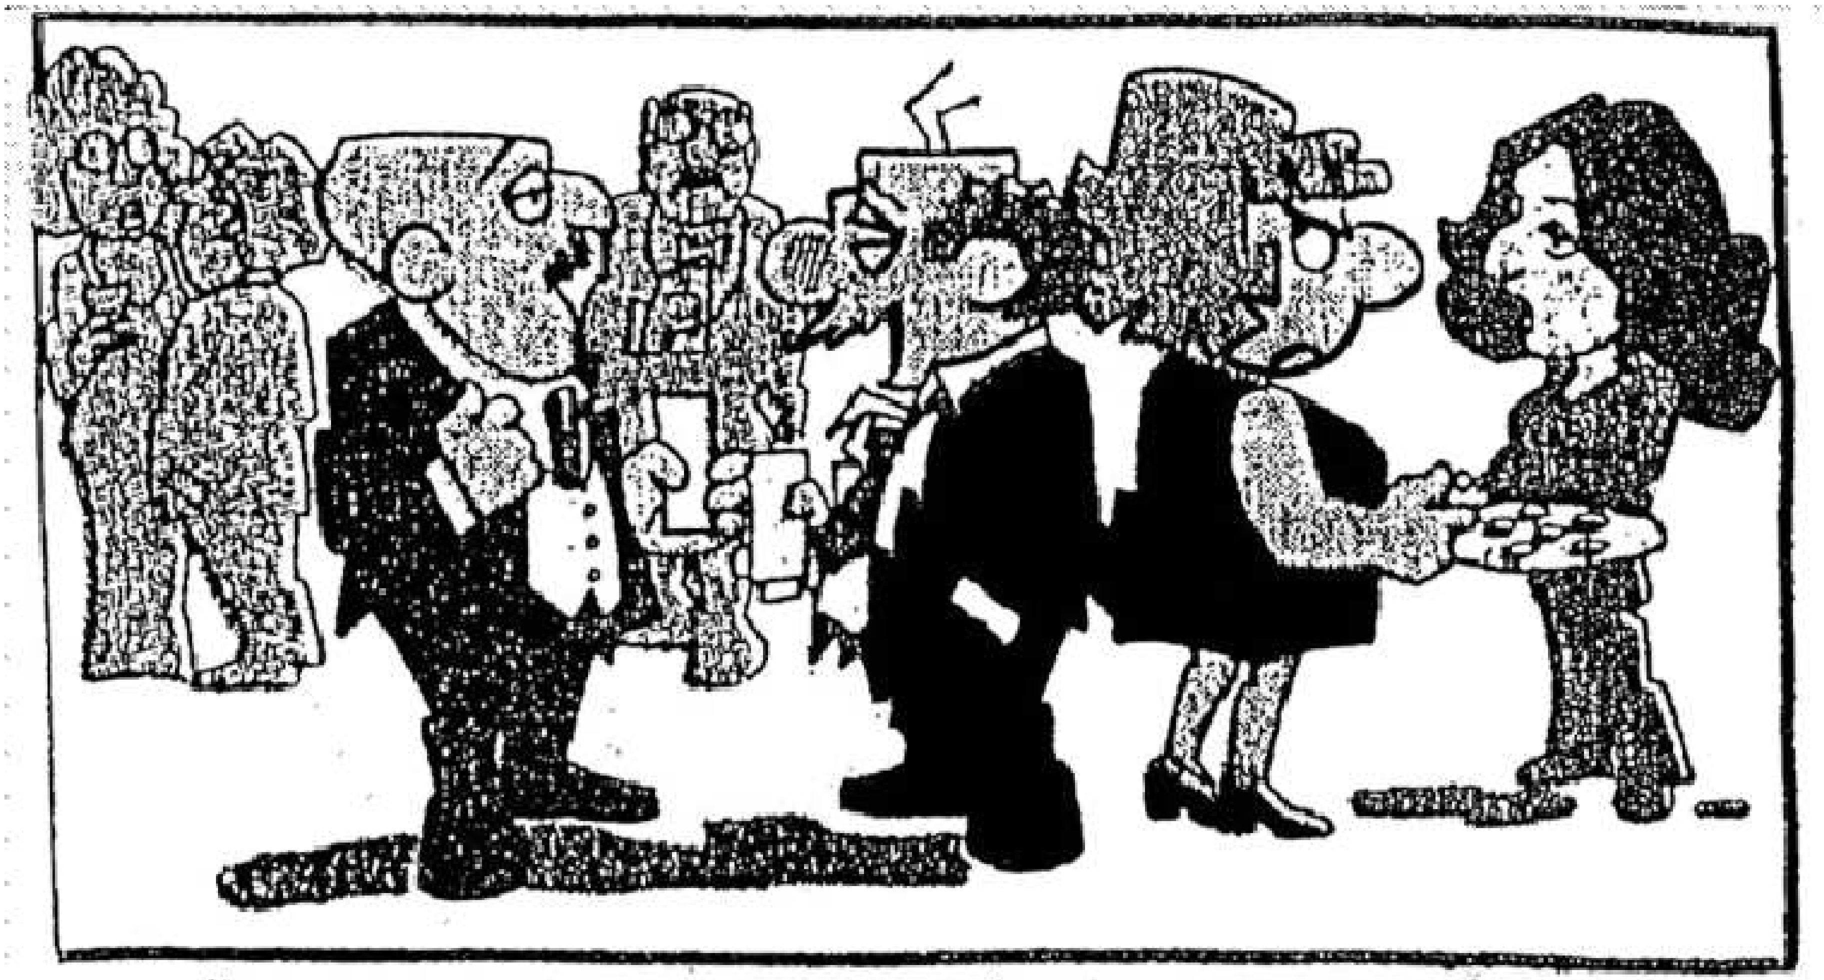

"LORETTA'S DRIVING BECAUSE I'M DRINKING, AND I'M DRINKING BECAUSE SHE'S DRIVING."

Dominant and Dominated Strategies

The simplest way to consider rational behavior by the players is that of *Dominant* and *Dominated* Strategies.

Consider the static game

$$\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$$

Definition. *A strategy $s_i \in S_i$ is a* Strictly Dominant Strategy *for player $i \in \mathcal{I}$ if, for any other $s' \in S_i$ ($s'_i \neq s_i$),*

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

for all $s_{-i} \in S_{-i}$

Example The strategy D (Defect) is a *Strictly Dominant Strategy* for both players in the Prisoners' Dilemma

Definition. *A strategy* $s_i \in S_i$ *is a* Strictly Dominated Strategy *for player* $i \in \mathcal{I}$ *if there exists another strategy* $s' \in S_i$ $(s'_i \neq s_i)$ *such that, for all* $s_{-i} \in S_{-i}$

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

In such case, we say that the strategy s'_i *strictly dominates* s_i for player i

Example The strategy C (Cooperate) is a *Strictly Dominated Strategy* for both players in the Prisoners' Dilemma. (The strategy D dominates C)

Note A strategy $s_i \in S_i$ is a *Strictly Dominant Strategy* for player $i \in \mathcal{I}$ if it *strictly dominates* every other strategy of player i

Definition. *A strategy* $s_i \in S_i$ *is a* Weakly Dominant Strategy *for player* $i \in \mathcal{I}$ *if, for any other* $s' \in S_i$ $(s'_i \neq s_i)$,

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$$

for all $s_{-i} \in S_{-i}$

Definition. *A strategy* $s_i \in S_i$ *is a* Weakly Dominated Strategy *for player* $i \in \mathcal{I}$ *if there exists another strategy* $s' \in S_i$ $(s'_i \neq s_i)$ *such that, for all* $s_{-i} \in S_{-i}$

$$u_i(s'_i, s_{-i}) \geq u_i(s_i, s_{-i})$$

In such case, we say that the strategy s'_i *weakly dominates* s_i for player i

Example

$1\backslash 2$	L	R
T	$1, 3$	$0, 2$
B	$0, 2$	$0, 3$

T weakly dominates B. That is, B is weakly dominated by T. Or, T is a *weakly dominant* strategy

Comments

- A strictly dominant strategy, if it exists, must be unique
- Weakly dominant strategies, if any, might not be unique
- Dominated strategies (both strictly and weakly) might not be unique

Iterated elimination of (strictly) dominated strategies

Strictly dominated strategies should never be used by a rational player. The assumption of *common knowledge* allows players to effectively "delete" such strategies. Sometimes, this leads to some reduction of the possible outcomes of the game

Example

$1\backslash 2$	L	C	R
T	$3,0$	$0,-5$	$0,-4$
M	$1,-1$	$3,3$	$-2,4$
B	$2,4$	$4,1$	$-1,8$

R strictly dominates C. Hence,

$1\backslash 2$	L	R
T	$3,0$	$0,-4$
M	$1,-1$	$-2,4$
B	$2,4$	$-1,8$

B strictly dominates M. Hence

$1\backslash 2$	L	R
T	$3,0$	$0,-4$
B	$2,4$	$-1,8$

T strictly dominates B. Hence

$1\backslash 2$	L	R
T	$3,0$	$0,-4$

Finally, L strictly dominates R. Thus, we finally have

$1\backslash 2$	L
T	$3,0$

The final outcomes of the game are reduced to (only) (T,L) which produces payoffs $(3,0)$

Comments

- Common knowledge is needed for this to make sense
- In many cases it is not of any help (Battle of the Sexes)
- The order of elimination is not important for the final result
- Strategies that are dominant at the origin will remain dominant through all the process

Iterated elimination of (weakly) dominated strategies

The "deletion" of *weakly* dominated strategies does not have as much sense as in the case of strictly dominated strategies. The reason is that, in this case, we can not say that "a rational player will never use such strategy" as before. Furthermore, the process might produce a strange result

Example

$1\backslash 2$	L	R
T	$5,1$	$4,0$
M	$6,0$	$3,1$
B	$6,4$	$4,4$

Strategies T and M are weakly dominated

a) If T is removed,

$1\backslash 2$	L	R
M	$6, 0$	$3, 1$
B	$6, 4$	$4, 4$

Now L is weakly dominated,

$1\backslash 2$	R
M	$3, 1$
B	$4, 4$

and then M can be eliminated to finally obtain

$1\backslash 2$	R
B	$4, 4$

b) If M is removed,

$1\backslash 2$	L	R
T	$5, 1$	$4, 0$
B	$6, 4$	$4, 4$

Now R is weakly dominated

$1\backslash 2$	L
T	$5, 1$
B	$6, 4$

And now T can be deleted to finally get

$1\backslash 2$	L
B	$6, 4$

Rationalizable Strategies

Consider the following game

$1\backslash 2$	L	R
T	$4,0$	$0,2$
M	$1,2$	$2,1$
B	$0,1$	$3,2$

Notice that,

- No player has a Dominant Strategy
- No player has a Dominated Strategy
- Yet, strategy M for the first player is never going to be used (It is never a good option by a rational player)

We could extend the concept of *Dominated Strategy* to strategies that are not going to be used by a rational player.

Definition. *A strategy* $s_i \in S_i$ *is a* best response *for player* $i \in \mathcal{I}$ *to the strategies* s_{-i} *of his opponents if*

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$$

for any $s'_i \in S_i$

Definition. *A strategy* $s_i \in S_i$ *is* never a best response *if there is no* s_{-i} *for which* s_i *is a* best response

We can, therefore, replace the process of iterated elimination of *dominated strategies* by a process of iterated elimination of *strategies that are never a best response.* Notice that:

- A *strictly dominated strategy* is *never a best response*
- A strategy that is *never a best response* might not be a *strictly dominated strategy*

Thus, more strategies are eliminated in this process o elimination of strategies that are never a best response

Definition. *The strategies in* S_i *that survive the process of iterated elimination of strategies that are never a best response are called* player's i rationalizable strategies

Alternatively, we can say that player's i rationalizable strategies *are those strategies in* S_i *that are a best response to some* s_{-i} *by player's* i *opponents.*

In the previous example, if we eliminate M *for player 1 we get*

$1\backslash 2$	L	R
T	$4,0$	$0,2$
B	$0,1$	$3,2$

Now L is strictly dominated for player 2

$1\backslash 2$	R
T	$0,2$
B	$3,2$

Now we can eliminate T

$1\backslash 2$	R
B	$3,2$

Thus, rationalizable strategies are B for player 1 and R for player 2

Nash Equilibrium

It is clear that in many cases the process of iterated elimination of strategies cannot be started and, thus, no prediction can be made about the outcome of the game

The Battle of the Sexes

$1\backslash 2$	F	O
F	$2,1$	$0,0$
O	$0,0$	$1,2$

Yet, it is clear that *rational* outcomes are (F,F) and (O,O), and that none of them is "more rational" that the other.

The two outcomes might derive from a process of *self-fulfilling conjectures* made by *rational players* under *common knowledge*

- If my partner chooses F, F is my best option, and if F is my best option, my partner will chose F (Same goes for O) [Stable]

- If my partner choose F,O is NOT my best option ... [Unstable]

Definition. *Given a game* $\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$, *a* strategy profile $s^* \in S$ *is a* Nash Equilibrium *if for every* $i \in \mathcal{I}$,

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*)$$

for any other $s_i \in S_i$

Thus, given that all are playing as in s^*, everybody's best option is to play as in s^*.

Definition. *Given a game* $\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$, *the* best response correspondence *for player* i, $b_i : S_{-i} \to S_i$, *is given by*

$$b_i(s_{-i}) = \{s_i \in S_i \,|\, u_1(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \text{ for all } s_i' \in S_i\}$$

Consequently,

Definition. *Given a game* $\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$, *a* strategy profile $s^* \in S$ *is a* Nash Equilibrium *if for every* $i \in \mathcal{I}$,

$$s_i^* \in b_i(s_{-i}^*)$$

Examples

Battle of the Sexes

$1\backslash 2$	F	O
F	$2,1$	$0,0$
O	$0,0$	$1,2$

Two Nash Equilibria: (F, F) and (O, O)

Meeting in Barcelona

$1\backslash 2$	P	R
P	$2,2$	$0,0$
R	$0,0$	$1,1$

Two Nash Equilibria: (P, P) and (R, R)

Comments

- A Nash equilibrium might not exist (for the moment ...)
- A Nash equilibrium might not be unique
- A Nash equilibrium needs not to be *efficient*
- In a Nash equilibrium no player has incentives to deviate
- Nash equilibrium are *minimal* requirements for strategic stability and self-enforcement
- Strictly dominant strategies are always present in a Nash equilibrium.
- No strictly dominated strategy can be part of a Nash equilibrium.
- Weakly dominated strategies may be part of a Nash equilibrium (!)

Mixed Strategies

How to play ?

$1\backslash 2$	R	S	P
R	$0,0$	$1,-1$	$-1,1$
S	$-1,1$	$0,0$	$1,-1$
P	$1,-1$	$-1,1$	$0,0$

If my opponent plays ...

$$p(R) = p(S) = p(P) = \frac{1}{3}$$

I get ...

$$\begin{aligned} u(R) &= \frac{1}{3}0 + \frac{1}{3}(-1) + \frac{1}{3}1 = 0 \\ u(S) &= \frac{1}{3}(-1) + \frac{1}{3}0 + \frac{1}{3}1 = 0 \\ u(P) &= \frac{1}{3}1 + \frac{1}{3}(-1) + \frac{1}{3}0 = 0 \end{aligned}$$

Thus, *any* strategy is a best response to $p(R) = p(S) = p(P) = \frac{1}{3}$.

But ... Is $p(R) = p(S) = p(P) = \frac{1}{3}$ my opponent's best response to any of my strategies ?

Only playing $p(R) = p(S) = p(P) = \frac{1}{3}$ is best response to playing $p(R) = p(S) = p(P) = \frac{1}{3}$.

That is, only $p(R) = p(S) = p(P) = \frac{1}{3}$ is a Nash Equilibrium.

It's an equilibrium in *mixed strategies*

Notation

- A *mixed strategy* for player $i \in \mathcal{I}$ is a probability distribution over $S^i = \{s_1^i, s_2^i, \ldots\}$

$$\sigma^i = (\sigma_i^i, \sigma_2^i, \ldots),\ \sigma_j^i \geq 0,\ \sum \sigma_j^i = 1$$

- The set of all *mixed strategies* of player i is the *Simplex* on S^i, denoted $\Delta(S^i)$

- $\Delta(S) = \Delta(S^1) \times \Delta(S^2) \times \cdots \times \Delta(S^I)$ is the *Set of strategy profiles*

- $(\sigma^{i*},\ \sigma^{-i}) = (\sigma^1, \sigma^2, \ldots, \sigma^{i-1}, \sigma^{i*}, , \ldots, \sigma^I)$

- $u^i : S \rightarrow \mathbb{R}$ is the *(von Neumann - Morgenstern) Payoff function of player* $i \in \mathcal{I}$

$$u^i(\sigma^i) = \sigma_1^i u^i(s_1^i) + \sigma_2^i u^i(s_2^i) + \cdots$$

Example: The Battle of the Sexes

$1\backslash 2$	F	O
F	$2, 1$	$0, 0$
O	$0, 0$	$1, 2$

We know that $\sigma^1 = \sigma^2 = (1, 0)$ and $\sigma^1 = \sigma^2 = (0, 1)$ are two Nash equilibria.

For another strategy profile (σ^1, σ^2) to be an equilibrium, it must be the case that

$$\begin{aligned} u^1(F, \sigma^2) &= u^1(O, \sigma^2) \\ u^2(\sigma^1, F) &= u^2(\sigma^1, O) \end{aligned}$$

That is,

$$\begin{aligned} 2\sigma_1^2 + 0(1 - \sigma_1^2) &= 0\sigma_1^2 + 1(1 - \sigma_1^2) \\ 1\sigma_1^1 + 0(1 - \sigma_1^1) &= 0\sigma_1^1 + 2(1 - \sigma_1^1) \end{aligned}$$

Solving we get that

$$\sigma^1 = (\frac{2}{3}, \frac{1}{3}) \text{ and } \sigma^2 = (\frac{1}{3}, \frac{2}{3})$$

is also a Nash equilibrium.

Interestingly...

		$\frac{1}{3}$	$\frac{2}{3}$
	$1\backslash 2$	F	O
$\frac{2}{3}$	F	$\frac{2}{9}$	$\frac{4}{9}$
$\frac{1}{3}$	O	$\frac{1}{9}$	$\frac{2}{9}$

there will be more miss-coordinations that coordinations

Nash Equilibrium

Definition. *Given a game* $\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$*, a* strategy profile $\sigma^* \in \Delta(S)$ *is a* Nash Equilibrium *if for every* $i \in \mathcal{I}$*,*

$$u^i(\sigma^{i*}, \sigma^{-i*}) \geq u^i(\sigma^i, \sigma^{-i*})$$

for any other $\sigma^i \in S^i$

Thus, given that all are playing as in σ^*, everybody's best option is to play as in σ^*.

Definition. *Given a game* $\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$*, the* best response correspondence *for player* i*,* $b^i : \Delta(S^{-i}) \to \Delta(S^i)$*, is given by*

$$b^i(\sigma^{-i}) = \{\sigma^i \in \Delta(S^i) \,|\, u^i(\sigma^i, \sigma^{-i}) \geq u^i(\sigma'^i, \sigma^{-i}) \text{ for all } \sigma'^i \in \Delta(S^i)\}$$

Consequently,

Definition. *Given a game* $\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$*, a* strategy profile $\sigma^* \in S$ *is a* Nash Equilibrium *if for every* $i \in \mathcal{I}$*,*

$$\sigma^{i*} \in b^i(\sigma^{-i*})$$

Existence of Nash Equilibrium

Theorem 1. *Any finite game has a mixed strategy Nash equilibrium*

Theorem 2. *(Kakutani's fix point theorem)*
Let $f : A \rightrightarrows A$ *be a correspondence, with* $x \in A \Rightarrow f(x) \subset A$*, satisfying the following conditions:*

(i) A *is a compact, convex, and non-empty subset of a finite dimensional Euclidean space*

(ii) $f(x)$ *is non-empty:* $\forall x \in A,\ f(x)$ *is well defined*

(iii) $f(x)$ *is convex:* $\forall x \in A,\ f(x)$ *is a convex-valued correspondence*

(iv) $f(x)$ *is an upper-hemicontinous correspondence: If* $\{x^n, y^n\} \rightarrow \{x, y\}$ *with* $y^n \in f(x^n)$*, then* $y \in f(x)$ *(the graph of* f *is closed), and the image of any compact set is bounded*

Then, $\exists x^* \in A$ *such that* $x^* \in f(x^*)$

Existence of Nash Equilibrium

Proof. Let $b : \Delta(S) \rightrightarrows \Delta(S)$ be a correspondence defined by:

$$b(\sigma) = (b^1(\sigma^{-1}), b^2(\sigma^{-2}), \ldots, b^I(\sigma^{-I})), \quad \forall \sigma \in \Delta(S)$$

The proof proceeds by showing that $b(\sigma)$ satisfies the conditions of Kakutani's theorem to conclude that $\exists \sigma^* \in \Delta(S)$ such that $\sigma^* \in b(\sigma^*)$

(i) $\Delta(S)$ is compact, convex, and non-empty

This is true by definition of $\Delta(S)$

(ii) $b(\sigma)$ is non-empty

By definition, $b^i(\sigma^{-i}) = \{\sigma^i \in \Delta(S^i) \,|\, u^i(\sigma^i, \sigma^{-i}) \geq u^i(\sigma'^i, \sigma^{-i})$ for all $\sigma'^i \in \Delta(S^i)\}$, that is, $b^i(\sigma^{-i})$ is the solution to the maximization of a continuous function u^i on a non-empty and compact set. Then, by the Weirstrass theorem, $b^i(\sigma^{-i})$ is non-empty

(iii) $b(\sigma)$ is convex

We need to prove that $b^i(\sigma^{\cdot i})$ is convex $\forall i \in \mathcal{I}$. Let σ'^i, $\sigma''^i \in b^i(\sigma^{-i})$. Then,

$$\begin{aligned} u^i(\sigma'^i, \sigma^{-i}) &\geq u^i(\mu^i, \sigma^{-i}) \quad \forall \mu^i \in \Delta(S^i) \\ u^i(\sigma''^i, \sigma^{-i}) &\geq u^i(\mu^i, \sigma^{-i}) \quad \forall \mu^i \in \Delta(S^i) \end{aligned}$$

Then, $\forall \lambda \in [0, 1]$,

$$\lambda u^i(\sigma'^i, \sigma^{-i}) + (1-\lambda) u^i(\sigma''^i, \sigma^{-i}) \geq u^i(\mu^i, \sigma^{-i}) \quad \forall \mu^i \in \Delta(S^i)$$

Now, since $u^i(\cdot, \sigma^{-i})$ is linear in σ^i,

$$u^i(\lambda\sigma'^i + (1-\lambda)\sigma''^i, \sigma^{-i}) \geq u^i(\mu^i, \sigma^{-i}) \quad \forall \mu^i \in \Delta(S^i)$$

Therefore, $\lambda\sigma'^i + (1-\lambda)\sigma''^i \in b^i(\sigma^{-i})$ and thus $b(\sigma)$ is convex

(iv) $b(\sigma)$ is an upper-hemicontinous correspondence

Since $b(\sigma) : \Delta(S) \rightrightarrows \Delta(S)$, and $\Delta(S)$ is a compact set, it is clear that the image of any compact set will be bounded. Hence, to show that $b(\sigma)$ is upper-hemicontinuous we only need to prove that the graph is closed. Assume to the contrary that there exists a convergent sequence of pairs of mixed strategy profiles $(\sigma^n, \hat{\sigma}^n) \rightarrow (\sigma, \hat{\sigma})$ such that $\forall n \in \{1, 2, \ldots\} \quad \hat{\sigma}^n \in b(\sigma^n)$ but $\hat{\sigma} \notin b(\sigma)$. Along the sequence, because $\hat{\sigma}^n \in b(\sigma^n)$, for each player $i \in \mathcal{I}$ we have that

$$\hat{\sigma}^{n^i} \in b^i(\sigma^{n^{-i}})$$

At the limit, because we are assuming that $\hat{\sigma}$ is not a best response to σ, at least one player must have a better strategy than $\hat{\sigma}^i$ against σ^{-i}. That is, because $\hat{\sigma} \notin b(\sigma), \quad \exists i \in \mathcal{I}, \ \exists \bar{\sigma}^i \in \Delta(S^i)$ such that

$$u^i(\bar{\sigma}^i, \sigma^{-i}) > u^i(\hat{\sigma}^i, \sigma^{-i})$$

Now, since $\sigma^n \rightarrow \sigma$ we have that $\sigma^{n^{-i}} \rightarrow \sigma^{-i}$. Thus, since u^i is continuous in σ (it's linear), we have that for n large enough $u^i(\bar{\sigma}^i, \sigma^{n^{-i}})$ will be arbitrarily close to $u^i(\bar{\sigma}^i, \sigma^{-i})$. Analogously, since $(\sigma^n, \hat{\sigma}^n) \rightarrow (\sigma, \hat{\sigma})$, we have that for n large enough $u^i(\hat{\sigma}^{n^i}, \sigma^{n^{-i}})$ will be arbitrarily close to $u^i(\hat{\sigma}^i, \sigma^{-i})$. Alltogether we have that, for n sufficiently large,

$$u^i(\bar{\sigma}^i, \sigma^{n^{-i}}) > u^i(\hat{\sigma}^{n^i}, \sigma^{n^{-i}})$$

that contradicts $\hat{\sigma}^{n^i} \in b^i(\sigma^{n^{-i}})$. Therefore, $b(\sigma)$ must be upper-hemicontinuous.

Thus, we have verified that the mixed strategy profiles set $\Delta(S)$ together with the best response correspondence b satisfy the hypothesis of the Kakutani's fixed-point theorem. Therefore, $\exists \sigma^* \in \Delta(S)$ such that

$$\sigma^* \in b(\sigma^*)$$

that is, σ^* is a Nash equilibrium

□

Correlated Equilibrium

Recall the *mixed strategy* equilibrium in the *Battle of the Sexes* game.

$1\backslash 2$	F	O
F	$2,1$	$0,0$
O	$0,0$	$1,2$

$$\sigma^{*1} = (\frac{2}{3}, \frac{1}{3}) \text{ and } \sigma^{*2} = (\frac{1}{3}, \frac{2}{3})$$

A possible interpretation of this equilibrium is that is corresponds to a *steady state* of a stochastic process according to which each player plays a best response to the beliefs about the future behavior of other players based on some "signal" received from nature (for instance, the empirically observed frequency of the play of each player in the past) In equilibrium these frequencies must remain constant over time and are stable in the sense that any action taken with positive probability by a player is optimal given the steady state beliefs.

An assumption that underlies this steady state interpretation is that no player detects any correlation among the other players' actions or between the other players' actions and his own behavior.

full correlation

One of the main assumptions of non-cooperative game theory is that players take actions *independently*, without communicating, without any sort of coordination. On the other extreme would be the case when players can talk to coordinate their actions. Aumann (1974) suggested that between "no coordination" (no communication) and "full coordination" (talking) there might exist a range of correlations between their actions. This is the idea of the *correlated equilibrium.*

Back to the Battle of the Sexes example, imagine that there are only two possible *states of nature*, x and y (for instance, *sunny* and *rainy*) and that the realized state can be fully observed by the two players. Suppose further that Player 1 announces that he is going to go to the Football if x occurs and to the Opera if y occurs. Then the best response for Player 2 is to mimic what Player 1 is doing. The combination of these behaviors is an *equilibrium* in the sense that, given that the other player is doing, no one wants to change his or her behavior. The equilibrium is achieved by means of some "coordination device" (a weather report, for example) that implements a random variable whose realizations are observable to the two players. This is the case of "perfect correlation" (or perfect coordination), when all the players have access to exactly the same information.

less than full correlation

Consider, on the contrary, the following case. There are three possible states of nature: x, y, and z (sunny, cloudy, and rainy), but players have access to the information in different ways (different weather reports). For instance, imagine that Player 1 can tell whether the state of nature is x or an element of $[y, z]$ (sunny or not), and that Player 2 can distinguish only between $[x, y]$ and z (rainy or not). Then, the strategy for Player 1 has two elements: what to do if he learns that the state of nature is x and what to do if the state of nature is in $[y, z]$. Similarly for Player 2: what to do when the state of nature is in $[x, y]$ and what to do if it is x.

A player's strategy is optimal, given the strategy of the other(s) player(s), if for any realization of his information he can do no better by choosing an action different from that dictated by his strategy

Imagine, for example, that the probabilities of the states x, y, and z are p_x, p_y, and $p_z$respectively, and that Player 2 announces that she is going to the Football if it is not going to rain (that is, $[x, y]$ occurs) and to the Opera if it is going to rain (x occurs). Then, if Player 1 is informed that it is not going to be sunny (that is, $[y, z]$ occurs)[1], he will have to choose the action that is optimal against his partner going to the Football with probability

[1] What will be Player 1's optimal choice if he is informed that x occurs ?

$p^2(F|[y,z])$ and to the Opera with probability $p^2(O|[y,z])$, where:

$$\begin{aligned} p^2(F|[y,z]) &= p([x,y]|[y,z]) = \frac{p_y}{p_y + p_z} \\ p^2(O|[y,z]) &= p(z|[y,z]) = \frac{p_z}{p_y + p_z} \end{aligned}$$

This is one example in which full coordination is not possible (players observe different signals) yet there is some correlation between their actions.

Formally, we can give the following definition of a *correlated equilibrium*

Definition of Correlated Equilibrium

Definition. *Correlated Equilibrium*
A correlated equilibrium of a a game $\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$ *consists of*

(i) *A finite probability space* (Ω, p) *(*Ω *is the set of* states of nature *and* p *is a probability measure on* Ω*)*

(ii) *For each player* $i \in \mathcal{I}$ *a* partition $\mathcal{P}^i$ *of* Ω *(*$\mathcal{P}^i$ *represents Player* i*'s information structure)*

(iii) *For each player* $i \in \mathcal{I}$ *a function* $\bar{\sigma}^i : \Omega \to S^i$ *such that* $\bar{\sigma}^i(\omega) = \bar{\sigma}^i(\omega')$ *whenever* $\omega, \omega' \in P^i_k$ *for some* $P^i_k \in \mathcal{P}^i$ *(*$\bar{\sigma}^i$ *is the* strategy *of Player* i*)*

such that for every player $i \in \mathcal{I}$ *and for any other* $\sigma^i : \Omega \to S^i$ *such that* $\sigma^i(\omega) = \sigma^i(\omega')$ *whenever* $\omega, \omega' \in P^i_k$ *for some* $P^i_k \in \mathcal{P}^i$ *(that is, for any strategy of Player* i*) we have*

$$\sum_{\omega\in\Omega} p(\omega) u^i(\bar{\sigma}^i(\omega), \bar{\sigma}^{-i}(\omega)) \geq \sum_{\omega\in\Omega} p(\omega) u^i(\sigma^i(\omega), \bar{\sigma}^{-i}(\omega)) \quad (1)$$

✔ Notice that the *probability space* and the *information partitions* are part of the equilibrium

Proposition. *Any Nash equiulibrium "is" a Correlated equilibrium*
For every Nash equilibrium σ^* of the game $\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$ there is a correlated equilibrium $\langle(\Omega, p), \{\mathcal{P}^i\}_{i\in\mathcal{I}}, \bar{\sigma}\rangle$ in which for each player $i \in \mathcal{I}$ the distribution on S^i induced by $\bar{\sigma}^i$ is σ^{*i}

Proof. Let $\Omega = S = S^1 \times S^2 \times \cdots \times S^I$ and define p by $p(s^1, s^2, \ldots, s^I) = \sigma^{*1}(s^1) \times \sigma^{*2}(s^2) \times \cdots \times \sigma^{*I}(s^I)$. For each player $i \in \mathcal{I}$ and for all strategy $s^i \in S^i$ define $P^i(s^i) = \{s \in S \,|\, s_i = s^i\}$ and let $\mathcal{P}^i = \{P^i(s^i_1), P^i(s^i_2), \ldots, P^i(s^i_{|S^i|})\}$. Finally, define $\bar{\sigma}^i(s^1, s^2, \ldots, s^I) = s^i$.

Then, $\langle(\Omega, p), \{\mathcal{P}_i\}_{i\in\mathcal{I}}, \bar{\sigma}\rangle$ is a *correlated equilibrium* since for each player $i \in \mathcal{I}$ (1) is satisfied for every other strategy σ^i:

- The left hand side is the payoff of player i in the mixed strategy Nash equilibrium σ^*
- The right hand side is the payoff of player i when using the mixed strategy that chooses the pure strategy $\sigma^i(s^1, s^2, \ldots, s^I)$ with probability $\sigma^{*i}(s^i)$ and everybody else is playing according to the mixed strategy σ^{*j}

Finally, the distribution on S^i induced by $\bar{\sigma}^i$ is σ^{*i} $\square$

Example: The Battle of the Sexes

Remember, again, the *mixed* strategy Nash equilibrium in the Battle of the Sexes

		$\frac{1}{3}$	$\frac{2}{3}$
	$1\backslash 2$	F	O
$\frac{2}{3}$	F	$2, 1$	$0, 0$
$\frac{1}{3}$	O	$0, 0$	$1, 2$

A correlated equilibrium can be constructed as follows

Example: The Battle of the Sexes

- $\Omega = \{(F,F),(F,O),(O,F),(O,O)\}$

- $p(F,F) = \frac{2}{3} \times \frac{1}{3} = \frac{2}{9}$

  $p(F,O) = \frac{2}{3} \times \frac{2}{3} = \frac{4}{9}$

  $p(O,F) = \frac{1}{3} \times \frac{1}{3} = \frac{1}{9}$

  $p(O,O) = \frac{1}{3} \times \frac{2}{3} = \frac{2}{9}$

- $\mathcal{P}^1 = \{\{(F,F),(F,O)\},\{(O,F),(O,O)\}\}$
  $\mathcal{P}^2 = \{\{(F,F),(O,F)\},\{(F,O),(O,O)\}\}$

- $$\begin{array}{ll} \bar{\sigma}^1(F,F) = F & \bar{\sigma}^2(F,F) = F \\ \bar{\sigma}^1(F,O) = F & \bar{\sigma}^2(F,O) = O \\ \bar{\sigma}^1(O,F) = O & \bar{\sigma}^2(O,F) = F \\ \bar{\sigma}^1(O,O) = O & \bar{\sigma}^2(O,O) = O \end{array}$$

Example: The Battle of the Sexes

In this case, the equilibrium condition

$$\sum_{\omega \in \Omega} p(\omega) u^i(\bar{\sigma}^i(\omega), \bar{\sigma}^{-i}(\omega)) \geq \sum_{\omega \in \Omega} p(\omega) u^i(\sigma^i(\omega), \bar{\sigma}^{-i}(\omega))$$

for player 1 boils down to

$$\frac{2}{9}2 + \frac{4}{9}0 + \frac{1}{9}0 + \frac{2}{9}1 \geq \frac{2}{9}0 + \frac{4}{9}1 + \frac{1}{9}2 + \frac{2}{9}0$$

Example: The Battle of the Sexes

It is also easy to check that there is another correlated equilibrium in the Battle of the Sexes game that yields a payoff of $(\frac{3}{2}, \frac{3}{2})$.

Indeed, consider

- $\Omega = \{x, y\}$
- $p(x) = \frac{1}{2}$

  $p(y) = \frac{1}{2}$
- $\mathcal{P}^1 = \{\{x\}, \{y\}\}$
  $\mathcal{P}^2 = \{\{x\}, \{y\}\}$
- $\begin{array}{ll} \bar{\sigma}^1(x) = F & \bar{\sigma}^2(x) = F \\ \bar{\sigma}^1(y) = O & \bar{\sigma}^2(y) = O \end{array}$

✔ Notice that $p(x)$ and $p(y)$ are not important for the equilibrium strategies, only for the final payoff

Proposition. *NOT every Correlated equilibrium is a Nash Equilibrium*
Any convex combination of correlated equilibrium payoffs profiles of a game $\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$ *is a correlated equilibrium payoff profile of* Γ

Proof. Let $(u^1, u^2, \ldots, u^K)$ a collection of K profiles of payoffs corresponding to K correlated equilibria

For each value of $k \in [1, K]$, let $\left\langle (\Omega^k, p^k), \{\mathcal{P}_i^k\}_{i\in\mathcal{I}}, \bar{\sigma}^k \right\rangle$ be the correlated equilibrium that generates the payoff profile u^k.

The following defines a correlated equilibrium whose payoff is $\sum_k \lambda^k u^k$:

- $\Omega = \Omega^1 \cup \ldots \cup \Omega^K$
- For each $\omega \in \Omega$, $p(\omega) = \lambda^k p^k(\omega)$, where $\lambda^k \geq 0$, $\sum \lambda^k = 1$ and k is such that $\omega \in \Omega^k$ (assume that all Ω^k are disjoint)
- For each $i \in \mathcal{I}$, $\mathcal{P}^i = \cup_k \mathcal{P}^{ik}$ and define σ^i by $\sigma^i(\omega) = \sigma^{ik}(\omega)$, where k is such that $\omega \in \Omega^k$

□

It is like if first we randomly select the correlated equilibrium to play, and second we implement it

In the case of the Battle of the Sexes we have

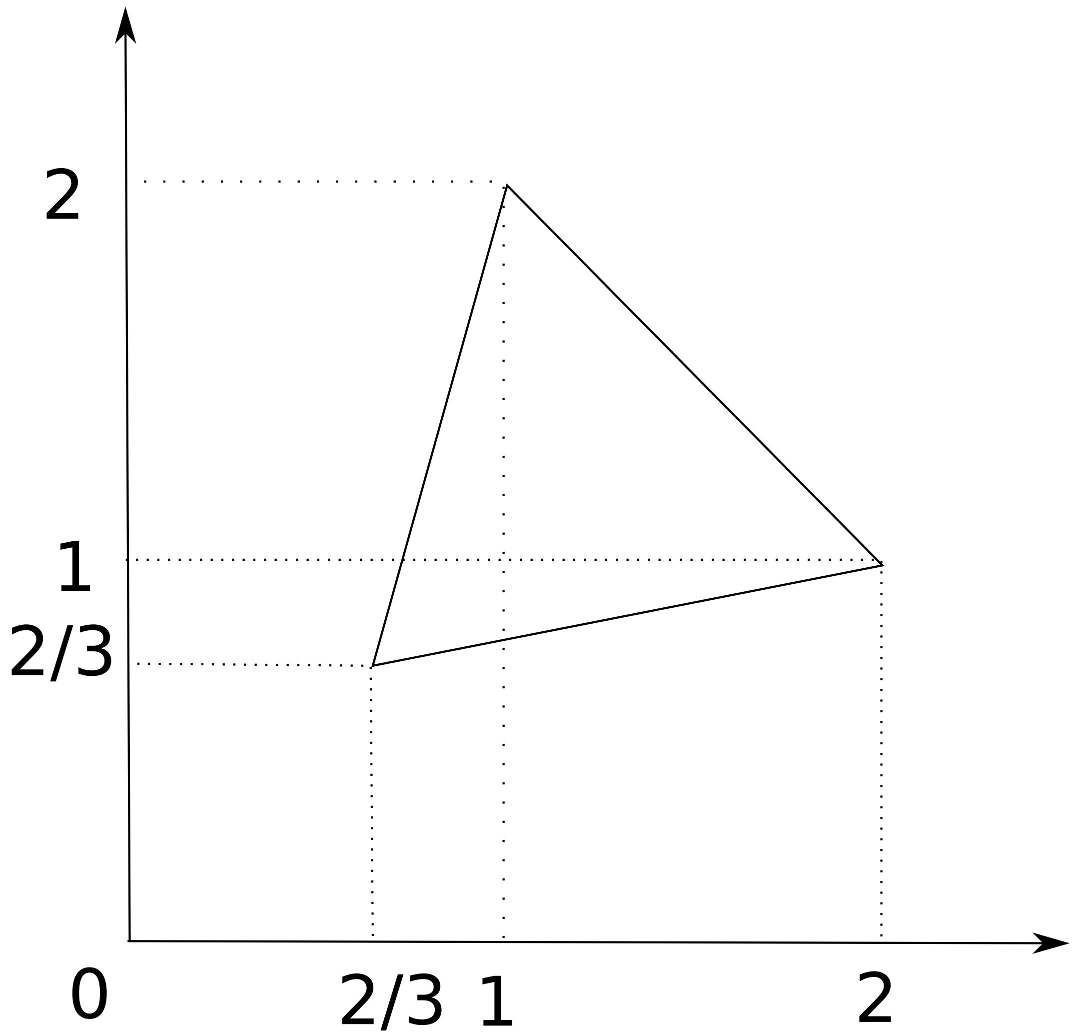

Example: The Game of Chicken (Wikipedia)

Consider the game of chicken

$1\backslash 2$	D	C
D	$0,0$	$7,2$
C	$2,7$	$6,6$

In this game two individuals are challenging each other to a contest where each can either dare or chicken out. If one is going to Dare, it is better for the other to chicken out. But if one is going to chicken out it is better for the other to Dare. This leads to an interesting situation where each wants to dare, but only if the other might chicken out.

In this game, there are three Nash equilibria. The two pure strategy Nash equilibria are (D, C) and (C, D). There is also a mixed strategy equilibrium where each player Dares with probability $\frac{1}{3}$.

Now consider a third party (or some natural event) that draws one of three cards labeled: (C, C), (D, C), and (C, D). After drawing the card the third party informs the players of the strategy assigned to them on the card (but not the strategy assigned to their opponent). Suppose a player is assigned D, he would not want to deviate supposing the other player played their assigned strategy since he will get 7 (the highest payoff possible). Suppose a player is assigned C. Then the other player will play C with probability $\frac{1}{2}$ and D with probability $\frac{1}{2}$. The expected utility of Daring is $0\frac{1}{2} + 7\frac{1}{2} = 3.5$ and the expected utility of chickening out is $2\frac{1}{2} + 6\frac{1}{2} = 4$. So, the player would prefer to Chicken out.

Since neither player has an incentive to deviate, this is a correlated equilibrium. Interestingly, the expected payoff for this equilibrium is $7\frac{1}{3} + 2\frac{1}{3} + 6\frac{1}{3} = 5$ which is higher than the expected payoff of the mixed strategy Nash equilibrium.

Trembling-Hand Equilibrium

- Rationality does not rule out weakly dominated strategies
- In fact, NE can include weakly dominated strategies
- Example: (B, R) in

$1\backslash 2$	L	R
T	$1,1$	$0,-3$
B	$3,0$	$0,0$

- But should we expect players to play weakly dominated strategies?
- Players should be completely sure of the choice of the others
- But, what if there is some risk that another player makes a "mistake"?

Trembling-Hand Perfection

Trembling hand perfect equilibrium is a refinement of Nash Equilibrium due to Reinhard Selten (1975). A trembling hand perfect equilibrium is an equilibrium that takes the possibility of off-the-equilibrium play into account by assuming that the players, through a "slip of the hand" or tremble, may choose unintended strategies, albeit with negligible probability.

We first define what is a "Perturbed Game"

Definition. *Perturbed Game*
For any game $\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$, *we can create a Perturbed Game* $\Gamma_\varepsilon = (\mathcal{I}, \{\Delta_\varepsilon(S^i)\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$ *by endowing each player* $i \in \mathcal{I}$ *with numbers* $\varepsilon^i(s^i) \in (0,1)$ *for any* $s^i \in S^i$ *in such a way that* $\sum_{s^i\in S^i} \varepsilon^i(s^i) < 1$ *and*

$$\Delta_\varepsilon(S^i) = \{\sigma^i \in \Delta(S^i) \,|\, , \sigma^i(s^i) \geq \varepsilon^i(s^i) \text{ for all } s^i \in S^i, \text{ and } \sum_{s^i\in S^i} \sigma^i(s^i) = 1\}$$

- Interpretation:

  Each strategy s^i is played with some minimal probability: this is the unavoidable probability of a mistake

Trembling-Hand Perfect Equilibrium

Definition. *Trembling-Hand Perfect Equilibrium (Selten 1975)*
A Nash equilibrium σ^* *of the game* $\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$ *is a Trembling-Hand Perfect Equilibrium if there is a sequence of perturbed games* $\{\Gamma_{\varepsilon^k}\}_{k=1}^{\infty}$ *that converges to* Γ *for wich there is some associated sequence of mixed strategy Nash equilibria* $\left\{\sigma^k\right\}_{k=1}^{\infty}$ *that converges to* σ^*

Proposition. *Computation of Trembling-Hand Perfect Equilibrium*
A Nash equilibrium σ^* *of the game* $\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$ *is a Trembling-Hand Perfect Equilibrium if and only if there is a sequence of totally mixed strategies* $\left\{\sigma^k\right\}_{k=1}^{\infty}$ *converging to* σ^* *such that, for any player* $i \in \mathcal{I}$, σ^{i*} *is a best response to every element of the sequence* $\left\{\sigma_{-i}^k\right\}_{k=1}^{\infty}$

- NOTICE:

  This rules out (B, R) in the previous example

Corollary.
σ^i in a trembling-hand perfect equilibrium cannot be weakly dominated. No weakly dominated pure strategy can be played with positive probability

Proposition. *Existence of Trembling-Hand Perfect Equilibrium*
Every finite game $\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$ has a Trembling-Hand Perfect Equilibrium

Corollary.
Every finite game $\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$ has at least a Nash equilibrium in which no player plays a weakly dominated strategy with possitive probability

Proper Equilibrium

Conside the game

$1\backslash 2$	L	R
T	$1,1$	$0,0$
B	$0,0$	$0,0$

It has 2 Nash Equilibria: (T,L) and (B,R) but only one Trembling-hand perfect equilibrium (T,L)

Let us add two *weakly dominated strategies*

$1\backslash 2$	L	R	R^+
T	$1,1$	$0,0$	$-9,-9$
B	$0,0$	$0,0$	$-7,-7$
B^+	$-9,-9$	$-7,-7$	$-7,-7$

Now it has 3 Nash Equilibria: (T,L), (B,R), and (B^+,R^+) and 2 Trembling-hand perfect equilibrium (T,L) and (B,R)!!!

(B, R) is indeed a Trembling-Hand perfect equilibrium. Consider the totally mixed strategy $(\varepsilon, 1-2\varepsilon, \varepsilon)$ for both players. Deviating to T (or L) would yield $(\varepsilon - 9\varepsilon)$, while B has a payoff of -7ε. Thus, playing B $((0,1,0))$ is a best response to $(\varepsilon, 1-2\varepsilon, \varepsilon)$ for any ε and, clearly, $(\varepsilon, 1-2\varepsilon, \varepsilon) \rightarrow_{\epsilon \rightarrow 0} (0,1,0)$

- Idea of *proper equilibrium*: more likely to tremble to better strategies:
  - second-best actions assigned at most ε times the probability of third-best actions,
  - fourth-best actions assigned at most ε times the probability of third-best actions,
  - etc.

- (B, R) is not a *proper equilibrium*: if player 2 puts weight ε on L and ε^2 on R then deviating to T for player 1 will produce a payoff of $(\varepsilon - 9\varepsilon^2)$, while B has a payoff of -7ε. In this case, $(\varepsilon - 9\varepsilon^2) > -7\varepsilon$ for ε small enough

Definition. *ε-proper strategies (Myerson 1978)*
A totally mixed strategy profile σ of the game $\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$ is said to be an ε-proper if there is a parameter $\varepsilon > 0$ such that for any player $i \in \mathcal{I}$ and for any pair of pure strategies $s_i, s_i' \in S^i$ for which

$$u^i(s_i, \sigma^{-i}) < u^i(s_i', \sigma^{-i})$$

we have that

$$\sigma^i(s_i) \leq \varepsilon\sigma^i(s_i')$$

Definition. *Proper Equilibrium (Myerson 1978)*
A Nash equilibrium σ^* of the game $\Gamma = (\mathcal{I}, \{S^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$ is said to be a Proper Equilibrium if it is the limit point of a sequence of ε-proper totally mixed strategy profiles

2.3 Dynamic Games of Complete Information

Players take decisions one after another. Thus, some players may know the actions taken by other players before choosing its own

Basic Assumptions

1. Players are rational and self-interested (Rationality)

2. Players have full information about all the elements of the game (Complete Information)

3. Players know that players are rational and fully informed, and they know that others know that players are rational and fully informed, and they know that others know that they know that players are rational and fully informed, and ... (Common Knowledge)

Extensive Form Game Representation

Dynamic games, because of its sequential structure, are better represented using trees. Formally, it is called *Extensive Form Representation* and consists of:

(i) A finite set of nodes N , a finite set of possible actions A, and a finite set of players $\mathcal{I} = \{1, ..., I\}$.

(ii) A function $p : N \rightarrow N \cup \emptyset$ specifying a single immediate predecessor of each node x;

- $p(x) \neq \emptyset$ for all $x \in N$ but one, 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.
- The set of terminal nodes is $T = \{x \in N | s(x) = \emptyset\}$. All other nodes in$N \backslash T$ are decision nodes.

(iii) A function $\alpha : N\backslash\{x_0\}\rightarrow A$ 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 $\alpha(x') \neq \alpha(x'')$.

- The set of choices/actions available at decision node x is $c(x) = \{a \in A \,|\, a = \alpha(x') \text{ for some } x' \in s(x)\}$.

(iv) A collection of information sets $\mathcal{H}$, and a function $H : N \rightarrow \mathcal{H}$ assigning each decision node x to an information set $H(x) \in \mathcal{H}$.

- Thus, the information sets in $\mathcal{H}$ form a partition of N .
- We require that all decision nodes assigned to a single information set have the same choices/actions available; formally, $c(x) = c(x')$ if $H(x) = H(x')$.

We can therefore write the choices/actions available at information set H as $C(H) = \{a \in A \,|\, a \in c(x) \text{ for all } x \in H\}$.

(v) A function $\iota : H \rightarrow \mathcal{I}$ assigning each information set in $\mathcal{H}$ to a player (or to nature: formally player 0) who moves at the decision nodes in the set.

- We can denote the collection of player i's information sets by $\mathcal{H}_i = \{H \in \mathcal{H} | i = \iota(H)\}$.

(vi) A function $\rho : \mathcal{H}_0 \times A \rightarrow [0,1]$ assigning probabilities to actions at information sets where nature moves and satisfying $\rho(H,a) = 0$ if $a \notin C(H)$, $\rho(H,a) \geq 0$ if $a \in C(H)$, and $\sum_{x \in C(H)} \rho(H,a) = 1$ for all $H \in \mathcal{H}_0$.

(vii) A collection of payoff functions $u = \{u^1(\cdot), ..., u^I(\cdot)\}$ assigning utilities to the players for each terminal node that can be reached,$u^i : T \rightarrow R$ (vNM utility functions!).

Thus, formally, $\Gamma_E = (N, A, I, p(\cdot), \alpha(\cdot), \mathcal{H}, H(\cdot), \iota(\cdot), \rho(\cdot), u)$ completely describes an extensive form game.

As long as the set N is finite, the extensive form game is *finite*.

Example: The Truth Game

There are two players, called 1 and 2, and a game-master. The game-master has a coin that is bent in such a way that, flipped randomly, the coin will come up "heads" 80% of the time. (The bias of the coin is known to both players.)

The game-master flips this coin, and the outcome of the coin flip is shown to player 1. Player 1 then makes an announcement to player 2 about the result of the coin flip; player 1 is allowed to say either "heads" or "tails" (and nothing else). Player 2, having heard what player 1 says but not having seen the result of the coin flip, must guess what the result of the coin flip was, either "heads" or "tails". That ends the game.

Payoffs are made as follows. For player 2 things are quite simple; player 2 gets $1 if his guess matches the actual result of the coin flip, and he gets $0 otherwise. For player 1 things are more complex. She gets $2 if player 2's guess is that the coin came up "heads", and $0 if player 2 guesses "tails", regardless of how the coin came up. In addition to this, player 1 gets $1 (more) if what she (player 1) says to player 2 matches the result of the coin flip, while she gets $0 more if her message to player 2 is different from the result of the coin flip.

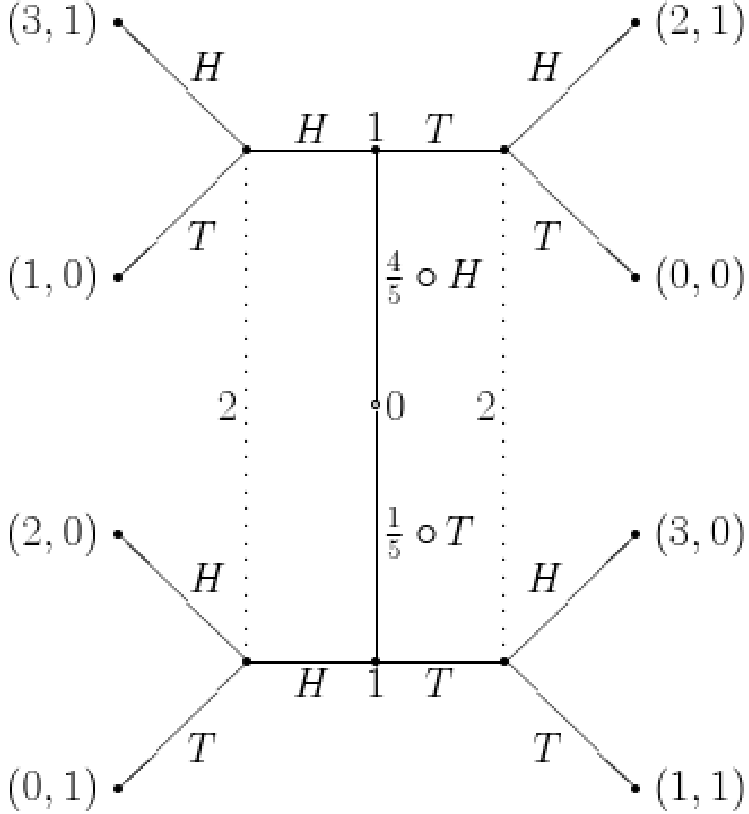

Figure 1: The Truth Game. Extensive Form Representation

Information

Notice, in the Truth Game, that Player 1 has full information on the move chosen by nature (the coin flip), whereas Player 2 doesn't. In other words, when its Player 1 turn to move, he has FULL information on the history of the game until that moment, while this is not true for Player 2.

This is shown in the tree by means of the dotted lines that joints the decision nodes of Player 2. Such "dotted lines" correspond to the information sets "owned" by Player 2

Definition. Perfect/Imperfect Information
An extensive form game Γ_E *is one of perfect information if each information set contains a single decision node. Otherwise it is a game of imperfect information.*

Thus, the Truth Game is a game of Imperfect Information as the information sets of Player 2 contain more than one decision node.

The difference between Perfect/Imperfect Information should not be mistaken by the difference between Complete/Incomplete Information

Definition. Perfect/Imperfect Recall
A game is of Perfect Recall if players never forget information acquired throughout the game, including the strategies he chose.

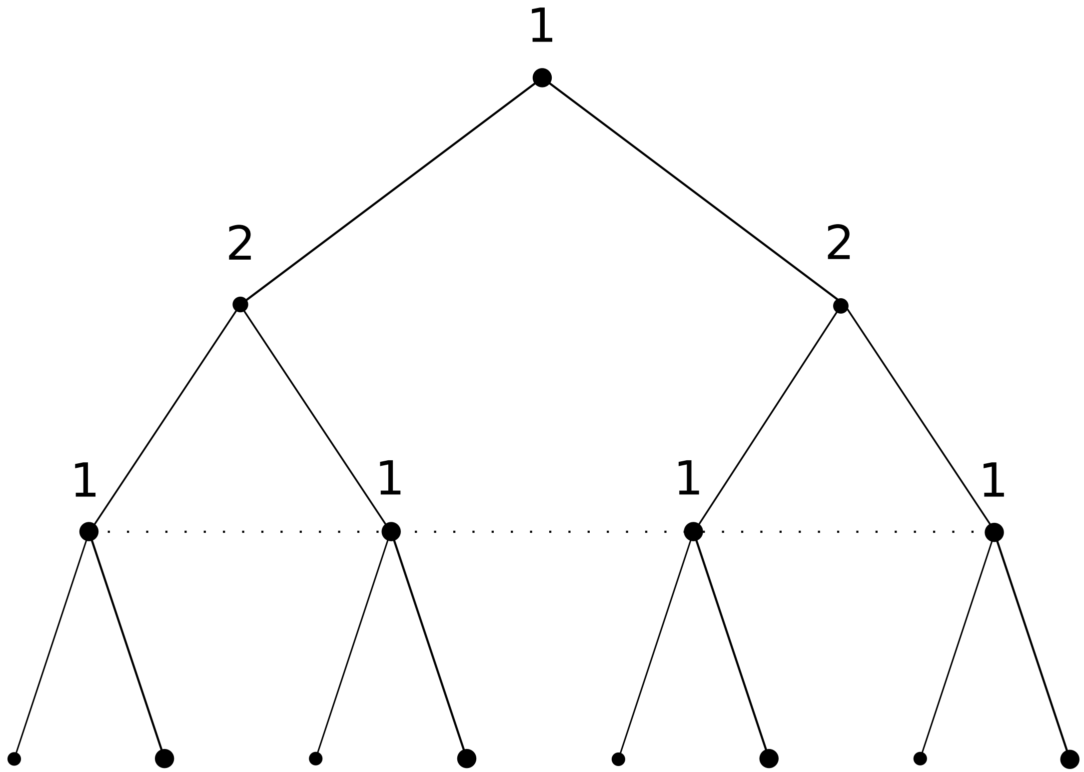

Figure 2: A game without Perfect Recall

We will consider games with Perfect Recall only

Strategies

Loosely speaking, a strategy for a player is a contingent plan, that is, a specification of what action to take at each of his/her information sets. Thus, a strategy in a dynamic game is a "plan" that says what to do at every point in the game the player has to choose an action.

This is (can be) very complicated (think Chess)

Definition. Pure and Mixed Strategies
A (pure) strategy *for player i is a function $s_i : \mathcal{H}_i \to A$ such that $s_i(H_i) \in C(H_i)$. Similarly as for normal form games a mixed strategy σ_i is a probability distribution over agent i's pure strategies.*

Example. The Truth Game
In The Truth Game players have the following (pure) strategies:

$$\begin{aligned} S_1 &= \{HH, HT, TH, TT\} \\ S_2 &= \{HH, HT, TH, TT\} \end{aligned}$$

where, for each player, each pair of actions indicates the action to choose at each of their information sets.

Note that any mixed strategy σ_i in an extensive form game will induce a probability distribution on the actions player i can choose in any of his information sets H_i. So, rather than randomizing over the potentially large set of pure strategies, player i could randomize separately over the possible actions in $C(H_i)$ at each of his information sets H_i.

Definition. Behavioral Strategies
A behavioral strategy *for player i specifies for every information set $H_i \in \mathcal{H}_i$ a probability distribution $b_i(\cdot, H_i)$ over $C(H_i)$, i.e., for each $H_i \in \mathcal{H}_i$ and $a \in C(H_i)$, $bi(a, H_i) \geq 0$, and $\sum_{a \in C(H_i)} b_i(a, H_i) = 1$.*

Formally, they are equivalent

Theorem. *In an extensive form game with perfect recall, mixed and behavioral strategies are equivalent (with respect to the probability distribution on the terminal nodes they induce).*

Example. The Truth Game
Notice, for instance, that the *mixed strategy* $(\frac{1}{2}, \frac{1}{2}, 0, 0)$ for Player 2 is equivalent to the *behavioral strategy*

$$
\begin{aligned}
b_2(H, H_{21}) &= 1, \quad b_2(T, H_{21}) = 0 \\
b_2(H, H_{22}) &= \frac{1}{2}, \quad b_2(T, H_{22}) = \frac{1}{2}
\end{aligned}
$$

where $H_{21}$is the information set owned by Player 2 following the choice of $Heads$ by Player 1 (Left on Figure 1) and $H_{22}$is the information set owned by Player 2 following the choice of $Tails$ by Player 1 (Right on Figure 1)

Normal Form Game Representation

As well as Static Games, Dynamic Games can be represented in *Normal (Strategic) Form.* The only drawback is that with this representation we loose all the information regarding the sequentiality and informational structure of the game.

Claim

- Any Dynamic Game can be represented in Normal (Strategic) Form in a unique way (up to re-labeling of the strategies) .
- Any Static Game can be represented in Extensive Form, but the representation might not be unique.

Example. The Truth Game

Recall that in The Truth Game players have the following (pure) strategies:

$$S_1 = \{HH, HT, TH, TT\}$$
$$S_2 = \{HH, HT, TH, TT\}$$

where, for each player, each pair of actions indicates the action to choose at each of their information sets.

The table below represents the same game in *Normal Form*

	HH	HT	TH	TT
HH	$\frac{14}{5}, \frac{4}{5}$	$\frac{14}{5}, \frac{4}{5}$	$\frac{4}{5}, \frac{1}{5}$	$\frac{4}{5}, \frac{1}{5}$
HT	$\frac{15}{5}, \frac{4}{5}$	$\frac{13}{5}, \frac{5}{5}$	$\frac{7}{5}, 0$	$\frac{5}{5}, \frac{1}{5}$
TH	$\frac{10}{5}, \frac{4}{5}$	$\frac{2}{5}, 0$	$\frac{8}{5}, \frac{1}{5}$	$0, \frac{1}{5}$
TT	$\frac{11}{5}, \frac{4}{5}$	$\frac{1}{5}, \frac{1}{5}$	$\frac{11}{5}, \frac{4}{5}$	$\frac{1}{5}, \frac{1}{5}$

Table 1: Normal Form Representation of the Truth Game

- The fact that any dynamic game can be represented in Normal Form guarantees that any dynamic game has a Nash equilibrium (at least in mixed strategies)

- The problem, though, is that typically dynamic games have "too many" Nash equilibria, some of which do not make sense.

It's easy to check that The Truth Game has infinitely many Nash Equilibria:

$$\begin{aligned} Player\,1 &: \quad HH \\ Player\,2 &: \quad anything \end{aligned}$$

Example. The Market Entry Game

A firm wants to enter a new market, and its main concern is about the reaction of an incumbent company that currently is making a profit \$100,000. If the reaction is aggressive, the challenger will suffer a loss of \$10,000, and the incumbent's profits will be reduced to only \$20,000 (because of the costs of the fight). On the other hand, if the incumbent chooses to accommodate to the new market scenario, then the two companies will share the \$100,000 profit (\$50,000 each) Obviously, the challenger can always choose to stay out if that seems to be the best choice (with a profit of \$0)

The tree in Figure 3 corresponds to the Extensive Form Representation of this game

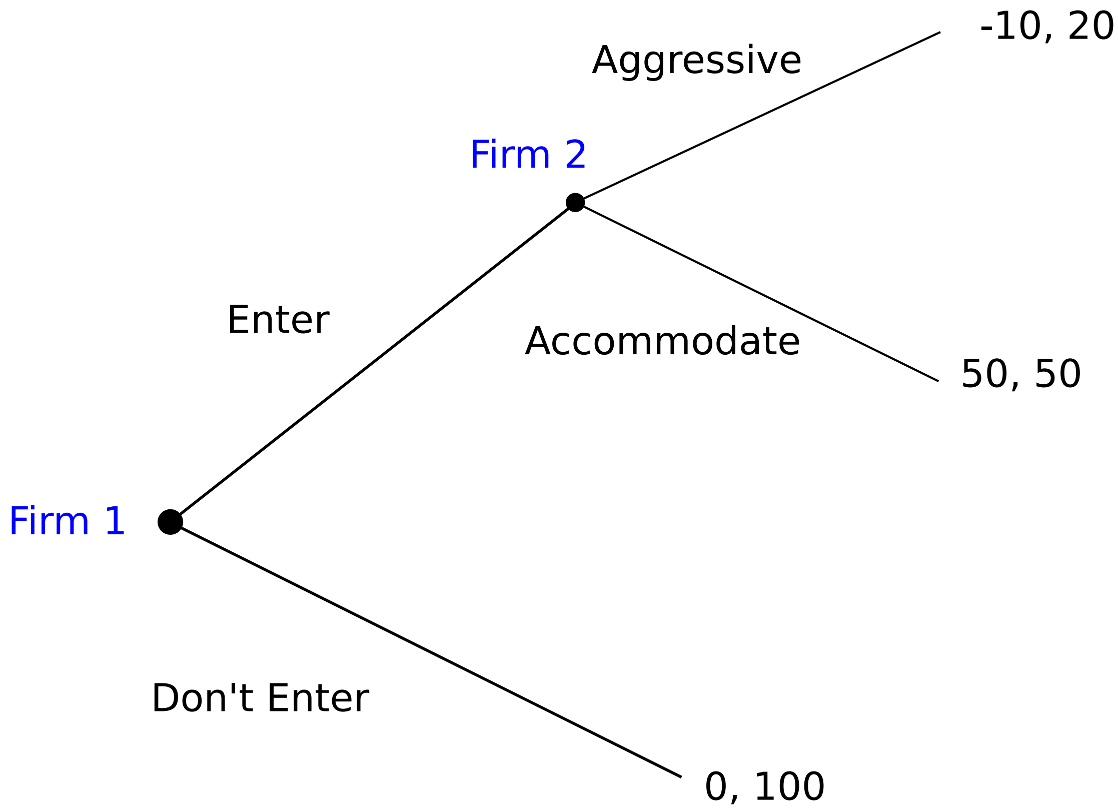

Figure 3: The Market Entry Game

In this case each player only has one *information set* and therefore *strategies* and *actions* coincide in this game. Thus, the *Normal Form Representation* is given in the table below

	Aggressive	Accommodate
Enter	-10,20	50,50
Don't	0,100	0,100

Table 2: Normal Form Representation of the Market Entry Game

Clearly, the game has 2 equilibria in pure strategies (as well as another one in mixed strategies)

$$EQ1 \quad : \quad (Enter,\ Accommodate)$$
$$EQ2 \quad : \quad (Don't,\ Aggressive)$$

- $EQ1$ is *sequentially rational* in the sense that if the players actually were to follow the strategies prescribed by this equilibrium, it would be rational for the two of them to do so in every step of the game

- $EQ2$ is **NOT** *sequentially rational* as if the players actually were to follow the strategies prescribed by this equilibrium, then it wouldn't be rational for Player 2 to play *Aggressive* once Player 1 has chosen *Enter* (Player 2 would get a strictly lower payoff that playing *Accommodate*) Being "Aggressive" in step 2 is a threat to Player 1, but it's an **incredible threat** since it would also harm Player 2 and therefore a "rational" Player 2 would never carry it out.

Definition. Sequential Rationality
A profile of strategies satisfies sequential rationality *if at any information set players following those strategies specify optimal actions at that point in the game.*

We will consider different *equilibrium refinements* to eliminate those Nash equilibria that are not *sequentially rational:*

- **Backward Induction:** When the game is of Perfect Information

- **Subgame Perfection:** When the game is of Imperfect Information

Equilibrium by Backward Induction

Definition. Equilibrium by Backward Induction
The (pure) strategy profile s is a backward induction strategy profile for the finite extensive form game Γ_E of perfect information if it is derived as follows.

1. *Call a node x penultimate in Γ_E if all nodes immediately following it are terminal nodes.*

2. *For every penultimate node x, let $s_{\iota(x)}(x)$ be an action leading to a terminal node that maximizes player $\iota(x)$'s payoff on $c(x)$.*

3. *Let u_x denote the resulting payoff vector.*

4. *Remove the nodes and actions strictly following every penultimate node x in Γ_E and assign the payoff u_x to x, which then becomes a terminal node.*

5. *Repeat the above described process until an action has been assigned to each decision node*

Example

Consider the game in the tree below. At the first “penultimate” node, Player 1 will choose R'

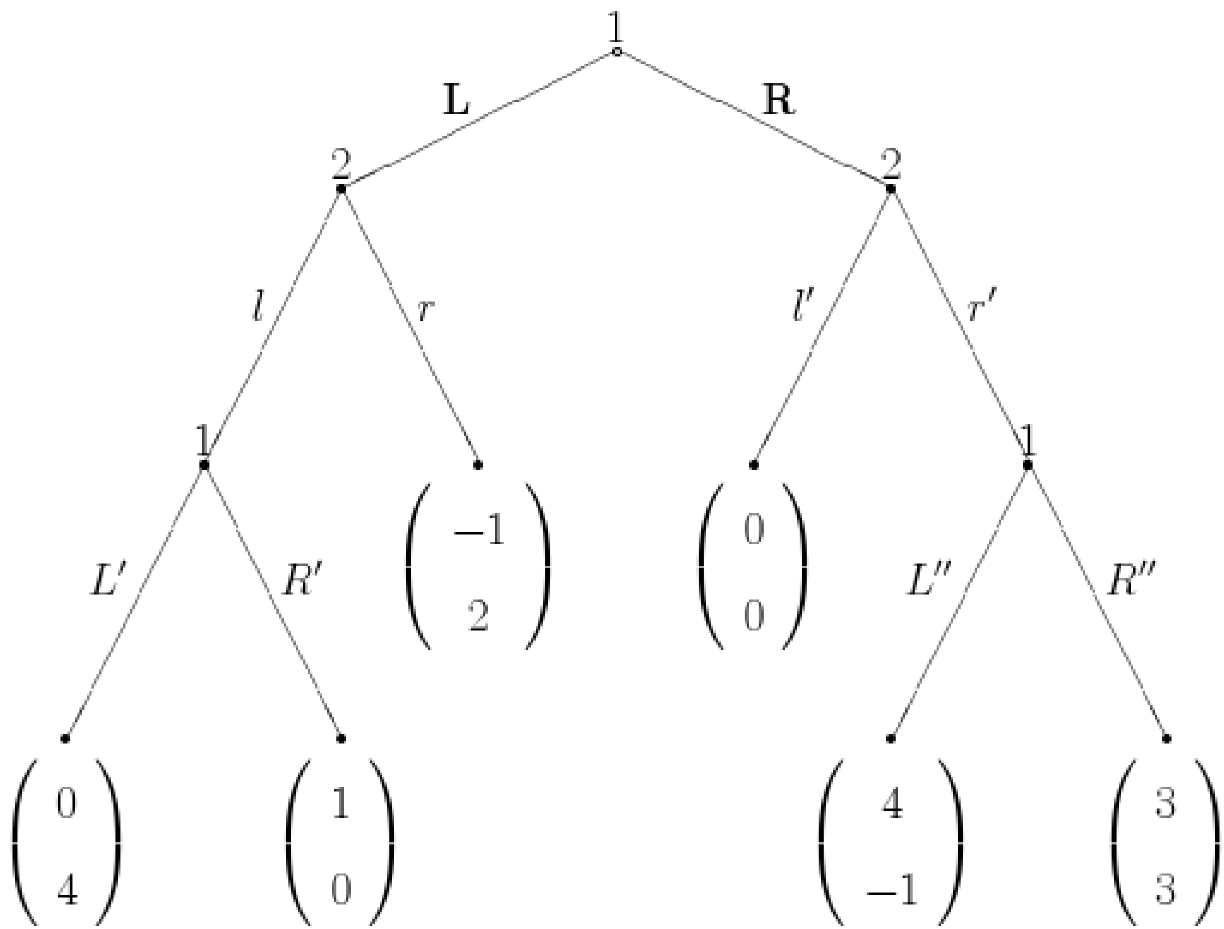

Example

The game reduces to the tree below. At the second penultimate node, Player 1 will chose L''

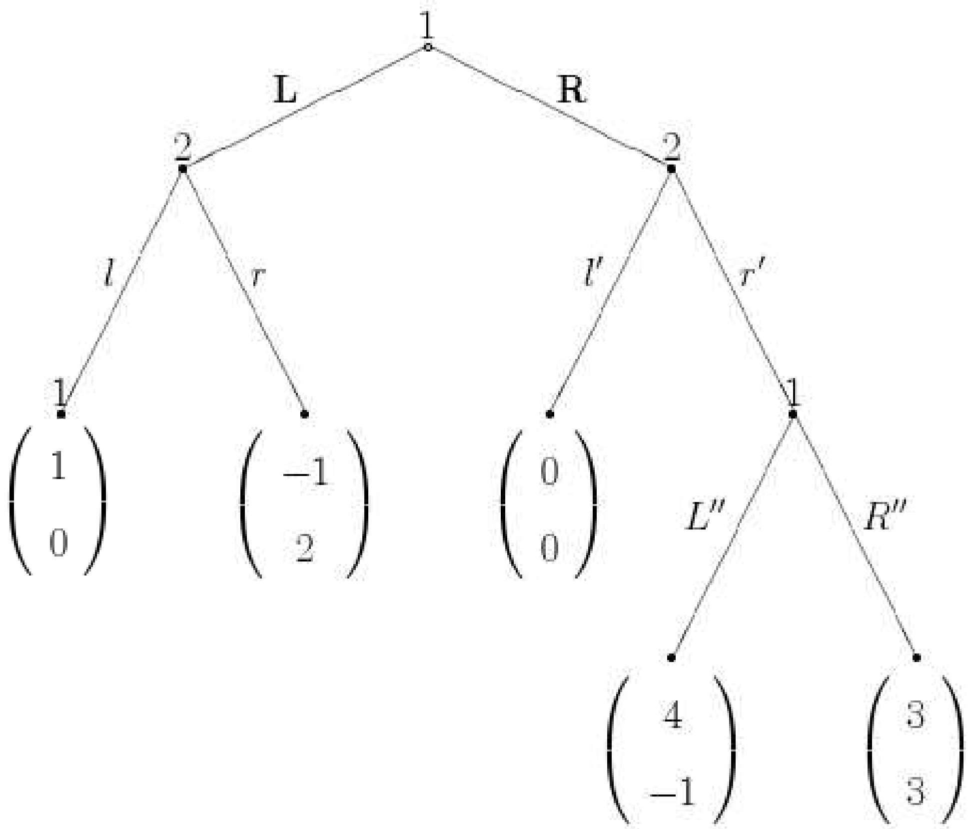

Example

The game reduces to the tree below. Now, at the new penultimate node, Player 2 will chose r

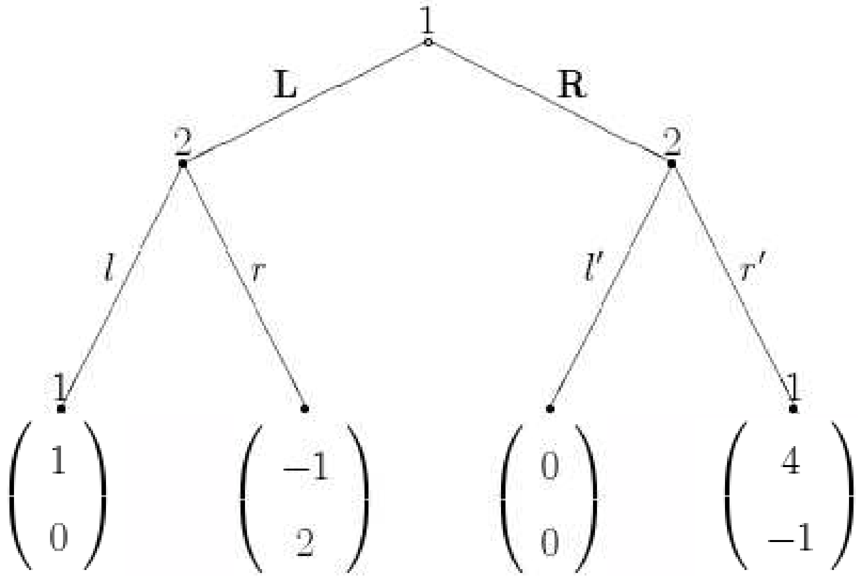

Example

The game reduces to the tree below. At the second new penultimate node, Player 2 will chose l'

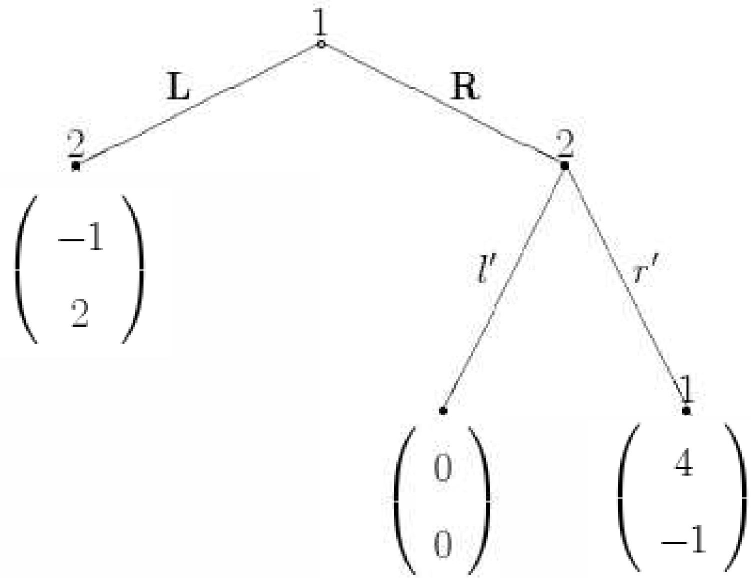

Example

The game reduces to the tree below.

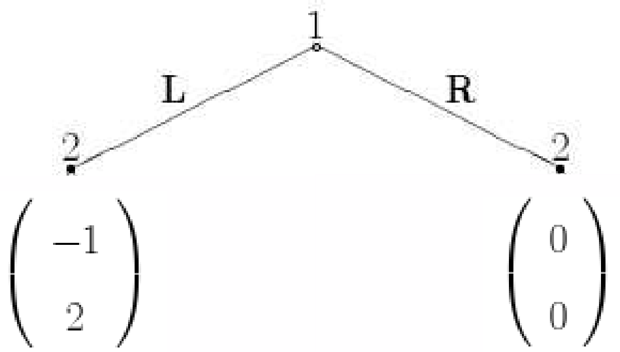

Example

Finally, Player 1 will chose R

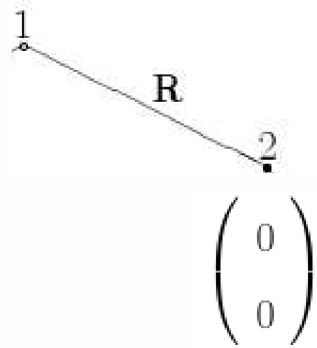

Thus, the Equlibrium by Backward Induction corresponds to:

$$\{RR'L'', rl'\}$$

Theorem. *Kuhn/Zermelo*
Every finite game of perfect information has a pure strategy Nash equilibrium that can be derived through backward induction.

Notice that only one of the two Nash Equilibria of the Market Entry Game (Figure 3) is obtained by Backward Induction, $EQ1 : (Enter,\ Accommodate)$

Subgame Perfect Nash Equilibrium

In games of Imperfect Information, like the Truth Game, Backward Induction is not possible.The concept of *Subgame* is needed to extend the principle of Backward Induction to this class of games

Definition. Subgame
A subgame of an extensive form game Γ_E 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 of this node, and contains only these nodes.*

(ii) *If decision node x is in the subgame, then every $x \in H(x)$ also is; e.g., there are no "broken" information sets.*

Definition. Subgame Perfect Nash Equilibrium
A strategy profile σ for Γ_E is a subgame perfect Nash equilibrium (SPNE) if it induces a Nash equilibrium in every subgame of Γ_E .

Example

Consider the game represented in the tree below

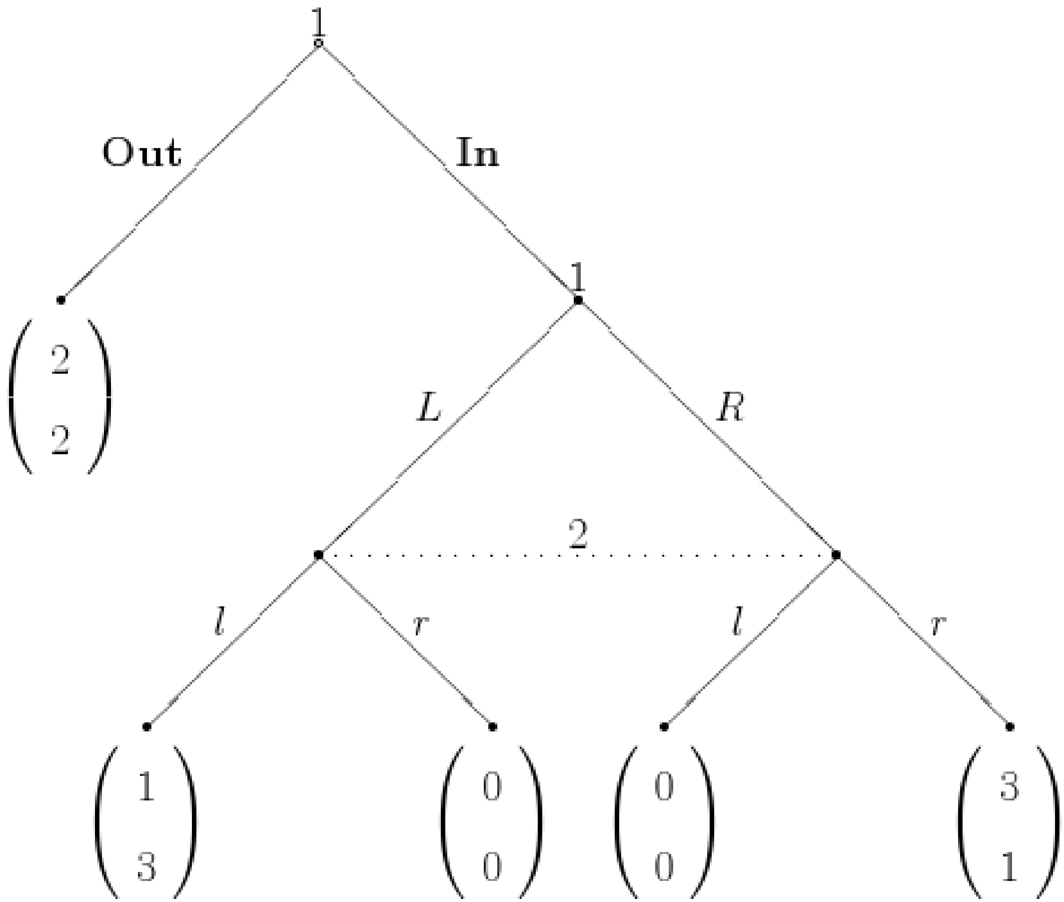

Example

There are two subgames in this game:\Gamm

1. The one that starts at the second information set of Player 1, which is a “proper” subgame
2. The game itself

Thus, to find the Subgame Perfect Nash Equilibrium we must

1. Find all the Nash Equilibria of the static game induced by the “proper” subgame
2. Proceed by Backward Induction to find all the Nash Equilibria of the game that results from the substitution of that “proper” subgame by each of its Nash Equilibria

Example

The Normal Form Representation of the proper subgame is given in the table 3 below

	l	r
L	1,3	0,0
R	0,0	3,1

Table 3: Normal From Representation of the proper subgame

This (sub)game has two Nash Equilibria in pure strategies: $\{(L,l),(R,r)\}$ and one in mixed strategies: $(\sigma^1,\sigma^2) = ((\frac{1}{4},\frac{3}{4}),(\frac{3}{4},\frac{1}{4}))$

Thus, moving by Backward Induction we find three possible games (one for each of the three Nash equilibria of the subgame)

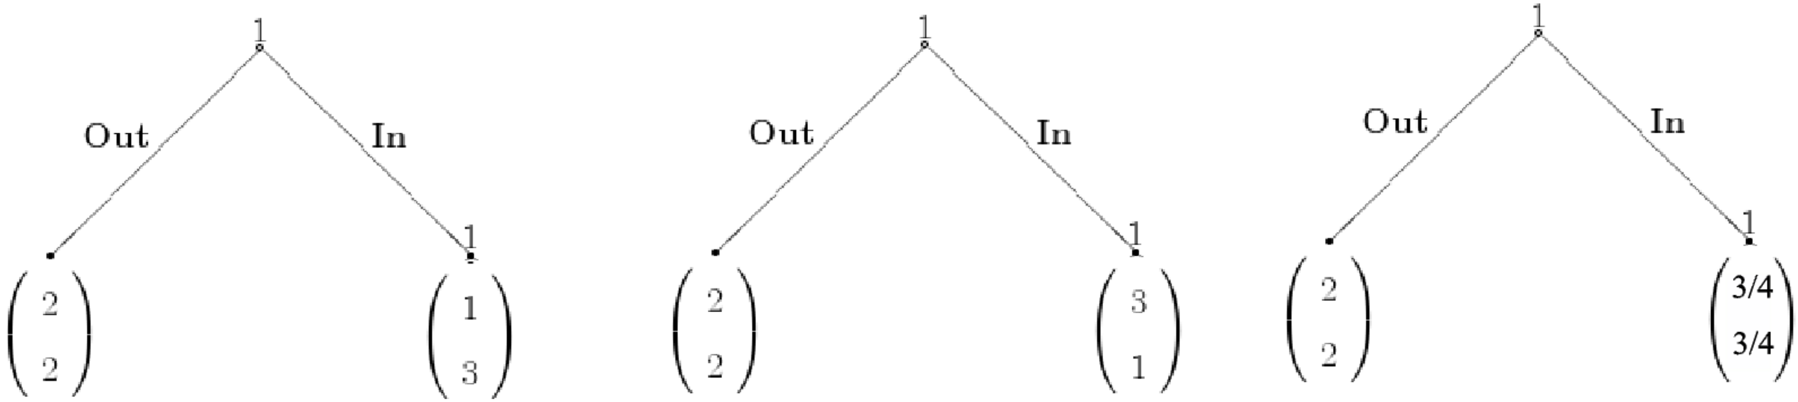

Hence, we have three Subgame Perfect Nash Equilibria of this game

$$
\begin{aligned}
EQ1 &: \{Out\, L, l\} \\
EQ1 &: \{In\, R, r\} \\
EQ3 &: \{Out\, (\frac{1}{4}, \frac{3}{4}), (\frac{3}{4}, \frac{1}{4})\}
\end{aligned}
$$

Question:
How many Subgame Perfect Nash Equilibria has The Truth Game ?

Repeated Games

- A special class of Dynamic Games is that of Repeated Games, when players play the same (static) *stage game* in every period
- When the actions of the players are observed after each period, players can condition their play on the past history of the game, which can lead to equilibrium outcomes that do not arise when the game is played only once
- In the Prisoners' Dilemma, for example, one might think that cooperation is possible based on some strategy of *punishment and reward* (stick and carrot)
- Repeated games can be a good representation of long-term relationships

Notation

- Let $\Gamma = (\mathcal{I}, \{A^i\}_{i\in\mathcal{I}}, \{u^i\}_{i\in\mathcal{I}})$ be the <u>*stage game*</u>. For convenience, strategies in the stage-game will be called *actions* $(A = A^1 \times \cdots \times A^I)$ in the repeated-game

- Let $h_t \in (A)^t$ be a <u>*history of the game*</u> up to period t. For completeness, $h_0 = \emptyset$ represents the history when the game starts at $t = 0$ (null history). Thus, $H_t = (A)^t \cap \emptyset$ represents the space of all possible period-t histories

- A pure strategy s^i for player i is a sequence of maps $\left\{s_t^i\right\}_{t=0}^{\infty}$ (one for each period t) that map possible histories at period t $(h_t \in H_t)$ to actions $a^i \in A^i$

$$s_t^i : H_t \longrightarrow A^i$$

- Consequently, a mixed (behavioral) strategy in the repeated game is a sequence of maps $\left\{\beta_t^i\right\}_{t=0}^{\infty}$ that map possible histories at period t $(h_t \in H_t)$ to mixed actions $\sigma^i \in \Delta(A^i)$

$$\beta_t^i : H_t \longrightarrow \Delta(A^i)$$

- Notice that a player's strategy can not depend on the past values of his opponents mixed actions σ_{-i}, it can depend only on the past values of a_{-i}.

- Notice that each period begins a proper subgame. Moreover, since the stage game is static, these are the only subgames of the game.

- Games can be repeated a finite (fixed) number of times T or an infinite (undecided or unknown) number of times

Infinitely Repeated Games

Correspond to the case when players think that the game can always extend one more period with positive probability

The payoff that the players (expect to) obtain at the end of the game is "some" weighted average of the per period payoffs

$$U^i = E_\sigma (1-\delta) \sum_{t=0}^{\infty} \delta^t u^i(\sigma_t(h_t))$$

where

- $\delta \in (0,1)$ is the *discount factor* that represents the fact that payoffs lose value as time passes (because they actually do or because players are impatient)

- $(1-\delta)$ is the *normalization factor* that serves to measure the stage-game and the repeated-game payoffs in the same units

Continuation Payoffs

Since each period begins a proper subgame, for any strategy profile β and history h_t we can compute the players' expected payoff from period t on. This is the *continuation payoff*, which will be re-normalized so that it is measured in period-t units

$$U_t^i = E_\sigma(1-\delta)\sum_{\tau=t}^{\infty}\delta^{\tau-t}u^i(\sigma_t(h_t))$$

Alternative Payoff specification

When players are "infinitely patient" $(\delta \to 1)$, we can use the limit of the average criterion

$$U^i = \lim\inf_{T\to\infty} E_\sigma(\frac{1}{T})\sum_{t=0}^{T} u^i(\sigma_t(h_t))$$

Example: The Repeated Prisoners' Dilemma

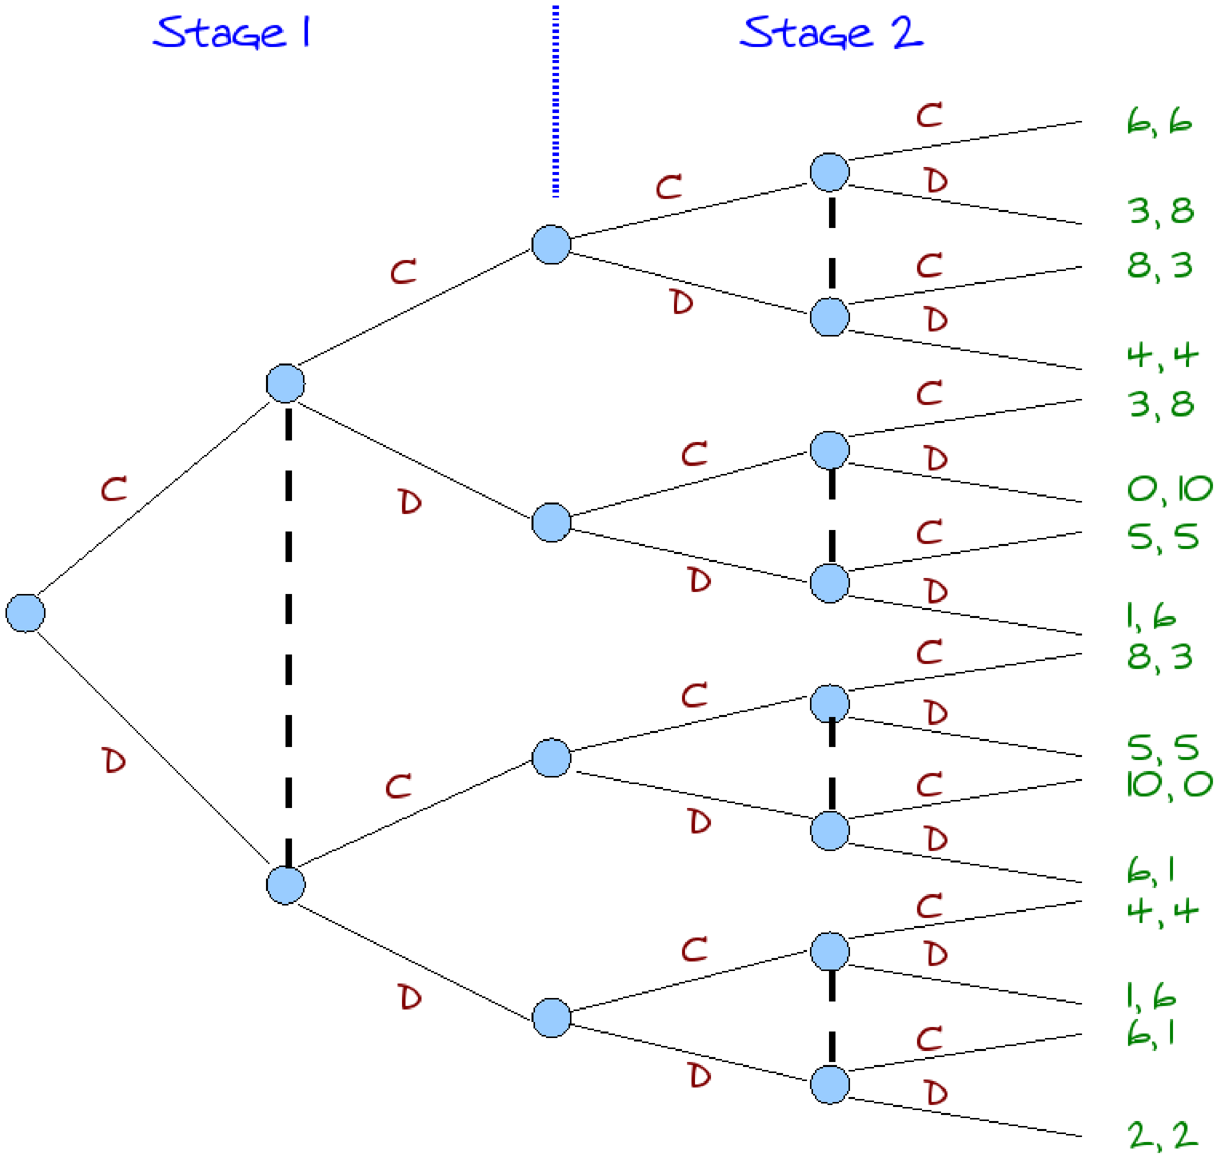

Basic Results (i)

The main results from the theory of infinitely repeated games are

Claim

(i) If σ^* is a Nash equilibrium of the stage game, then the strategies $\left\{\beta_t^{i*}\right\}_{t=0}^{\infty} = \left\{\sigma^{i*}\right\}_{t=0}^{\infty}$ for every player $i \in \mathcal{I}$ consitute a subgame-perfect Nash equilibrium of the repeated game

(ii) If $\left\{\sigma_j^*\right\}_{j=1}^{k}$ are k different Nash equilibria of the stage game, then for any *assginment map* $\rho : \{0, 1, 2, \ldots\} \to \{1, 2, \ldots, k\}$ (from *time periods* to *indices*) the strategies $\left\{\beta_t^{i*}\right\}_{t=0}^{\infty} = \left\{\sigma_{\rho(t)}^{i*}\right\}_{t=0}^{\infty}$ for every player $i \in \mathcal{I}$ consitute a subgame-perfect Nash equilibrium of the repeated game

To see that this Claim is true, note that with these strategies the future play of player i's opponents is independent of how he plays today. Thus, it is optimal for player i to maximize his current period's payoff, that is, to play a *static best response* to $\sigma_{\rho(t)}^{-i*}$

Thus, the repeated play of a game <u>does not</u> reduce the set of equilibrium payoffs

Basic Results (ii)

As a matter of fact, the only reason not to play *static best responses* is the concern about the future.

Hence, if the discount factor is small enough (the future does not matter much), only strategies that specify a *static equilibrium* at every history to which the equilibrium gives positive probability are subgame perfect Nash equilibria of the repeated game. On the other hand, for a large enough discount factor *almost* any payoff can be achieved by means of some subgame perfect Nash equilibrium (using strategies that punish out-of-equilibrium behaviors for long enouch periods)

These observations are formalized in the "Folk Theorem(s) for Infinitely Repeated Games"

Basic Results (iii)

Definition. *Reservation Utility (minmax payoff) and Individualy Rational Payoffs*

- *Define player i's* reservation utility *(or* minmax payoff*) to be*

$$\underline{v}^i = \min_{\sigma^{-i}} \left[\max_{\sigma^i} u^i(\sigma^i, \sigma^{-i}) \right]$$

That is, the lowest payoff player i's opponents can hold him to by any choice of σ^{-i}, provided that player i correctly foreses σ^{-i} and plays a best response to it (Worst Case Scenario).

- *A payoff v^i is* individualy rational *for player i if $v^i \geq \underline{v}^i$*

- *If $m^{-i,i}$ denotes the* minmax *profile agains player i, let $m^i \in A^i$ be a strategy for player i such that $u^i(m^i, m^{-i,i}) = \underline{v}^i$*

Example: The Repeated Prisoners' Dilemma

$1\backslash 2$	C	D
C	$3,3$	$0,5$
D	$5,0$	$1,1$

In this case we have:

$$\underline{v}^1 = \underline{v}^2 = 1 \text{ and } m^i = m^{-i} = D \; (i = 1, 2)$$

Basic Results (iv)

Definition. *Feasible payoff*
Under some technical assumptions (see Fudenberg-Tirole, pp. 151-152 for details), the set of feasible payoffs *vectors of the repeated game is*

$$V = \textit{convex hull}\left\{v \in \Re^I \mid \exists a \in A \,\textit{with}\, u^i(a) = v^i\, \forall i \in \mathcal{I}\right\}$$

Example: The Repeated Prisoners' Dilemma

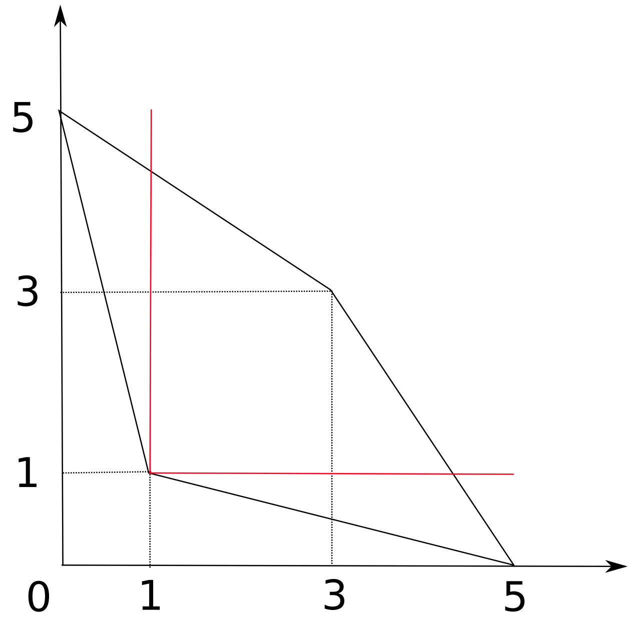

Any payoff combination inside the figure is feasible. Only payoff vectors above $(\underline{v}^1, \underline{v}^2) \geq (1,1)$ are individually rational

Basic Results (v)

Theorem. The Folk Theorem
For every possible payoff vector v with $v^i \geq \underline{v}^i$ for all player $i \in \mathcal{I}$, there exists a $\underline{\delta} < 1$ such that for all $\delta \in (\underline{\delta}, 1)$ there is a Nash equilibrium of the repeated game with payoffs v

Proof. Assume that there is a pure action profile $a \in A$ such that $u^i(a) = v^i \, \forall i$ and consider the following strategies for each player i

"Play a^i in period $t = 0$ and continue playing a^i as long as:

(i) the realized action in the previous period was a

(ii) the realized action in the previous period differed from a in two or more components

Otherwise, if player j was the only one not to follow profile a, then play $m^{i,j}$ (the *minmax* strategy against j) for the rest of the game"

Can player i gain by deviating from this strategy ?. If he does:

- He will get at most $\max_{b \in A^i} u^i(b, a^{-i})$ in the period he deviates

- He will get at most $\underline{v}^i$ in all periods after his first deviation since his opponents will *minmax* him afterwards

Thus, if player i deviates in period t he obtains at most

$$(1-\delta^t)v^i + \delta^t(1-\delta)\max_{b\in A^i} u^i(b,a^{-i}) + \delta^{t+1}\underline{v}^i$$

which will be less than v^i if

$$\begin{aligned}
(1-\delta^t)v^i + \delta^t(1-\delta)\max_{b\in A^i} u^i(b,a^{-i}) + \delta^{t+1}\underline{v}^i &\leq v^i \\
-\delta^t v^i + \delta^t(1-\delta)\max_{b\in A^i} u^i(b,a^{-i}) + \delta^{t+1}\underline{v}^i &\leq 0 \\
-v^i + (1-\delta)\max_{b\in A^i} u^i(b,a^{-i}) + \delta\underline{v}^i &\leq 0 \\
(\max_{b\in A^i} u^i(b,a^{-i}) - v^i) &\leq \delta(\max_{b\in A^i} u^i(b,a^{-i}) - \underline{v}^i) \\
\frac{(\max_{b\in A^i} u^i(b,a^{-i}) - v^i)}{(\max_{b\in A^i} u^i(b,a^{-i}) - \underline{v}^i)} &\leq \delta
\end{aligned}$$

Call

$$\underline{\delta}^i = \frac{(\max_{b \in A^i} u^i(b, a^{-i}) - v^i)}{(\max_{b \in A^i} u^i(b, a^{-i}) - \underline{v}^i)}$$

Notice that since $v^i > \underline{v}^i$ we have that $\underline{\delta}^i < 1$.

Let $\underline{\delta} = \max_i \underline{\delta}^i$. Then, it is clear that for $\delta > \underline{\delta}$, no player wants to deviate from the strategy $\square$

2.4 Static Games of Incomplete Information

Players take decisions without knowing the actions taken by other players

Basic Assumptions

1. Players are rational and self-interested (Rationality)

2. Some players don't have full information about all the elements of the game (Incomplete Information)

3. Players know that players are rational and fully informed, and they know that others know that players are rational and fully informed, and they know that others know that they know that players are rational and fully informed, and ... (Common Knowledge)

Notation

- $\mathcal{I} = \{1, 2, \ldots, I\}$ is the *Set of Players*
- $S^i = \{s_1^i, s_2^i, \ldots\}$ is the *Set of Pure Strategies of player* $i \in \mathcal{I}$
- $S = S^1 \times S^2 \times \cdots \times S^I$ is the *Set of pure strategy profiles*
- T^i is the set of possible "types" for player $i \in \mathcal{I}$
- $T = T^1 \times T^2 \times \cdots \times T^I$ is the set of types profiles (usually assume that T is finite)
- $u^i : S \times T \to \mathbb{R}$ is the *(von Neumann - Morgenstern) Payoff function of player* $i \in \mathcal{I}$ *that depends on both, the strategy profile and the types profile*
- p is the *common prior* distribution (known to all the players) over the set of possible types profiles T such that $\forall t \in T,\ p(t) > 0$

- Each player $i \in \mathcal{I}$ knows his type t_i, but there is uncertainty about all other agents types t_{-i}. Let agent i be of type $\bar{t}_i$. Now, based on the common prior p and $\bar{t}_i$, we can capture agent i's belief about t_{-i} by the conditional probability distribution that is induced by p and $\bar{t}_i$:

$$p(t_{-i}|\bar{t}_i) = \frac{p(\bar{t}_i, t_{-i})}{p(\bar{t}_i)} = \frac{p(\bar{t}_i, t_{-i})}{\sum_{t_{-i} \in T_{-i}} p(\bar{t}_i, t_{-i})}$$

A *Bayesian Game* is fully characterized

$$\Gamma_{BN} = (\mathcal{I}, \{S^i\}_{i \in \mathcal{I}}, \{T^i\}_{i \in \mathcal{I}}, \{u^i\}_{i \in \mathcal{I}}, p)$$

This representation of a game is known as the "normal form **bayesian** game" or "strategic form **bayesian** game".

Equilibria in Games of Incomplete Information

- What solution concept can we use for incomplete information games?

- We simply can use the solution concepts that we developed so far.

- Trick: Interpret each type as a separate player. This implies a larger set of players $\tilde{\mathcal{I}}$ and an enlarged game $\tilde{\Gamma}_{BN}$. Nature, represented by p chooses which types are actually realized. The strategies for the large game $\tilde{\Gamma}_{BN}$ are obtained from the original game.

Notice

- Any solution for $\tilde{\Gamma}_{BN}$ implies a solution for any game that is actually realized.

- A Nash equilibrium in $\tilde{\Gamma}_{BN}$ induces a Nash equilibrium in the game that actually is to be played after "nature's move".

The Associate Strategic Form Game

- $\mathcal{J} = \{j \mid j = (i, t_i)$ where $i \in \mathcal{I}$ and $t_i \in T^i\}$ is the *extended* set of players
- $R^j = S^i$ is the set of strategies for player $j = (i, t_i) \in \mathcal{J}$
- $R = R^1 \times R^2 \times \cdots \times R^J$ is the *Set of pure strategy profiles* (where $J = |\mathcal{J}| = |\mathcal{I}| \times |T|$)
- $\tilde{u}^j : R \to \mathbb{R}$ is the *(von Neumann - Morgenstern) Payoff function* of player $j \in \mathcal{J}$ that is given by

$$\tilde{u}^j(r) = \sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) u(r_{|\text{agents in type profile } (t_i, t_{-i})}; (t_i, t_{-i}))$$

The game $\tilde{\Gamma}_{BN} = (\mathcal{J}, \{R^j\}_{j \in \mathcal{J}}, \{\tilde{u}^j\}_{j \in \mathcal{J}})$ is the strategic form game associated with the game of incomplete information Γ_{BN}

Example

MWG, p. 254, Example 8.E.1.

With probability μ prisoner 2 is of type I ("normal") and with probability $1-\mu$ prisoner 2 is of type II. Payoffs are:

1\(2,I)	C	D
C	0,-2	-10,-1
D	-1,-10	-5,-5

1\(2,II)	C	D
C	0,-2	-10,-7
D	-1,-10	-5,-11

In this case,

- $\mathcal{J} = \{1, (2, I), (2, II)\}$
- $R^1 = S^1 = \{C, D\}$, $R^{(2,I)} = S^2 = \{C, D\}$, $R^{(2,II)} = S^2 = \{C, D\}$
- $\tilde{u}^1(r_1, r_{(2,I)}, r_{(2,II)}) = \mu u^1(r_1, r_{(2,I)}; 1) + (1-\mu)u^1(r_1, r_{(2,II)}; 1)$
- $\tilde{u}^{(2,I)}(r_1, r_{(2,I)}, r_{(2,II)}) = u^2(r_1, r_{(2,I)}); I)$
- $\tilde{u}^{(2,II)}(r_1, r_{(2,I)}, r_{(2,II)}) = u^2(r_1, r_{(2,II)}); II)$

Bayesian Nash Equilibrium

Definition. Bayesian Nash Equilibrium
A Bayesian-Nash equilibrium of a game of incomplete information Γ_{BN} *is a Nash equilibrium of the associated strategic form game* $\tilde{\Gamma}_{BN}$.

Note that in a Bayesian Nash equilibrium, each player must be playing a best response to the conditional distribution of his opponent strategies for each type that he might end up having.

Theorem. Existence of Bayesian Nash equilibria
Every finite strategic game of incomplete information has at least one Bayesian Nash equilibrium.

Proof. Trivial □

Example: Mixed strategies revisited

Reconsider the Battle of the Sexes Game, including the following incomplete information (t_i's are independently drawn from the uniform distribution on $[0, \varepsilon]$):

	F	O
F	$2 + t_1, 1$	$0, 0$
O	$0, 0$	$1, 2 + t_2$

- Find a pure strategy Bayesian Nash equilibrium
- Show that if the incomplete information disappears $(\varepsilon \to 0)$ the players behavior in the pure BNE approaches their behavior in the mixed-strategy NE in the original complete information game.

Harsanyi (1973) showed that this result is quite general: a mixed strategy Nash equilibrium in a game of complete information can (almost always) be interpreted as a pure strategy Bayesian Nash equilibrium in a closely related game with a *small* amount of incomplete information.

Notice that in both cases one player is uncertain about the other player's choice, but there may be different reasons for the uncertainty: it can arise because of randomization or because of a little incomplete information.

2.5 Dynamic Games of Incomplete Information

Players take decisions one after another. Thus, some players may know the actions taken by other players before choosing its own

Basic Assumptions

1. Players are rational and self-interested (Rationality)

2. Some players don't have full information about all the elements of the game (Incomplete Information)

3. Players know that players are rational and fully informed, and they know that others know that players are rational and fully informed, and they know that others know that they know that players are rational and fully informed, and ... (Common Knowledge)

Similarly as before, we can apply the solution concept developed for normal form games to extensive form games as well. However, from our experience with complete information dynamic games, we know that some information is lost when considering an extensive form game in its normal form.

To rule out "unreasonable" equilibria we considered refinements using backwards induction and subgame perfection. Therefore, when extending the normal form solution concept of Bayesian Nash equilibria to the extensive form, we at the same time try to extend the idea of subgame perfection.

In order to capture sequential rationality in dynamic games with incomplete information, we will use the concept of a *Perfect Bayesian equilibrium*. To develop this notion, we start with a weaker version that also applies to dynamic games with complete information.

Weak Perfect Bayesian Equilibrium

Consider the game below

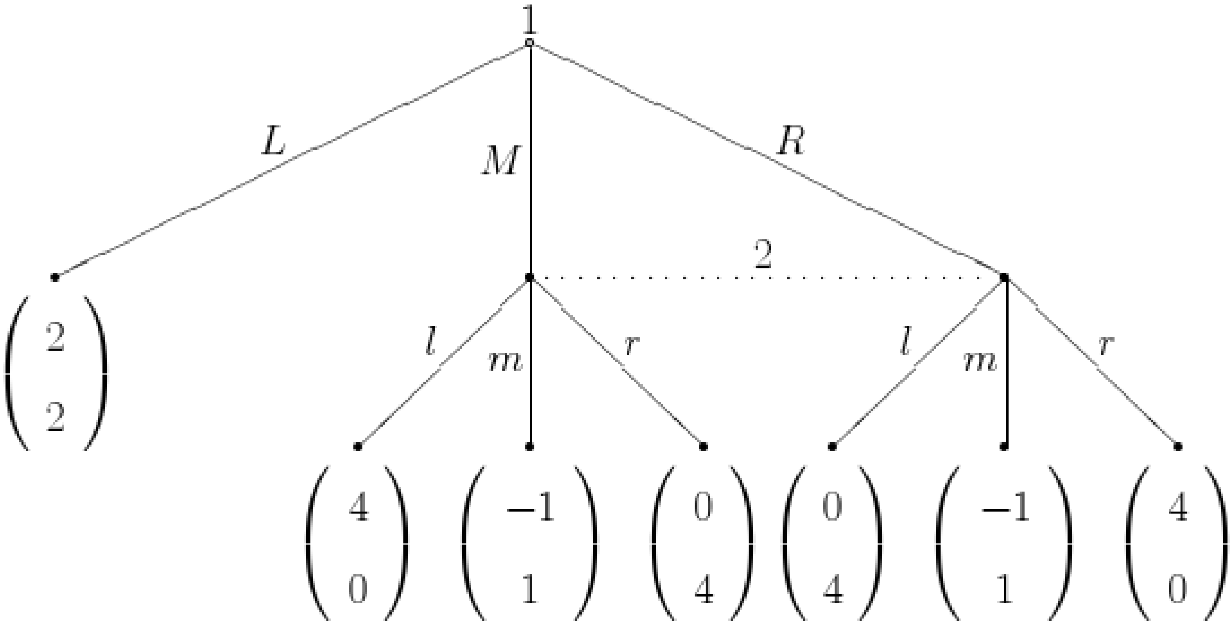

Figure 4:

Notice that there is a unique subgame: the game itself. Thus, any Nash equilibrium of the game will be a Subgame perfect Nash equilibrium as well.

Using the normal form representation

	l	m	r
L	2,2	2,2	2,2
M	4,0	-1,1	0.4
R	0,4	-1,1	4.0

Table 4: The strategic form representation of the game in Figure4

We find that (L, m) is the unique Nash equilibrium in pure strategies and, thus, the unique Subgame perfect Nash equilibrium (in pure strategies)

Suppose now that Player 2 "believes" that the game is at the right node of her information set with probability μ(and at the left node with probability $1 - \mu$)[2]

For what values of μ is m a "sequentially rational" strategy for Player ? That is, for what values of μ is the expected payoff of m larger than the expected payoff of l and r?

[2]Also, M and R can be interpreted as the same strategy for 2 different types of Player 1. Them, μ is the common prior distribution over the set of types

$$\begin{aligned} E_\mu(u^2(l)) &= 4(1-\mu) \\ E_\mu(u^2(m)) &= 1 \\ E_\mu(u^2(r)) &= 4\mu \end{aligned}$$

Now,

$$\begin{aligned} E_\mu(u^2(m)) &> E_\mu(u^2(l)) \Rightarrow 1 > 4(1-\mu) \Rightarrow \mu > \frac{3}{4} \\ E_\mu(u^2(m)) &> E_\mu(u^2(r)) \Rightarrow 1 > 4\mu \Rightarrow \mu < \frac{1}{4} \end{aligned}$$

Thus, m is never sequentially rational ! It cannot be "rationalized" by any "beliefs" Player 2 might have regarding her position in his information set.

We will require that a player's strategy is optimal for some belief that he or she might have about his opponents strategies. This criterium leads to Weak perfect Bayesian equilibria.

Definition. System of Beliefs
A system of beliefs μ *in an extensive form game* Γ_E *is a function* $\mu : N \to [0, 1]$ *that specifies a probability for each decision node such that for all information sets* H, $\sum_{x \in H} \mu(x) = 1$. *Such a system of beliefs* μ *embodies the beliefs of all players at each of their information sets about the history of play up to this point.*

By $u^i(H, \mu, \sigma)$ we denote player i's expected utility starting at information set H if his beliefs regarding the conditional probabilities of being at the various nodes of H are given by μ, he plays strategy σ_i, and all other players play according to σ_{-i}.

Definition. Sequential Rationality
A strategy profile σ *in* Γ_E *is sequentially rational at information set* H *for a system of beliefs* μ *if for the player* $\iota(H)$ *who moves at* H *and for any of his strategies* $\tilde{\sigma}_{\iota(H)}$ *we have*

$$u^{\iota(H)}(H, \mu, (\sigma_{\iota(H)}, \sigma_{-\iota(H)})) \geq u^{\iota(H)}(H, \mu, (\tilde{\sigma}_{\iota(H)}, \sigma_{-\iota(H)}))$$

The definition of a Weak perfect Bayesian equilibrium requires that

- strategies are sequentially rational given beliefs μ
- whenever possible, beliefs are consistent with the given strategies: in equilibrium, players should have correct beliefs about their opponents' strategy choices.

Definition. Weak perfect Bayesian equilibrium
A profile of strategies and a system of beliefs (σ, μ) *is a Weak perfect Bayesian equilibrium (weak PBE) for* Γ_E *if it has the following properties:*

(i) the strategy profile is sequentially rational given belief system μ and

(ii) the system of beliefs μ is derived through Bayes' rule whenever possible. That is, for any information set H that can be reached with positive probability under σ (i.e., $p(H|\sigma) > 0$) we have that for all $x \in H$:

$$\mu(x) = \frac{p(x|\sigma)}{p(H|\sigma)}$$

Theorem. Existence of Weak perfect bayesian equilibria
Every finite strategic game of incomplete information has at least one Weak perfect Bayesian equilibrium.

The following result sheds some light on the conditions that a Nash equilibrium must verify to be a Weak perfect bayesian equilibrium. It also shows that Weak perfect bayesian equilbria are refinements of Nash equilibria

Proposition. Relation between Weak perfect bayesian equilibrium and Nash equilibrium
A strategy profile is a Nash equilibrium of Γ_E if and only if there exists a system of beliefs μ such that

(i) the strategy profile is sequentially rational given belief system μ at all information sets H such that $p(H|\sigma) > 0$ and

(ii) the system of beliefs μ is derived through Bayes' rule whenever possible.

Example

Consider the game in Figure 4 whose strategic form representation is at Table 4

The game has a total of 6 Nash equilbria

	σ^1_L	σ^1_M	σ^1_R	σ^2_l	σ^2_m	σ^2_r
1	1	0	0	0	1	0
2	1	0	0	$\frac{3}{5}$	$\frac{2}{5}$	0
3	1	0	0	$\frac{1}{2}$	0	$\frac{1}{2}$
4	1	0	0	0	$\frac{2}{5}$	$\frac{3}{5}$
5	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{2}$	0	$\frac{1}{2}$
6	0	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	0	$\frac{1}{2}$

Which one/ones is/are Weak perfect bayesian equilibria ?

Let's start with the strategies of Player 2 that are sequentially rational at her information set. For convenience, let's call $x(M)$ the left node on the information set of Player 2 and $x(R)$ the right node. Let the *beliefs* μ of Player 2 be such that $p(x(M)|H) = \mu$ (and $p(x(R)|H) = 1 - \mu$). Remember that m is never sequentially rational and thus we must consider only l and r with possitive probability. The corresponding expected payoffs (according to μ) are

$$\begin{aligned} E_\mu(l) &= 4(1-\mu) \\ E_\mu(r) &= 4\mu \end{aligned}$$

clearly,

$$\begin{aligned} E_\mu(l) &> E_\mu(r) \Leftrightarrow \mu < \frac{1}{2} \\ E_\mu(l) &< E_\mu(r) \Leftrightarrow \mu < \frac{1}{2} \\ E_\mu(l) &= E_\mu(r) \Leftrightarrow \mu = \frac{1}{2} \end{aligned}$$

From here we can consider 3 cases

a) $\mu < \frac{1}{2}$ In this case, $\sigma^2 = (1,0,0)$. Then, given σ^2, the best reply for Player 1 is $\sigma^1 = (0,1,0)$. But his can not be an equilibrium because in such case we would have that

$$\frac{p(x(M)|\sigma)}{p(H|\sigma)} = \frac{1}{1} \neq \mu$$

b) $\mu > \frac{1}{2}$ In this case, $\sigma^2 = (0,0,1)$. Then, given σ^2, the best reply for Player 1 is $\sigma^1 = (0,0,1)$. But his can not be an equilibrium because in such case we would have that

$$\frac{p(x(M)|\sigma)}{p(H|\sigma)} = \frac{0}{1} \neq \mu$$

c) $\mu = \frac{1}{2}$ In this case Player 2 is indiferent between l and r. Thus, any σ^2 such that ${\sigma_m}^2 = 0$ is sequentially rational for Player 2. Now, for the belief $\mu = \frac{1}{2}$ to be consistent with the equilibrium path it must be the case that $\sigma^1_M = \sigma^1_R$, and this happens at all Nash equilibria. Thus, the Weak perfect bayesian equilibria are all the Nash equilibria such that ${\sigma_m}^2 = 0$, that is:

	σ^1_L	σ^1_M	σ^1_R	σ^2_l	σ^2_m	σ^2_r
3	1	0	0	$\frac{1}{2}$	0	$\frac{1}{2}$
5	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{2}$	0	$\frac{1}{2}$
6	0	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	0	$\frac{1}{2}$

Notice that all other Nash equilibria are sequentially rational along the equilibrium path, but not off the equilibrium path.

Strengthening the Weak perfect bayesian equilibrium concept

One problem of weak PBE is that no restrictions are placed on beliefs "off the equilibrium path" Thus, extra consistency restrictions on beliefs are often added, which result in Perfect Bayesian equilibria (Fudenberg and Tirole 1991). Other refinements were obtained by Kreps and Wilson (1982, Sequential equilibria) and Selten (1975, Trembling hand perfection).

Definition. Sequential Equilibrium
A profile of strategies and a system of beliefs (σ, μ) *is a Sequential Equilibrium for* Γ_E *if it has the following properties:*

(i) the strategy profile is sequentially rational given belief system μ and

(ii) there exists a sequence of completely mixed strategies $\{\sigma^k\}_{k=1}^{\infty} \to \sigma$ such that $\mu = \lim_{k\to\infty} \mu^k$, where

$$\mu^k(x) = \frac{p(x|\sigma^k)}{p(H|\sigma^k)}$$

Proposition.
Every sequential equilibrium is both Weak perfect bayesian equilibrium and Subgame perfect Nash equilibrium

Example: Signaling Games

These sender-receiver game features two players. Player 1, the sender, has private information regarding his own type and sends a signal to Player 2, the receiver, who observes Player 1's action (but not his type) before choosing her own

Examples

- The sender is a worker who knows his productivity and must choose a level of education. The receiver is a firm, which observes the workers education, but not the productivity level, and then decides on the wage to offer.

- The sender is a driver who knows his probability of having an accident and must choose an insurance policy. The receiver is an insurance company, which observes the policy proposed by the sender, but no the probability of an accident, and then decides whether to accept it or not

A signaling game consists of

- Player 1, who is the sender, has private information on his type $\theta \in \Theta$ (the use of θ instead of t_1 is common in signaling games) and chooses an action $a_1 \in A_1$

- Player 2, whose type is common knowledge, is the receiver who observes a_1 and chooses $a_2 \in A_2$

- Mixed strategies and their spaces are denoted by $\alpha_i \in \mathcal{A}_i$. The payoff associated with mixed strategy profile (α_1, α_2) and type θ is $u_i(\alpha_1, \alpha_2, \theta)$

- Before the game begins, it is common knowledge that Player 2 holds prior beliefs p over Player 1's type θ

- Player 1's strategies:$\sigma_1 : \Theta \rightarrow \mathcal{A}_1$; i.e., a strategy prescribes a probability distribution $\sigma_1(\cdot|\theta)$ over actions in A_1 for each type $\theta \in \Theta$.

- Player 2's strategies: :$\sigma_2 : A_1 \rightarrow \mathcal{A}_2$; i.e., a strategy prescribes a probability distribution $\sigma_2(\cdot|a_1)$ over actions in A_2 for each action a_1.

- Payoffs given $\sigma = (\sigma_1, \sigma_2)$ and type θ are

Player 1

$$U^1(\sigma|\theta) = \sum_{a_1} \sum_{a_2} \sigma_1(a_1|\theta)\sigma_2(a_2|a_1)u^1(a_1, a_2, \theta)$$

Player 2

$$U^2(\sigma|\theta) = \sum_{\theta} p(\theta) \left(\sum_{a_1} \sum_{a_2} \sigma_1(a_1|\theta)\sigma_2(a_2|a_1)u^2(a_1, a_2, \theta) \right)$$

Perfect Bayesian Equilibrium

Definition. *Perfect Bayesian Equilibrium of a signaling game*
A profile of strategies σ^* *and a system of (posterior) beliefs* $\mu(\cdot|a_1)$ *is a Perfect Bayesian Equilibrium of a signaling game if it has the following properties:*

(i) for all θ,

$$\sigma_1^*(\cdot|\theta) \in \arg\max_{\alpha_1} u^1(\alpha_1, \sigma_2^*, \theta)$$

(ii) for all a_1,

$$\sigma_2^*(\cdot|a_1) \in \arg\max_{\alpha_2} \sum_{\theta} \mu(\theta|a_1) u^2(a_1, \alpha_2, \theta)$$

(iii) If $\sum_{\theta' \in \Theta} p(\theta')\sigma_1^*(a_1|\theta') > 0$, then

$$\mu(\theta|a_1) = \frac{p(\theta)\sigma_1^*(a_1|\theta)}{\sum_{\theta' \in \Theta} p(\theta')\sigma_1^*(a_1|\theta')}$$

In practical terms

Definition. *Perfect Bayesian Equilibrium of a signaling game*
A profile of pure strategies σ^* and a system of (posterior) beliefs $\mu(\cdot|a_1)$ is a Perfect Bayesian Equilibrium of a signaling game if it has the following properties:

(i) for all θ,

$$\sigma_1(\theta) \in \arg\max_{a_1} u^1(a_1, \sigma_2(a_1), \theta)$$

(ii) for all a_1,

$$\sigma_2(a_1) \in \arg\max_{a_2} \sum_{\theta} \mu(\theta|a_1) u^2(a_1, a_2, \theta)$$

(iii) If $\sum_{\theta' \in \{\Theta|\sigma_1(\theta')=a_1\}} p(\theta') > 0$, then

$$\mu(\theta|a_1) = \frac{p(\theta)}{\sum_{\theta' \in \{\Theta|\sigma_1(\theta')=a_1\}} p(\theta')}$$