Advanced Microeconomics (ES30025)

Topic One: Introduction to Game Theory (i)

I. Introduction

We use game theory to analyze situations in which agents are conscious that the strategies they adopt affect each other, i.e., for problems of interaction between (rational) agents (*the players*) where the *payoff* of each player depends on the *strategies* of the other(s). For simple examples, the word strategy could be replaced here by action to help understanding. But as the theory is developed, the distinction between actions and strategies will become important.

Modern studies of industrial organisation rely heavily on the ideas of game theory because they investigate situations where the outcome is determined by the actions of more than one agent. Note that in monopoly only one seller is involved, while in perfect competition the price-taking assumption guarantees independence of players' choices. If we want to move away from these simple cases, we need some heavier technical artillery. This is what game theory provides. It is also a useful organising principle.

II. Describing Games

We describe games in two ways - *extensive* or *normal* form.

Extensive Form Representation

These specify the order of play, the information and choices available to every player when it is his turn to play, the payoffs to the players, and (sometimes) the probability of moves by nature – see Figure 1: Note: Each 'node' represents a point in the game where a player needs to make a decision. The payoffs to each player are represented at the bottom of the game tree. Figure 1 represents a *dynamic game* - there is a clear temporal sequence with Player 2 moving after Player 1 and thus knowing what Player 1 has played (decided) before he makes his decision. We assume *common knowledge* - the information in the game is known by all players, and all players in the game know that everybody else knows that, and so on.

In the game depicted in Figure 1 the possible actions by either player is to play *L* or *R*. One particular *strategy* for Player 2 is to play *L* irrespective of what Player 1 does. A *Pure Strategy* is a choice by a player of a given action with certainty (probability 1).[1] For example, Player 2 *always* plays *R* if Player 1 plays *L*, and *always* plays *L* if Player 1 plays *R*.

[1] That is, as opposed to a *mixed strategy* which is a randomisation over pure strategies - for example, Player 1 plays L with probability p and R with probability $(1 - p)$.

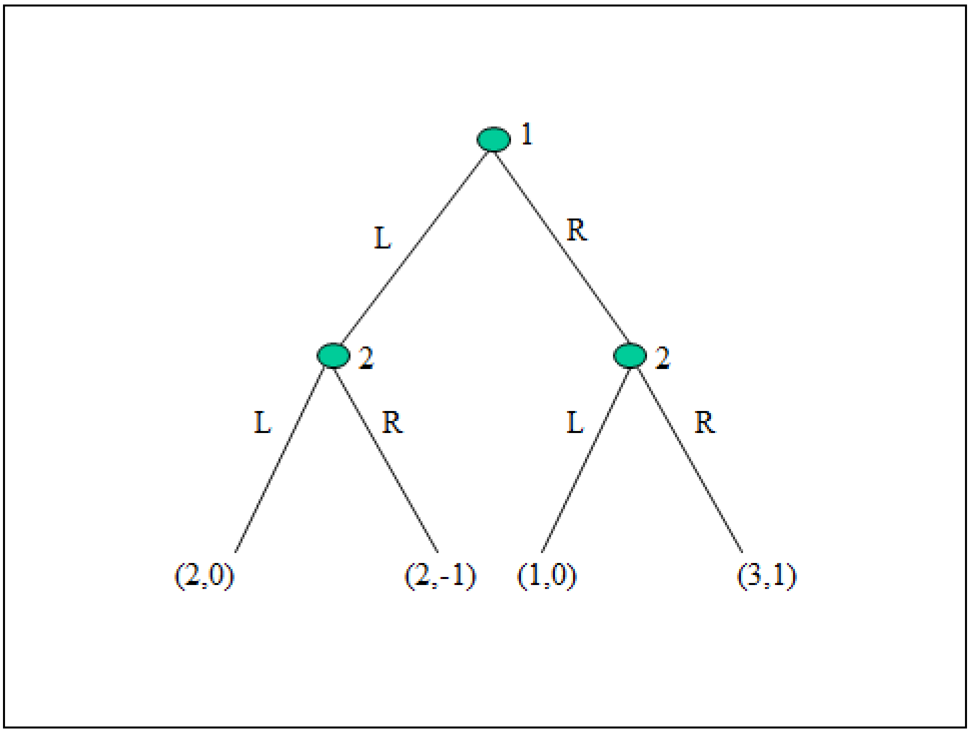

Figure 1: A Sequential Game (Game One)

The static (or simultaneous move) version of the previous game is depicted in Figure 2 following:

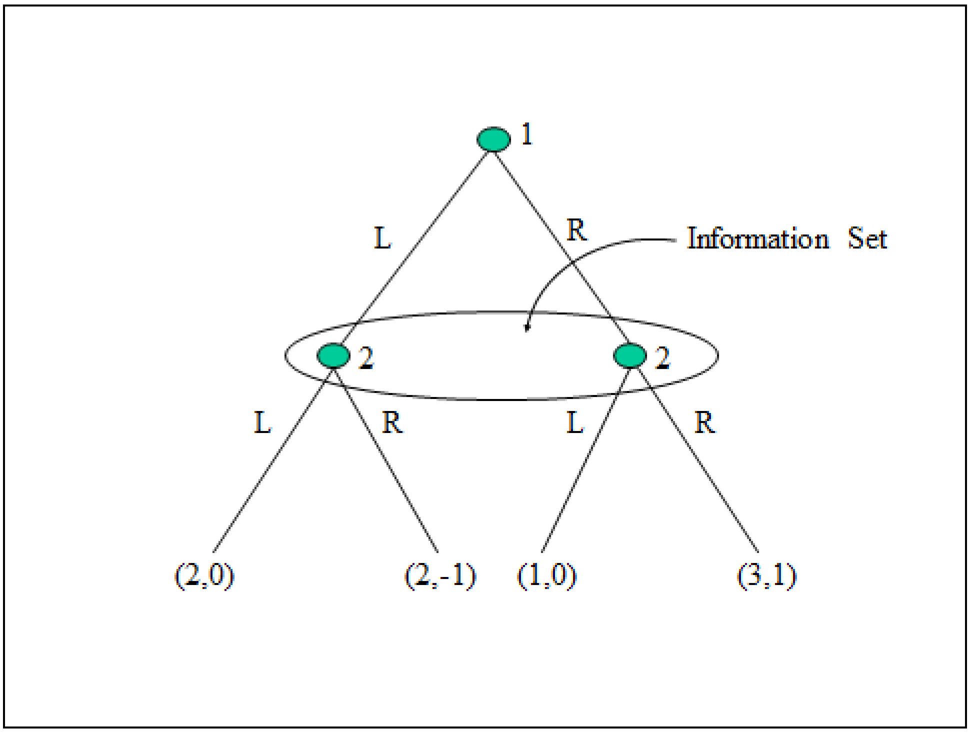

Figure 2: A Static Game (Game Two)

An *Information Set* describes the state of knowledge of a particular player when he is making a choice. A player cannot distinguish between particular nodes within an information set. Thus the game represented in Figure 2 is *static* in the sense that Player 2 does not know what

Player 1 has done until after he chooses his strategy. In this information sense, Player 2 moves simultaneously with Player 1, even though temporally they may move in sequence.

Normal Form Representation

A game in *normal form* consists of: (i) a group of *N* players; (ii) a description of the actions they can take; (iii) a description of the payoffs that they each receive under all possible configurations of chosen actions. Normal form games take no account of the order in which actions are taken. They can thus be seen as useful in situations where the order of play is unimportant, or where choices of actions are simultaneous, or where each player only learns of his antagonists' choices after he has made his own choice.

A normal form representation is therefore the collection of pure strategies available to each player at each of the information sets in the extensive form and the associated payoffs. The normal form representation of the game set our in extensive form in Figure 2 is given by:

Player 2 / Player 1	$s_2^1 = L$	$s_2^2 = R$
$s_1^1 = L$	2, 0	2, -1
$s_1^2 = R$	1, 0	3, 1

Table 1: A Static Game (Game Two)

The table shows the actions of the two players and the payoffs they receive. The first element of each pair in the table indicates the payoff of Player 1, the second the payoff of Player 2. Thus if 1 chooses *L* and 2 chooses *R*, Player 1 gets 2 and Player 2 gets -1. The strategies available to each player are:

$$\begin{aligned} s_1 &= \left(s_1^1, s_1^2\right) = (L, R) \\ s_2 &= \left(s_2^1, s_2^2\right) = (L, R) \end{aligned} \tag{1}$$

where s_i^j denotes the *j*th strategy available to player *i*. In the dynamic game (i.e. the game set out in extensive form in Figure 1) the actions of Player 2 can be contingent on what Player 1 does. Thus we have:

$$\begin{aligned} s_1 &= \left(s_1^1, s_1^1\right) = (L, R) \\ s_2 &= \left(s_2^1, s_2^2, s_2^3, s_2^4\right) = \{(L, L), (R, R), (L, R), (R, L)\} \end{aligned} \tag{2}$$

Player 2 / Player 1	s_2^1 (*L*, *L*)	s_2^2 (*R*, *R*)	s_2^3 (*L*, *R*)	s_2^1 (*R*, *L*)
$s_1^1 = L$	2, 0	2,-1	2, 0	2, -1
$s_1^2 = R$	1, 0	3, 1	3, 1	1, 0

Table 2: A Sequential Game (Game One)

In words, $s_2^3 = (L, R)$ is Player 2's third strategy, which tells him to 'follow' Player 1 - i.e. to play *L* if Player 1 plays $s_1^1 = L$ and to play *R* if Player 2 plays $s_1^2 = R$.

Thus, we can define an *action* as the set of choices available to a player at each decision node and a *strategy* as a rule that tells a player what action to take at each information set in the game. A *strategy* thus specifies the order of play, the information and choices available to every player when it is his turn to play, the payoffs to the players, and (sometimes) the probability of moves by nature; it is a detailed plan that tells a player that action to take at every contingency.

III. Computing Outcomes of Games

Consider the following ways of computing the outcome of games.

Elimination of Dominated Strategies

A *dominated strategy* is one that always gives a lower payoff to a player than another strategy, irrespective of the actions the other players choose. Consider the following game:

$$\begin{aligned} s_1 &= \left(s_1^1, s_1^2\right) = \left(War, Peace\right) \\ s_2 &= \left(s_2^1, s_2^2\right) = \left(War, Peace\right) \end{aligned} \tag{3}$$

Player 2 / Player 1	$s_2^1 = War$	$s_2^2 = Peace$
$s_1^1 = War$	**$\underline{1}$, $\underline{1}$**	$\underline{3}$, 0
$s_1^2 = Peace$	0, $\underline{3}$	2, 2

Table 3: War and Peace[2]

This game has dominant *strategy equilibrium (DSE)*. *War* is the best strategy for Player 1 irrespective of what Player 2 chooses. If Player 2 chooses *War*, Player 1 gets 1 by choosing *War* and 0 by choosing *Peace*. If Player 2 chooses Peace, Player 1 gets 3 by choosing *War* and 2 by choosing *Peace*. Check that Player 2's choice of *War* is similarly better for her for both 1's possible choices.

Note: A strategy is said to be dominant for player i if, no matter what other players do, i is best off using it. An equilibrium in dominant strategies exists if all players have a dominant strategy.

Consider now the case of two prisoners arrested and questioned separately for a crime they actually committed. There is insufficient evidence against them and if neither confesses to the crime they will both set free after some minimal incarceration. If one prisoner confesses and the other one does not (i.e. 'splits') then the splitter is rewarded (i.e. let off) and the sticker is punished severely. If both confess they are both punished. Assume the payoffs (i.e. punishments) are as follows:

[2] Note - underlined numbers indicate best responses.

Player 2 / Player 1	s_2^1 *Confess*	s_2^2 *Deny*
s_1^1 = *Confess*	-8, -8	0, -10
s_1^2 = *Deny*	-10, 0	-1, -1

Table 4: The Prisoners' Dilemma

It is apparent that 'deny' is a dominated strategy - whatever each player anticipates the other player doing, he will never choose 'deny'.[3] Since a rational player would never play a dominated strategy we could remove these from the game with the result that the outcome of the game must be amongst the remaining possibilities.

Player 2 / Player 1	s_2^1 *Confess*	~~s_2^2~~ ~~*Deny*~~
s_1^1 = *Confess*	**-8, -8**	~~0, -10~~
~~s_1^2 = *Deny*~~	~~-10, 0~~	~~-1, -1~~

Table 5: Equilibrium in the Prisoners' Dilemma

In the Prisoners' Dilemma, the elimination of dominated strategies yields a unique predicted dominant-strategy equilibrium $\left(s_1^1, s_2^1\right)$ = (*Confess*, *Confess*). But this is more the exception than the rule - not all games have an equilibrium in dominant strategies:

Comments:

- The Prisoners' Dilemma is a static, one-shot, nonzero sum, non-cooperative, game with common knowledge.
- $\left(s_1^2, s_2^2\right)$ = (Deny, Deny) is a (Pareto) superior outcome, but cannot be achieved in the absence of binding commitments.
- If, for a given player, a particular strategy gives a payoff that is better than the payoff from every other strategy for that player, the strategy is said to be *strictly* (or *strongly*) *dominant*. If, instead, the strategy gives a payoff that is better than or equal to the payoff from every other strategy for that player, the strategy is said to be *weakly dominant*.
- If all players have strictly dominant strategies, we have a *Strictly Dominant Strategy Equilibrium* (*SDSE*). If at least one player has a weakly dominant strategy and any other player has a strictly dominant strategies, then we have a *Weakly Dominant Strategy Equilibrium* (*WDSE*).

Iterated Elimination of (Strictly) Dominated Strategies

Sometimes, although a *SDSE* may not exist, we may be able to eliminate from consideration some strategies by another method. Once these strategies are eliminated, we may then be able to find a *SDSE*. Consider the following normal form game (like the Prisoners' Dilemma, the game is non-zero sum, non-cooperative, common knowledge, static and one-shot):

[3] Note that in this simple game the strategies and actions of each player are the same. This is not always the case. In the game depicted in Figure 1 the actions and strategies of Player 2 are not the same.

Player 2 Player 1	s_2^1 *Left*	s_2^2 *Middle*	s_2^3 *Right*
s_1^1 = *Up*	1, 0	1, 2	0, 1
s_1^2 = *Down*	0, 3	0, 1	2, 0

Table 6

If player 2 were choosing *Left* or *Middle*, player 1's best strategy, would be to choose *Up*. But if player 2 were choosing *Right*, player 1's best strategy would be to choose *Down*. Player 1 does not therefore have a (weakly or strictly) dominated strategy. If player 1 were choosing *Up*, player 2's best strategy would be to choose *Middle*. If player 1 were choosing *Down*, player 2's best strategy would be to choose *Left*. Player 2 therefore does not have a (weakly or strictly) dominated strategy. However, comparing the last two columns, it can be seen that irrespective of what player 1 chooses, *Right* is strictly dominated by *Middle* for player 2. Suppose both players recognize this - i.e. that *Right* would never be chosen by Player 2. They can therefore eliminate it from consideration. Deleting this column from the normal form gives

Player 2 Player 1	s_2^1 *Left*	s_2^2 *Middle*	$\cancel{s_2^3}$ ~~*Right*~~
s_1^1 = *Up*	1, 0	1, 2	~~0, 1~~
s_1^2 = *Down*	0, 3	0, 1	~~2, 0~~

Table 7

Now we find that player 1 has a strictly dominant strategy (*Up*), though player 2 still does not have one. However, since *Up* is a strictly dominant strategy for player 1, *Down* can be eliminated from consideration. Deleting this row yields:

Player 2 Player 1	s_2^1 *Left*	s_2^2 *Middle*	$\cancel{s_2^3}$ ~~*Right*~~
s_1^1 = *Up*	1, 0	1, 2	~~0, 1~~
~~s_1^2 = *Down*~~	~~0, 3~~	~~0, 1~~	~~2, 0~~

Table 8

Or:

Player 2 Player 1	s_2^1 *Left*	s_2^2 *Middle*
s_1^1 = *Up*	1, 0	**1, 2**

Table 9

Now *Middle* is a strictly dominant strategy for player 2. We have therefore derived an *iterated strictly dominant strategy equilibrium* (*ISDSE*) to the original game *vis*: (s_1^1, s_2^2) = (*Up*, *Middle*).

Comments:

- Note that without the elimination of *Right*, the elimination of *Down* would not have been possible.
- That players should eliminate strictly dominated strategies seems reasonable. However, each player is assumed to apply accurately some quite sophisticated logic and to rely on the other player applying the same logic. If there are many strategies (i.e. rows and/or columns) in a game, this may be an unrealistic assumption.
- Especially in games with many strategies, it is possible to eliminate strategies in different orders. However, this does not affect the final result.
- Suppose that the requirement for elimination is only that a strategy is weakly dominated. Then it is possible to follow a similar procedure to that above for any given game. It is possible, however, that with one order of elimination one solution will be obtained, whilst, with a different order of elimination a different solution is obtained. This is no help to the players who are trying to work out the solution so that they can pick their strategies. Because of this problem, iterated elimination of weakly dominant strategies is not usually considered.
- The principles of dominance and iterative dominance do not usually go very far in limiting the number of possible outcomes - 'too many things' can happen! Recall the game set out in Table 1:

Player 2 / Player 1	$s_2^1 = L$	$s_2^2 = R$
$s_1^1 = L$	2, 0	2, -1
$s_1^2 = R$	1, 0	3, 1

Table 10: Game Two Revisited

Here no strategy is dominated and so we make no progress in determining the outcome of the game. Similarly, consider the following game:

Player 2 / Player 1	s_2^1 = *Opera*	s_2^2 = *Football*
s_1^1 = *Opera*	2, 1	0, 0
s_1^2 = *Football*	0, 0	1, 2

Table 11: Battle of the Sexes

Here, whatever Player 1 chooses, Player 2 is best off choosing the same thing. The same is true of Player 2. Whereas in a game with dominant strategy equilibrium, there is no real interdependence between the behaviour of the players, in this game, the players' best choice depends on the choice of the other.

Thus, since not all games have an equilibrium in dominant strategies, we need a 'tighter' solution concept.

N.B. An *ISDSE* may not exist. For example, try the following game

Player 2 / Player 1	s_2^1 *Left*	s_2^2 *Middle*	s_2^3 *Right*
s_1^1 = *Up*	0, 4	4, 0	5, 3
s_1^2 = *Middle*	4, 0	0, 4	5, 3
s_1^3 = *Down*	3, 5	3, 5	6, 6

Table 12: Game Three

Equilibrium Outcomes Concepts

We require a criterion that eliminates as many outcomes as possible without being so tight that it eliminates outcomes that are plausible. The more outcomes that are eliminated, the more predictive power the theory has. But if too many plausible outcomes are eliminated, then this increase in the theory's predictive power is at a high price - the theory's relevance to the real world is reduced. 'Equilibrium refinement' - the process of finding grounds on which to reduce the set of possible outcomes of games – helps to strike the right balance between these two requirements. Perhaps the most powerful equilibrium concept is that devised by Nash.

IV. Nash Equilibrium in Pure Strategies

A Nash Equilibrium is a set of strategies from which no player, taking his opponents' strategies as given, wishes to deviate - there is no regret![4] Formally:

Definition

A set of strategies is a Pure Strategy Nash Equilibrium (PSNE) iff:

$$\pi_i\left(s_i^*, s_{-i}^*\right) \geq \pi_i\left(s_i, s_{-i}^*\right) \qquad \forall i \tag{4}$$

where $s_i = \left(s_i^1, s_i^2, \ldots, s_i^n\right)$ denotes the set of strategies available to player *i*, -*i* denotes 'all players *except* player *i*', and $\pi_i\left(s_i^*, s_{-i}^*\right)$ denotes the payoff to Player i as a function of his and every other player's strategy. In words, a *Nash Equilibrium* is a set of strategies such that no player, taking his opponents' strategy as given, wishes to change his own strategy. Thus an outcome is a Nash equilibrium if no player can gain by deviating from it, given that no other player deviates. For illustration, a Nash equilibrium in output between two oligopolists would be defined:

$$\begin{aligned} &\pi_1\left(q_1^*, q_2^*\right) \geq \pi_1\left(q_1, q_2^*\right) \\ &and \\ &\pi_2\left(q_1^*, q_2^*\right) \geq \pi_1\left(q_1^*, q_2\right) \end{aligned} \tag{5}$$

In words, q_1^* $\left(q_2^*\right)$ maximises Firm 1's (2's) profit given that Firm 2 (1) is producing q_2^* $\left(q_1^*\right)$.

Computation

In *Normal Form* games one can compute Nash Equilibria by finding the best response of each player to the strategies of the other players. A Nash Equilibrium, therefore, is a 'box' which is the best response for every player to the other players 'playing' the same box.

To check that a particular outcome is a Nash equilibrium take each outcome in turn and see if either player can gain by deviating. If neither can then the outcome is a Nash

[4] *Nash Equilibrium* is named for John Nash, who discussed it in a path-breaking paper in *Econometrica* (1951). *Subgame Perfect Equilibrium* is one of a large number of refinements of Nash equilibrium. It was described by Reinhardt Selten in 1965 (also *Econometrica*). These two shared the 1996 Nobel prize for Economics with John Harsanyi, who invented a third equilibrium concept (*Bayesian Perfect Equilibrium*) applicable to games with imperfect information. A biography of Nash has recently been published by Sylvia Nasir (*A Beautiful Mind*), the film of which, staring Russell 'Gladiator' Crowe, won several Oscars in 2002.

equilibrium. In the game depicted in Table 11, (*Opera*, *Opera*) is a Nash equilibrium because if Player 1 were to choose to go to *Football* alone, he would do worse than going to the *Opera* with Player 2. Similarly, Player 2 is better off going to *Opera* with Player 1, than going to *Football* alone. Player 2's best outcome of all is to go to the *Football* with Player 1, but this is irrelevant to the Nash equilibrium. The two players' optimal actions are rendered interdependent by the structure of the game. Recall also Game Two:

Player 2 / Player 1	$s_2^1 = L$	$s_2^2 = R$
$s_1^1 = L$	**2**, **0**	2, -1
$s_1^2 = R$	1, 0	**3**, **1**

Table 13: Nash Equilibria in Game Two

Consider Player 1; If he believes Player 2 will play *L*, his best response is to play *L* also; If he believes Player 2 will *R*, his best response is to play *R* also. Now consider Player 2; if he believes Player 1 will play *L*, his best response is to play *L* also; if he believes Player 1 will *R*, his best response is to play *R* also. There are two situations, then, in which neither player will regret his choice of action, namely if both players play *L* or both players play *R* - for example, if Player 1 plays *L* the best Player 2 could do is to play *L* also, and if Player 2 plays *L*, the best Player 1 can do is to play *L* also. Similarly if they both play *R*. There are thus two Nash Equilibria *vis*:

$$\text{Nash Equilibria: } s^* = \left\{\left(s_1^1 = L, s_2^1 = L\right);\left(s_1^2 = R, s_2^2 = R\right)\right\} \quad (6)$$

Note then that Nash Equilibria need not be unique - this example has two equilibria. Similarly, recall Game One:

Player 2 / Player 1	s_2^1 (*L*, *L*)	s_2^2 (*R*, *R*)	s_2^3 (*L*, *R*)	s_2^1 (*R*, *L*)
$s_1^1 = L$	**2**, **0**	2,-1	2, 0	2, -1
$s_1^2 = R$	1, 0	**3**, **1**	**3**, **1**	1, 0

Table 14: Multiple Equilibria in Game One

$$\textit{Nash Equilibria}: s^* = \left\{\left[s_1^1 = L, s_2^1 = (L,L)\right];\left[s_1^2 = R, s_2^2 = (R,R)\right];\left[s_1^2 = R, s_2^3 = (L,R)\right]\right\} \quad (7)$$

Note here that two of the three equilibria yield the same payoffs in equilibria.

It is also true that Nash equilibrium does not always exist. This is illustrated by the next example.

Player 2 / Player 1	$s_2^1 = Left$	$s_2^2 = Right$
$s_1^1 = Up$	0, 2	1, 0
$s_1^2 = Down$	2, 0	0, 1

Table 15

To be sure, then, consider the following game:

Player 2 / Player 1	s_2^1 *a*	s_2^2 *b*	s_2^3 *c*
$s_1^1 = A$	6, 4	**3**, **6**	0, 5
$s_1^2 = B$	4, 3	2, 0	2, 2
$s_1^3 = C$	7, 0	1, 1	4, 0

Table 16

Player 1 reasons as follows: If player 2 chooses *a* my best reply would be *C;* If player 2 chooses *b* my best reply would be *A;* If player 2 chooses *c* my best reply would be *C*. Player I therefore does not have a (weakly or strictly) Dominant Strategy. Similarly for Player 2: If player 1 chooses *A* my best reply would be *b*; If player 1 chooses *B* my best reply would be *a*; If player 1 chooses *C* my best reply would be *b*. Player 2 does not have (weakly or strictly) Dominant Strategy either. But in this example there is a single Nash Equilibrium (i.e. a combination of strategies from which either player has an incentive to deviate, given that the other does not deviate. In the example $\left(s_1^1, s_2^2\right)$ = (*A*, *b*) is a Nash Equilibrium because if Player 2 chooses *b*, player 1's best response is *A*, and if Player 1 chooses *A*, player 2's best response is *b*. Note:

- Nash Equilibrium relates to weakly preferred strategies - if I get a payoff of 5 in a Nash Equilibrium, and there is another (non-Nash Equilibrium) strategy that would also give me 5, I do not have an incentive to deviate from the Nash Equilibrium.
- Every *DSE* (strict or weak) is a Nash Equilibrium, but a Nash Equilibrium is not necessarily a *DSE*.
- Is a Nash Equilibrium reasonable? Any alternative pair of strategies would have at least one of the players choosing a response that is not best for that player.
- There is no dynamic path to the Nash Equilibrium - the game is *static*.
- There may be multiple Nash Equilibria, or there may be no Nash Equilibrium (N.B. this comment applies particularly to *pure* strategies)

Consider the co-ordination to equilibrium - even when there are multiple equilibria, one of them may be more likely to occur than the others, especially if that equilibrium is Pareto optimal. Netherthreless, if players get 'stuck' into a 'bad' equilibrium, they may find it difficult to co-ordinate a change to a 'good' equilibrium. A puzzling question, then, is when no Nash Equilibrium are Pareto Optimal how do players co-ordinate onto one of them? For example, the following game has two identical Nash equilibria:

Player 2 / Player 1	$s_2^1 = L$	$s_2^2 = R$
$s_1^1 = L$	**1**, **1**	0, 0
$s_1^2 = R$	0, 0	**1**, **1**

Table 17: Game 4

Since the equilibria cannot be ranked in the Pareto sense, we cannot claim that one of them is more reasonable than the other - indeed, we cannot even claim that the outcome will be one of the two equilibria! There is nothing in a 'one-shot' game (i.e. a game played just once) preventing Player 1 playing *R* in an attempt to get to the (*R*, *R*) equilibrium, and Player 2 playing *L* in an attempt to get to the (*L*, *L*) equilibrium, thereby leading to the (*R*, *L*) outcome.

Appendix: Some Classifications

Strategy: A choice of action by a player, possibly contingent on other players' actions.

Zero-Sum v Non-Zero-Sum: In zero-sum games the payoffs sum to the same amount whatever strategies are chosen.

Cooperative v Non-Cooperative: In cooperative games the players can make binding agreements.

Common Knowledge v Asymmetric Information: In common knowledge games each player has the same information and each knows that the others have this information.

Static (Simultaneous) v Dynamic (Sequential): In static games each player makes his decision before finding out the other player(s) decision(s). In dynamic games strategic behavior is possible. (N.B. Chess is a dynamic game.)

One-Shot v Repeated: A one-shot game is played just once; in a repeated game the players face the same situation more than once. Repeated games can be finite or infinite.