A game is a formal representation of a situation in which individuals interact in a setting of strategic interdependence.

Strategic interdependence

⇒each individual's utility depends not only on his own actions but on the actions chosen by other individuals.

What action is best or "optimal" for each agent may depend on what others choose.

Therefore, decision making must take into account expectation of how other players act.

Four basic elements of a game:

Players

- agents that interact

Rules

- who moves when, what do they know or observe at each point of move, what they choose from....

Outcomes

- for each possible configuration of actions by all players what is the eventual outcome of the interaction - may not be quantifiable

Payoffs

- the players' preferences or utility function defined over possible outcomes.

Games may involve randomness (exogenous uncertainty or randomization in choice of actions)

Players may need to evaluate probability distributions or lotteries over outcomes.

Assume: each agent has preferences over all lotteries over outcomes of the game that are representable by an expected utility function.

The *payoff function* of a player: her Bernoulli utility: $\{$space of outcomes of the game$\} \to \mathbb{R}$.

The actual utility levels are called payoffs.

Games may involve direct conflict of interest or objectives.

Ex. Matching pennies (zero sum game).

Games may involve no conflict of interest.

Ex. Pure coordination game.

Games may involve both conflict of interest and coordination problems.

Ex. Battle of Sexes.

Extensive Form Representation of a Game.

Captures

- who moves when (the sequencing of moves),
- what actions each player may choose from at each point of decision making
- what they know about other players and previous actions chosen by others at each point where they have to move in the game,
- how each configuration of action choices by players through the game generates an outcome....

Finite games: finite number of players, finite number of possible actions, finite number of moves.

Can use game tree to depict the extensive form.

Elements of a game tree:

* Decision nodes (points at which players are required to make decisions):

- Initial Nodes

- Successor Nodes

* Each action at a decision node leads to a distinct branch of the tree.

* Terminal nodes: where game terminates and an outcome of the game is realized.

* Payoff vectors at each terminal node indicating payoffs realized at that outcome.

Exogenous Uncertainty in the play of the game: modeled as move of nature.

Games of Perfect Information:

Games where at each point of decision, every player observe all *prior* decisions made in course of the play of the game:

In terms of the game tree, at every decision node, players observe every action chosen in prior decision nodes that lead up to that decision node - a player knows exactly which decision node she is at.

Game of Imperfect Information.

May not observe action chosen by a previous mover in the game.

A player may not therefore know which decision node she is at.

She may know that is anywhere among a set of multiple nodes: Information Set.

In games of imperfect information, players make decisions at information sets consisting possibly of multiple nodes.

Singleton information set: Just one decision node.

Though actions chosen at an information set can lead to different outcomes depending on which node the player is really at (i.e., what unobservable actions were actually chosen in prior moves by other players)

- the player herself does not know which decision node she is at.

The set of actions she chooses from when she is at an information set must be independent of the true decision node she happens to be in.

One Shot Simultaneous Move Game: is a game of imperfect information.

No player observes the action chosen by other players when she makes her decision.

Assume: perfect recall.

Player does not forget what she observed at an earlier stage of the game.

Assume: Common knowledge of the structure of the game.

In an extensive form game, this implies all players know the extensive form.

Strategy:

A complete contingent plan or decision rule that specifies how the player will act in each possible *distinguishable* circumstance in which she might be called upon to move i.e., in each information set where she is may be possibly required to make a choice.

Given the strategies of all players, the actual play of the game may not require the players to face all contingencies that their strategy covers - all information sets may not be reached.

Definition. Let $\mathcal{H}_i$ denote the collection of information sets where player i can possibly be required to make a decision, $\mathcal{A}$ the set of possible actions in the game and $C(H) \subset \mathcal{A}$ the set of actions possible at an information set H.

A strategy for player i is a function $s_i : \mathcal{H}_i \to \mathcal{A}$ such that $s_i(H) \in C(H)$ for all $H \in \mathcal{H}_i$.

A strategy profile in a game with I players is a vector $s = (s_1, ..., s_I)$ where s_i is the strategy chosen by player i.

Also denoted sometimes as (s_i, s_{-i}) where s_{-i} is a $(I-1)$ vector consisting of a strategy choice for each player other than player i.

Normal Form Representation of a Game:

Every profile of strategies $s = (s_1, ..., s_I)$ induces an outcome of the game:

- a sequence of moves actually taken

$\Rightarrow$a probability distribution over terminal nodes of the game

$\Rightarrow$a probability distribution over payoff realizations of the game

$\Rightarrow$ *expected* payoff (utility) $u_i(s_1, ..., s_I)$ for each player i.

Definition: For a game with I players, the normal form representation Γ_N specifies for each player i a set of strategies S_i (with $s_i \in S_i$) and a payoff function $u_i(s_1, ..., s_I)$ giving the VNM utility levels associated with the (possibly random) outcomes arising from strategies $(s_1, ..., s_I)$.

Formally, $\Gamma_N = [I, \{S_i\}_{i=1}^I, \{u_i(.)\}_{i=1}^I]$.

Normal form: no information about moves, order of moves, sequencing, how the "strategy" of each player is composed or played or even what it means.

Can be seen as a simultaneous move game where players choose their strategies (rather than actions at various decision nodes),

For any extensive form game, unique normal form representation.

Converse not true.

Players may randomize over actions at any decision node.

Choose probability distributions over deterministic or *pure* strategies.

Such randomized strategies are called *mixed* strategies.

Suppose that the S_i, the (pure) strategy set of each player i is finite.

A mixed strategy by player i denoted by $\sigma_i : S_i \to [0, 1]$ assigns to each pure strategy $s_i \in S_i$ a probability $\sigma_i(s_i)$ that it will be played where $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$.

The set of all possible mixed strategies of player i is denoted by $\triangle(S_i)$.

Every profile of mixed strategies (one for each player) generates a probability distribution over outcomes and payoffs of the game.

As players have VNM utility on the space of lotteries over outcomes, we payoff to each player from a mixed strategy profile is the expected utility (or payoff) generated.

Let $S = S_1 \times S_2 \times ... \times S_I$.

Let $\sigma = (\sigma_1, ..., \sigma_I)$ be a profile of mixed strategies where players randomize independently (not correlated strategies).

Player i's VNM utility or payoff from this mixed strategy profile, denoted by $u_i(\sigma)$, is given by

$$u_i(\sigma) = \sum_{(s_1,...,s_I) \in S} [\sigma_1(s_1)....\sigma_I(s_I)] u_i(s_1, ..., s_I)$$

If strategy set is not finite, each mixed strategy is captured by a probability distribution function and the payoffs can be similarly defined.

Normal form game allowing for mixed strategies: denoted by $\Gamma_N = [I, \{\triangle(S_i)\}, \{u_i\}]$

In extensive form games, we can allow players to randomize over actions at each information set where she is required to act.

Sometimes called *behavior strategies*.

Simultaneous Move Games (Normal Form Games).

Consider normal form game $\Gamma_N = [I, \{S_i\}_{i=1}^I, \{u_i(.)\}_{i=1}^I]$ where we confine players to use only pure strategies.

Prisoner's Dilemma

$$\begin{bmatrix} 1\downarrow, 2\rightarrow & Not\ Confess & Confess \\ Not\ Confess & -2,-2 & -10,-1 \\ Confess & -1,-10 & -5,-5 \end{bmatrix}$$

(Strictly) Dominant Strategy for each player: Confess.

Let $S_{-i} = S_1 \times S_{i-1} \times S_{i+1} ... \times S_I$ denote the product of strategy sets of all players other than player i.

Definition: A strategy $s_i \in S_i$ is a strictly dominant strategy for player i in a game $\Gamma_N = [I, \{S_i\}_{i=1}^I, \{u_i(.)\}_{i=1}^I]$ if for all $s_i' \neq s_i, s_i' \in S_i$, we have

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

If a player has a strictly dominant strategy, it is individually optimal for the player to play it irrespective of her belief about what other players play.

In fact, it is the unique individually optimal strategy.

If every player has a strictly dominant strategy, it is obvious that all players should play this.

However, the outcome obtained as a result may be "collectively or jointly suboptimal" or "Pareto inefficient" in the sense that all players could have been better off if they had played according to a different strategy profile.

An example of how self interested individual behavior may not be collectively good.

Reason: each player determines his or her "optimal" strategy by looking at his or her own payoff ignoring the payoffs of other players.

"Externality".

It is rare for strictly dominant strategies to exist.

What strategy is optimal for a player often depends on what other players play.

However, a rational player will never play a strategy that is dominated by some other strategy (i.e., leads to strictly lower payoff no matter what other players play).

Definition: A strategy $s_i \in S_i$ is a *strictly dominated* strategy for player i in a game $\Gamma_N = [I, \{S_i\}_{i=1}^I, \{u_i(.)\}_{i=1}^I]$ if there exists another strategy $s_i' \in S_i$ such that

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

In this case, we say s_i' strictly dominates s_i.

A strictly dominated strategy should not be played by a rational player no matter what he believes about the strategy choice of other players.

$$\begin{bmatrix} 1\downarrow,2\rightarrow & L & R \\ U & 1,-1 & -1,1 \\ M & -1,1 & 1,-1 \\ D & -2,5 & -3,2 \end{bmatrix}$$

Both U and M strictly dominate D.

Note that if there is a strictly dominant strategy for a player , it strictly dominates every other strategy of the player (and vice-versa).

Definition: A strategy $s_i \in S_i$ is a *weakly dominated* strategy for player i in a game $\Gamma_N = [I, \{S_i\}_{i=1}^I, \{u_i(.)\}_{i=1}^I]$ if there exists another strategy $s_i' \in S_i$ such that

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

and further, there exists $\widehat{s}_{-i} \in S_i$ such that

$$u_i(s_i', \widehat{s}_{-i}) > u_i(s_i, \widehat{s}_{-i}).$$

In this case, we say s_i' weakly dominates s_i.

$$\begin{bmatrix} 1\downarrow,2\rightarrow & L & R \\ U & 5,1 & 4,0 \\ M & 6,0 & 3,1 \\ D & 6,4 & 4,4 \end{bmatrix}$$

D weakly dominates U and M.

If a strategy for a player weakly dominates every other strategy in the strategy set of the player, we say it is a *weakly dominant* strategy.

Unlike a strictly dominated strategy, a rational player may play a weakly dominated strategy (if he/she has certain kind of belief about what the other players play).

Cannot be ruled out ex ante.

Rationality $\Rightarrow$ Rules out strictly dominated strategies.

Common knowledge of rationality

$\Rightarrow$ *Iterated Elimination of Strictly Dominated Strategies.*

Prisoner's Dilemma Modified (bias in favor of prisoner 1).

$$\begin{bmatrix} 1\downarrow,2\rightarrow & Not\ Confess & Confess \\ Not\ Confess & 0,-2 & -10,-1 \\ Confess & -1,-10 & -5,-5 \end{bmatrix}$$

$1\downarrow,2\rightarrow$	L	M	R
T	$-1,7$	$4,5$	$4,10$
C	$0,11$	$1,4$	$3,2$
B	$-1,19$	$2,10$	$1,-1$

Order of deletion does not affect the set of strategies that survive iterated elimination of strictly dominated strategies.

Can generalize strictly dominated and dominant strategy concepts to normal form games that allow for mixed strategies in a straightforward way.

$$\begin{bmatrix} 1\downarrow,2\rightarrow & L & R \\ U & 10,1 & 0,4 \\ M & 4,2 & 4,3 \\ D & 0,5 & 10,2 \end{bmatrix}$$

Playing U and D with probability $\frac{1}{2}$ each strictly dominates M.

RATIONALIZABILITY.

Pushes the idea of iterated deletion using common knowledge of rationality to its fullest possible extent.

Definition: In game $\Gamma_N = [I, \{\Delta(S_i)\}_{i=1}^{I}, \{u_i(.)\}_{i=1}^{I}]$, strategy σ_i is *a best response for player i to his rival' strategies σ_{-i}* if

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i}) \text{ for all } \sigma_i' \in \Delta(S_i).$$

Definition: Strategy σ_i is *never a best response for player i* if there is no σ_{-i} for which σ_i is a best response for player i.

i.e., there is no belief that player i may hold about his opponents' strategy choices that justifies choosing strategy σ_i.

A strictly dominated strategy is never a best response.

For two player games, a strategy is never a best response if and only if it is strictly dominated.

In games of more than two players there can be (mixed) strategies that are not strictly dominated but are never a best response.

Rationalizable Strategies (Bernheim & Pearce, 1984):

Definition: In game $\Gamma_N = [I, \{\Delta(S_i)\}_{i=1}^{I}, \{u_i(.)\}_{i=1}^{I}]$, the strategies in $\Delta(S_i)$ that survive iterated elimination of strategies that are never a best response are known as player i's rationalizable strategies.

Set of rationalizable strategies $\subset$ Set of strategies that survive iterated elimination of strictly dominated strategies.

$$\begin{bmatrix} 1\downarrow, 2\rightarrow & b_1 & b_2 & b_3 & b_4 \\ a_1 & 0,7 & 2,5 & 7,0 & 0,1 \\ a_2 & 5,2 & 3,3 & 5,2 & 0,1 \\ a_3 & 7,0 & 2,5 & 0,7 & 0,1 \\ a_4 & 0,0 & 0,-2 & 0,0 & 10,-1 \end{bmatrix}$$

b_4 is never a best response as it is strictly dominated by a strategy that plays b_1 and b_3 with probability $\frac{1}{2}$ each.

After eliminating b_4, in the remaining game a_4 is strictly dominated by a_2.

In the remaining game, every pure strategy is a best response to some other pure strategy.

Set of rationalizable pure strategies for player 1 is $\{a_1, a_2, a_3\}$

Set of rationalizable pure strategies for player 2 is $\{b_1, b_2, b_3\}$

For each rationalizable strategy, one can construct a valid chain of justification for choosing to play this strategy.

For example, player 1 can justify playing a_2 by the belief that player 2 will play b_2 and player 1 can justify this belief by believing that player 2 believes that player 1 will play a_2 which can be justified by the belief that player 2 thinks that player 1 thinks that player 2 will play b_2....

Infinite chain of justification: $(a_2, b_2, a_2, b_2, ...)$

Chain does not break down.

Another chain:

$(a_1, b_3, a_3, b_1, a_1, b_3, a_3, b_1, a_1...)$

Cannot justify playing a_4 (which is not a rationalizable strategy) with an infinite chain of justification.

Nash Equilibrium.

Consider normal form game $\Gamma_N = [I, \{S_i\}, \{u_i\}]$ where players restrict themselves to pure strategies.

Definition 1 *A strategy profile* $s^* = (s_1^*, s_2^*, ..., s_I^*) \in S$ *constitutes a Nash Equilibrium (NE) if for every* $i = 1, ...I$,

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

for all $s_i \in S_i$.

Each player's strategy is *a* best response to the strategies actually played by rivals.

$$s_i^* \in b_i(s_{-i}^*), i = 1, ...I$$

where $b_i(s_{-i}^*)$ is the best-response (or best reply or "reaction") correspondence defined by

$$b_i(s_{-i}^*) = \{s_i \in S_i : s_i solves \max_{s_i' \in S_i} u_i(s_i', s_{-i}^*)\}.$$

*No player has a (strict) incentive to *unilaterally* deviate from playing according to strategy profile s^* (does not rule out gainful deviation by a coalition of multiple players).In a NE, players play rationally holding correct conjectures (or forecasts) of rivals' play.

Therefore, NE strategies are rationalizable.

NE: stronger than rationalizability which only requires the players play rationally given some reasonable conjecture about rivals play (i.e., those that can be similarly justified).

Example:

$$\begin{bmatrix} & b_1 & b_2 & b_3 \\ a_1 & 0,7 & 2,5 & 7,0 \\ a_2 & 5,2 & 3,3 & 5,2 \\ a_3 & 7,0 & 2,5 & 0,7 \end{bmatrix}$$

Every pair of pure strategies (a_i, b_j) is a rationalizable - every strategy is a best response to some strategy.

However, there is a unique NE: (a_2, b_2).

* Let N denote the set of NE strategy profiles,

R the set of rationalizable strategy profiles,

IED the set of strategy profiles that survive iterated elimination of strictly dominated strategies and

U the set of strategy profiles consisting of strategies that are strictly undominated.

Then,

$$N \subset R \subset IED \subset U.$$

The concept of NE is based on the concept of mutually correct expectations.

Quite often, there can be multiple NE.

Coordination problems.

Example: Coordination game.

$$\begin{bmatrix} & L & R \\ U & 100,100 & 0,0 \\ D & 0,0 & 1000,1000 \end{bmatrix}$$

The two NE are Pareto-ranked (both players better off in (D, R) compared to (U, L).

Example: (Pure coordination game)

$$\begin{bmatrix} & L & R \\ U & 100,100 & 0,0 \\ D & 0,0 & 100,100 \end{bmatrix}$$

Example: Battle of Sexes

$$\begin{bmatrix} & Opera & Game \\ Opera & 100,1000 & 0,50 \\ Game & 50,0 & 1000,100 \end{bmatrix}$$

Example: Cake eating.

A cake is to be divided among two players.

Players 1 and 2 simultaneously choose the shares $(s_1, s_2), 0 \leq s_i \leq 1$, of the cake they demand.

The payoff of each player i is the share of the cake obtained be her and is given by:

$$\begin{aligned} x_i &= s_i, \text{ if } s_i + s_j \leq 1, \\ &= 0, \text{ if } s_i + s_j > 1. \end{aligned}$$

Set of NE $= \{(s_1, s_2) : s_1 + s_2 = 1, 0 \leq s_i \leq 1, i = 1, 2\}$

Continuum of NE. Conflict of objectives across NE.

Why should we expect conjectures to be correct?

Certainly not a necessary consequence of rationality or common knowledge of rationality and payoffs.

* If there is a unique predicted outcome for a game (a unique obvious way to play the game), then it must be a Nash equilibrium.

* If certain outcomes are *focal* (Schelling) for cultural or other reasons (having to do with information not contained within the description of the game), then such an outcome can be a prediction only if it is Nash equilibrium.

* If players make a *non-binding* agreement prior to play about how they are going to play the game, then such an agreement is credible only if it is a Nash equilibrium (the pre-game communication makes the agreement focal).

*Stable social convention (norm): If the game is played repeatedly, then some stable social convention about how to play the game may emerge (a limit of some dynamic adjustment process); such a stable social convention or norm must be a NE.

Mixed Strategy Nash Equilibrium.

Consider the normal form game $\Gamma_N = [I, \{\Delta(S_i)\}, \{u_i\}]$

Definition 2 *A strategy profile* $\sigma^* = (\sigma_1^*, \sigma_2^*, ..., \sigma_I^*) \in \prod_{i=1}^{I} \Delta(S_i)$ *constitutes a Nash Equilibrium (NE) if for every* $i = 1, ...I$*,*

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*)$$

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

Example (Matching Pennies)

$$\begin{bmatrix} & H & T \\ H & +1,-1 & -1,+1 \\ T & -1,+1 & +1,-1 \end{bmatrix}$$

There is no NE in pure strategies.

Each player playing H and T with probability $\frac{1}{2}$ each constitutes a mixed strategy NE.

Given this strategy of rival, each player indifferent between playing H or T.

In any mixed strategy NE, each player is indifferent between pure strategies that she plays with strictly positive probability

i.e., given the mixed strategies played by other players, all such pure strategies must yield her exactly her the same expected utility or payoff (which would also be her NE payoff).

Further, no pure strategy that is played with probability zero by a player can yield strictly higher payoff than the payoff from the pure strategies that are played with strictly positive probability.

[In case the strategy set is not finite, the above must be true for *almost every* strategy in *the support* of the mixed strategy of each player].

The following proposition is written for the case of finite strategy sets and shows that the above is both necessary as well as sufficient for a mixed strategy NE:

Proposition. Assume S_i is finite. Let $S_i^+ \subset S_i$ denote the set of pure strategies that player i plays with strictly positive probability in a mixed strategy profile $\sigma^* = (\sigma_1^*, ..., \sigma_I^*) \in \prod_{i=1}^{I} \Delta(S_i)$. Strategy profile σ^* is a NE *if and only if* for all $i = 1, ...I$

$$\begin{aligned} (i)\ u_i(s_i, \sigma_{-i}^*) &= u_i(s_i', \sigma_{-i}^*), \forall s_i, s_i' \in S_i^+ \\ (ii)\ u_i(s_i, \sigma_{-i}^*) &\geq u_i(s_i', \sigma_{-i}^*), \forall s_i \in S_i^+ \text{ and } \forall s_i' \in S_i - S_i^+. \end{aligned}$$

So to test whether a given mixed strategy profile is a NE we only need to test that all pure strategies played with strictly positive probability yield equal payoffs for each player and that no player can do better by playing some other *pure* strategy.

Example:

$$\begin{bmatrix} & L & R \\ U & 100,100 & 0,0 \\ D & 0,0 & 1000,1000 \end{bmatrix}$$

Suppose player 1 plays U and D with probability p and $1-p$, respectively.

For player 2, playing L yields expected payoff $100p$ and playing R yields $1000(1-p)$.These two expected payoffs are equal only if $p = \frac{1}{11}$. By a symmetric argument, player 1 is indifferent between U and D if and only if player 2 plays L and R with probabilities $\frac{1}{11}$ and $\frac{10}{11}$, respectively.

Thus, these mixed strategies constitute a NE.

Existence of NE:

* Every normal form game where the strategy sets of all players are finite has a mixed strategy NE.

* Every normal form game where:

(1) the strategy set of each player is a nonempty, convex and compact subset of $\mathbb{R}^n$

(2) payoff function $u_i(s_1, ...s_I)$ of each player is continuous in $(s_1, ...s_I)$ and quasi-concave in s_i

has a NE in *pure strategies*.

If quasi-concavity of u_i fails but continuity holds, there is mixed strategy NE.

Dasgupta and Maskin (1986).

Games of Incomplete Information: Bayesian Nash Equilibrium.

Games analyzed thus far: games of complete information (assumes common knowledge of players, payoffs, rules of the game etc.)

Games of incomplete information: players may not know other players' preferences over outcomes i.e. their payoff (or utility) function.

Harsanyi approach: imagine each player's preference structure or payoff function is randomly chosen by nature at the beginning of the game according to some *commonly known* probability distribution.

The actual realization of nature's draw is only observed by the player (private information) while others play the game only knowing the probability distribution used by nature.

This captures uncertainty about the preferences of other players.

More specifically, each possible preference structure or payoff function of a player is defined as a possible *type* of the player.

Nature (player 0) first chooses the realization of a random variable that determines the *type* of every player.

The realized type of player i is observed only by player i.

Example.

Consider prisoner's dilemma where player 1 is the DA's brother (this is known by both players).

The DA has some discretion and so if player 1 and 2 keep mum, he can let player 1 go free.

Otherwise, the punishments are same as in the usual prisoner's dilemma.

Suppose further that player 2 may either be purely selfish (type 1) or someone who hates to rat on his buddy (type 2).

In the latter case, player 2 gets a psychological dis-utility equivalent to 6 additional months in prison if he confesses.

Player 1 is purely selfish.

It is commonly known that player 2 is selfish with probability $\mu \in [0, 1]$.

Of course, player 2 knows whether or not he is selfish.

This can be formalized as the following extensive form game:

first nature chooses type of player 2 from a probability distribution that assigns probability μ to type 1 and probability $1-\mu$ to type 2.

This move of nature is not observed by player 1 but observed by player 2;

this determines their information sets.

Players than simultaneously choose whether to confess (C) or not confess (NC).

If player 2 is chosen by nature to be of type 1, the payoffs from this latter simultaneous move game are

$$\begin{bmatrix} & NC & C \\ NC & 0,-2 & -10,-1 \\ C & -1,-10 & -5,-5 \end{bmatrix}$$

while if player 2 is chosen by nature to be of type 2, then the payoffs are:

$$\begin{bmatrix} & NC & C \\ NC & 0,-2 & -10,-7 \\ C & -1,-10 & -5,-11 \end{bmatrix}.$$

Note player 2 will know which payoff matrix is relevant when he chooses whether or not to confess because he will know his preference, but player 1 will not.

A pure strategy for player 2 must specify what he is going to do for each choice by nature of his type (complete contingent plan)- his pure strategy set is {(C if type 1, C if type 2), (C if type 1, NC if type 2), (NC if type 1, C if type 2), (NC if type 1, NC if type 2)}.

Note that though we know that player 2 actually knows his type, to play this game it is important for player 1 to imagine how player 2 would play the game if he was of each possible type and to choose rationally accordingly.

This is the basic reason behind the Harsanyi formulation.

Let Θ_i be the set of all possible types of player i and $\Theta = \Theta_1 \times ... \times \Theta_I$.

The type of player i denoted by $\theta_i \in \Theta_i$ is a random variable chosen by nature whose realization is observed only by player i.

The joint probability distribution of the types of all players is given by $F(\theta_1, \theta_2, ...\theta_I)$ which is assumed to be common knowledge.

Note that the random variables $\theta_1, \theta_2, ...\theta_I$ need not be independent.

Each player i has a payoff function $u_i(s_1, ..s_I, \theta_i)$.

Here, $s_i \in S_i$,the set of all actions that can be chosen by player i.

The Bayesian game is summarized by $[I, \{S_i\}, \{u_i\}, \Theta, F]$.

Pure strategy of player i is function or decision rule $s_i(\theta_i)$: $\Theta_i \to S_i$.

Let Σ_i be the set of all such functions - the set of pure strategies of player i.

Player i's expected payoff from any profile of pure strategies $(s_1(.), ...s_I(.))$ is then given by

$$\begin{aligned} & \widetilde{u}_i(s_1(.), ...s_I(.)) \\ = \quad & E_\theta[u_i(s_1(\theta_1), ...s_I(\theta_I), \theta_i] \end{aligned}$$

A Nash equilibrium of the "reduced" normal form game $[I, \{\Sigma_i\}, \{\widetilde{u}_i\}]$ is called a Bayesian-Nash equilibrium.

Definition 3 *A (pure strategy) Bayesian Nash equilibrium (BNE) 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 the game* $\Gamma_N = [I, \{\Sigma_i\}, \{\widetilde{u}_i\}]$ *i.e., for every* $i = 1, ...I,$

$$\widetilde{u}_i(s_i(.), s_{-i}(.)) \geq \widetilde{u}_i(s_i'(.), s_{-i}(.)), \forall s_i'(.) \in \Sigma_i.$$

One implication of this:

for each possible type that he may have, in a BNE, a player plays an action that is a best response to the conditional distribution of his opponents' strategies .

Proposition 4 *A profile of decision rules* $(s_1(.), ...s_I(.))$ *is a BNE* if, and only if, *for all* i *and all* $\overline{\theta}_i \in \Theta_i$ *occurring with positive probability (or more generally,* almost every $\overline{\theta}_i \in \Theta_i$)

$$E_{\theta_{-i}}[u_i(s_i(\overline{\theta}_i), s_{-i}(\theta_{-i}), \overline{\theta}_i \quad | \quad \overline{\theta}_i] \geq$$
$$E_{\theta_{-i}}[u_i(s_i', s_{-i}(\theta_{-i}), \overline{\theta}_i \quad | \quad \overline{\theta}_i]$$

for all $s_i' \in S_i$ *where* $E_{\theta_{-i}}$ *is the conditional expectation taken over all possible realizations of other player's' types (conditional on player i's types).*

We can think of *each player of each type* as being a distinct player directly playing an action from S_i and maximizing his payoff given the conditional probability distribution over the strategy choices of rivals.

Earlier example:

Nature$\bullet$

$\mu \swarrow \qquad \searrow 1-\mu$

$$\begin{bmatrix} & NC & C \\ NC & 0,-2 & -10,-1 \\ C & -1,-10 & -5,-5 \end{bmatrix} \qquad \begin{bmatrix} & NC & C \\ NC & 0,-2 & -10,-7 \\ C & -1,-10 & -5,-11 \end{bmatrix}$$

BNE:

Player 2 plays (C if type 1, NC if type 2).

Player 1 plays NC (C) if $\mu \leq (\geq) \frac{1}{6}$.

Example.

First price, sealed bid auction with private independent valuations.

2 bidders $i = 1, 2$.

Valuation of bidder $i : v_i$ - known only by bidder i.

Valuations are independently and uniformly distributed on $[0, 1]$.

Net utility of a bidder with valuation $v_i : v_i - p$,if he gets the good and pays price p; otherwise, its zero.

Bidder i's bid: $b_i \geq 0$.

Both bidders simultaneously submit bids and highest bidder wins the good and pays the price she bids.

In case of tie, each bidder wins with probability $\frac{1}{2}$.
Bayesian game:
Type of a player i: v_i.
Types drawn independently by nature from the uniform distribution on $[0,1]$.
Action space of each player of each type: $\mathbb{R}_+$.
Payoff of player i :

$$\begin{aligned} u_i(b_1, b_2; v_i) &= v_i - b_i, \text{ if } b_i > b_j \\ &= \frac{v_i - b_i}{2}, \text{ if } b_i = b_j \\ &= 0, \text{ if } b_i < b_j. \end{aligned}$$

Strategy of player $i : b_i(v_i)$.
BNE:
Given $b_j(v_j)$, for each $v_i \in [0,1]$, $b_i(v_i)$ solves

$$\max_{b_i}[(v_i - b_i)\Pr\{b_i > b_j(v_j)\} + \frac{1}{2}(v_i - b_i)\Pr\{b_i = b_j(v_j)\}]. \tag{1}$$

Claim: $b_i(v_i) = \frac{v_i}{2}, i = 1, 2,$constitutes a BNE.
To see this, suppose that $b_j(v_j) = \frac{v_j}{2}$.
Then, b_j is uniformly distributed on $[0, \frac{1}{2}]$.
In particular, for any b_i,

$$\Pr\{b_i = b_j(v_j)\} = \Pr\{b_i = \frac{v_j}{2}\} = 0.$$

Also, for any $b_i \geq 0$,

$$\begin{aligned} \Pr\{b_i &> b_j(v_j)\} = \Pr\{b_i > \frac{v_j}{2}\} \\ &= \Pr\{v_j < 2b_i\} \\ &= \min\{1, 2b_i\}. \end{aligned}$$

Thus, the maximand in (1), reduces to

$$\begin{aligned} & [(v_i - b_i)\min\{1, 2b_i\}] \\ =& \ v_i - b_i, \text{ if } b_i \geq \frac{1}{2} \\ =& \ 2(v_i - b_i)b_i \text{ if } b_i \leq \frac{1}{2}. \end{aligned}$$

Check that for each type $v_i \in [0,1]$,the optimal solution is $b_i = \frac{v_i}{2}$ (which is also $\leq \frac{1}{2}$).

In fact, this is the unique BNE where the strategy of each player is linear in valuation.