EC220 - Mathematical Economics 1A
Notes[1]

Marco Del Vecchio[2]

Last revised on June 1, 2016

[1]Based on the offical lecture notes and *A Primer in Game Theory* by Prof. Robert Gibbons

[2]M.Del-Vecchio@Warwick.ac.uk

Contents

1 Game Theory **1**

1.1 Static Games of Complete Information 1

1.1.1 Normal Form Represenantion of Games 1

1.1.2 Intereted Elimination of Strictly Dominanted Strategies 1

1.1.3 Nash Equilibrium 2

1.1.4 Mixed Strategies 4

1.1.5 Existence of Nash Equilibrium 7

1.2 Dynamic Games of Complete Information 7

1.2.1 Dynamic Games of Complete and Perfect Information 7

1.2.2 Finitely Repeated Games 8

1.2.3 Infinitely Repeated Games 9

1.2.4 Repeated Games 10

1.2.5 Dynamic Games of Complete but Imperfect Information 11

1.2.6 Extensive-Form Representation of Games 11

1.2.7 Subgame-Perfect Nash Equilibrium 13

1.3 Static Games of Incomplete Information - Static Bayesian Games 14

1.3.1 Mixed Strategies 16

1.4 Dynamic Games of Incomplete Information - Dynamic Bayesian Games . . 16

1.4.1 Perfect Bayesian Equilibrium in Signalling Games 16

1.5 Evolutionary Game Theory 17

Chapter 1

Game Theory

1.1 Static Games of Complete Information

1.1.1 Normal Form Represenantion of Games

Definition ***Strategy Space*** *Let $\mathcal{S}_i$ be the set of strategies available to player i, and let s_i denote an arbitrary member of this set.*

Definition ***Payoff Function*** *Let u_i denote player i's payoff function such that $u_i(s_1, \dots, s_N)$ is the payoff of player i if the players choose the strategies $(s_1, \dots, s_N) \in \bigtimes_{i=1}^{N} \mathcal{S}_i$.*

Definition ***Normal-Form Representation*** *The normal-form representation of an n-player game specifies the players' strategy spaces $\mathcal{S}_1, \dots \mathcal{S}_N$ and their payoff functions $u_1 \dots, u_N$. We denote this game by $G = \langle \boldsymbol{S}, \boldsymbol{U} \rangle$ where $\boldsymbol{S} = \{\mathcal{S}_i\}_{i=1}^{N}$ and $\boldsymbol{U} = \{u_i\}_{i=1}^{N}$.*

Example: ***The Prisoner Dilemma***

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	-1,-1	-10,0
	Defect	0,-10	-5,-5

Here, N=2 since there are only two players. The strategy spaces for the two players are given by $\boldsymbol{S} = \{\mathcal{S}_1, \mathcal{S}_2\}$ where $\mathcal{S}_1 = \{\text{Cooperate}, \text{Defect}\}$ and $\mathcal{S}_2 = \{\text{Cooperate}, \text{Defect}\}$. Further, The payoffs of the two players when a particular pair of strategies is chosen are given in the appropriate cell of the bi-matrix. By convention, the payoff to the so-called row player (here, Player 1) is the first payoff given, followed by the payoff of the column player (here Player 2).

1.1.2 Intereted Elimination of Strictly Dominanted Strategies

Definition ***Strictly Dominated Strategy*** *in the normal-form game $G = \langle \boldsymbol{S}, \boldsymbol{U} \rangle$, let s_i' and s'' be feasible strategies for player i (i.e. s_i' and $s_i'' \in \mathcal{S}_i$). Then, we say that strategy s_i' is strictly dominated by strategy s_i'' if for each feasible combination of the other player's strategies, i's payoff from playing s_i' is strictly less than i's payoff from playing s_i'':*

$$u_i(s_1, \dots, s_{i-1}, s_i', s_{i+1}, \dots, s_N) < u_i(s_1, \dots, s_{i-1}, s_i'', s_{i+1}, \dots, s_N) \quad (SD)$$

for each $s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_N$ that can be constructed from the other player's strategy spaces $\mathcal{S}_1, \dots,$
$\mathcal{S}_{i-1}, \mathcal{S}_{i+1} \dots, \mathcal{S}_N$.

1.1.3 Nash Equilibrium

Definition ***Nash Equilibrium*** *in the n-player normal-form game $G = \langle \boldsymbol{S}, \boldsymbol{U} \rangle$, the strategies $(s_1^*, \dots, s_N^*)$ are a Nash equilibrium if, for each player i, s_i^* is player i's best response to the strategies specified for the $n-1$ other players, $(s_1^*, \dots, s_{i-1}^*, s_{i+1}^*, s_N^*)$:*

$$u_i(s_1^*, \dots, s_{i-1}^*, s_i^*, s_{i+1}^*, \dots, s_N^*) \geq u_i(s_1^*, \dots, s_{i-1}^*, s_i, s_{i+1}^*, \dots, s_N^*) \quad (NE)$$

for every feasible strategy $s_i \in \mathcal{S}_i$. That is to say, s_i^* solves:

$$\max_{s_i \in \mathcal{S}_i} u_i(s_1^*, \dots, s_{i-1}^*, s_i, s_{i+1}^*, \dots, s_N^*).$$

Proposition *In the n-player normal-form game $G = \langle \boldsymbol{S}, \boldsymbol{U} \rangle$, if iterated elimination of strictly dominated strategies eliminates all but the strategies $(s_1^*, \dots, s_N^*)$ then these strategies are the unique Nash equilibrium of the game.*

Proof. Suppose that iterated elimination of strictly dominated strategies eliminates all but the strategies $(s_1^*, \dots, s_N^*)$ yet, these strategies do not represent a Nash equilibrium. Then, there must exist some player i and some feasible strategy $s_i \in \mathcal{S}_i$ such that (NE) fails:

$$u_i(s_1^*, \dots, s_{i-1}^*, s_i^*, s_{i+1}^*, \dots, s_N^*) < u_i(s_1^*, \dots, s_{i-1}^*, s_i, s_{i+1}^*, \dots, s_N^*). \quad (1.1.1)$$

But, since $(s_1^*, \dots, s_N^*)$ are the only strategies surviving iterated elimination, there must exist $s_i' \in \mathcal{S}_\rangle$ which has strictly dominated s_i in the process

$$u_i(s_1, \dots, s_{i-1}, s_i, s_{i+1}, \dots, s_N) < u_i(s_1, \dots, s_{i-1}, s_i', s_{i+1}, \dots, s_N). \quad (1.1.2)$$

for each $(s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_N)$ that can be constructed from the strategies remaining in the other players' strategy spaces. Then, since the strategies $(s_1^*, \dots, s_{i-1}^*, s_{i+1}^*, \dots, s_N^*)$ survive the iterated elimination process by assumption, one implication of (1.1.2) is that

$$u_i(s_1^*, \dots, s_{i-1}^*, s_i, s_{i+1}^*, \dots, s_N) < u_i(s_1^*, \dots, s_{i-1}^*, s_i', s_{i+1}^*, \dots, s_N^*). \quad (1.1.3)$$

If s_i^* is the strategy that strictly dominates s_i (i.e. $s_i^* = s_i'$) then (1.1.3) contradicts (1.1.1), in which case the proof is complete (since the assumption that (NE) fails leads to a contradiction). If $s_i^* \neq s_i'$ then, there must exist some strategy $s_i'' \in \mathcal{S}_i$ which strictly dominates s_i' since s_i' does not survive the process. Thus, there would be inequalities similar to (1.1.1) and (1.1.3) with s_i' and s_i'' replacing s_i and s_i', respectively. Once again, if $s_i^* = s_i''$ then the proof is complete. Otherwise, two analogous inequalities can be constructed. Since, s_i^* is the only strategy in $\mathcal{S}_i$ which survives iterated elimination, repeating this argument in a finite game eventually completes the proof. □

Proposition *In the n-player game $G = \langle \boldsymbol{S}, \boldsymbol{U} \rangle$, if the strategies $(s_1^*, \dots, s_N^*)$ are a Nash equilibrium, then they survive iterated elimination of strictly dominated strategies.*

Proof. Suppose that the strategies $(s_1^*, \dots, s_N^*)$ are a Nash equilibrium of the normal-form game $G = \langle \boldsymbol{S}, \boldsymbol{U} \rangle$, but suppose s_i^* is the first of the strategies $(s_1^*, \dots, s_N^*)$ to be eliminated for being strictly dominated. Then, there must exist a strategy $s_i' \in \mathcal{S}_i$ which has not yet been eliminated such that

$$u_i(s_1, \dots, s_{i-1}, s_i^*, s_{i+1}, \dots, s_N) < u_i(s_1, \dots, s_{i-1}, s_i', s_{i+1}, \dots, s_N) \tag{1.1.4}$$

for each $(s_1, \dots s_{i-1}, s_{i+1}, \dots, s_N)$ that can constructed from the strategies that have not yet been eliminated form the other player's strategy spaces.

Now, since s_i^* was the first strategy to be eliminated, the other players' equilibrium strategies have not been eliminated, so one implication of (1.1.4) is

$$u_i(s_1^*, \dots, s_{i-1}^*, s_i^*, s_{i+1}^*, \dots, s_N^*) < u_i(s_1^*, \dots, s_{i-1}^*, s_i', s_{i+1}^*, \dots, s_N^*) \tag{1.1.5}$$

But (1.1.5) violates the definition of Nash equilibrium: s_i^* must the a best response to $(s_1^*, \dots, s_{i-1}^*, s_{i+1}^*, \dots, s_N^*)$, so there cannot exist a strategy $s_i' \in \mathcal{S}_i$ that strictly dominates s_i^*. This contradiction completes the proof. □

Example: Cournot Model of Duopoly

Let q_1 and q_2 denote the quantities (of homogeneous product) produced by firms 1 and 2, respectively. Let $P(Q) = a - Q$ be the market-cleaning price when the aggregate quantity on the market is $Q = q_1 + q_2$. Assume that the total cost to firm i of producing quantity q_i is $C(q_1) = cq_i$.

Let us now translate this informal statement into a normal-form representation of the game by individuating: (a) the players; (b) the strategy spaces; (c) the payoffs. There are two players in any duopoly game - the two firms. In the Cournot model, the strategies available to each firm are the different quantities it might produce. Naturally, firm's output cannot be negative, thus, each firm's strategy space can be represented as $\mathcal{S}_i = [0, \infty)$. We assume that the firm's payoffs are simply their profits. Thus, the payoff $u_i(s_i, s_j)$ in a general two-player game in normal form can be written here as

$$\pi_i(q_i, q_j) = u_i(q_i, q_j) = q_i P(Q) - C(q_i) = q_i(a - (q_i + q_j) - c).$$

We are ready to compute the game's Nash equilibrium. We have already saw that in a two-player game in normal form, the strategy pair (s_i^*, s_j^*) is a Nash equilibrium if for each player i,

$$u_i(s_i^*, s_j^*) \geq u_i(s_i, s_j^*)$$

for every possible strategy $s_i \in \mathcal{S}_i$. Equivalently, for each player i, s_i^* must solve the problem

$$\max_{s_i \in \mathcal{S}_i} u_i(s_i, s_j^*)$$

In the Cournot duopoly model, the analogous statement is that the quantity pair (q_i^*, q_j^*) is a Nash equilibrium if, for each firm i, q_i^* solves

$$\max_{q_i \in [0,\infty)} \pi_i(q_i, q_j^*) = \max_{q_i \in [0,\infty)} q_i(a - (q_i + q_j^*) - c)$$

Solving the optimisation problems by deriving the first order condition yields

$$\frac{\partial \pi_i}{\partial q_i} = 0 \implies q_i^* = \frac{1}{2}(a - q_j^* - c)$$

Thus, for the quantity pair (q_1^*, q_2^*) to be a Nash equilibrium, the equations

$$q_1^* = \frac{1}{2}(a - q_2^* - c)$$

and

$$q_2^* = \frac{1}{2}(a - q_1^* - c)$$

must be satisfied simultaneously. Solving this pair of equation yields

$$q_1^* = q_2^* = \frac{a-c}{3}.$$

Example: Bertrand Model of Duopoly

We now consider the scenario where firms choose prices, rather than quantities. Let us consider the case of differentiated products. If firms 1 and 2 will choose prices p_1 and p_2, respectively, the quantity demanded will be from firm i will be

$$q_i(p_i, p_j) = a - p_i + bp_j$$

Where $b > 0$ reflects the extent to which firm i's product is a substitute for firm j's product.Assume that the total cost to firm i of producing quantity q_i is $C(q_1) = cq_i$.

Let us now translate this informal statement into a normal-form representation of the game by individuating: (a) the players; (b) the strategy spaces; (c) the payoffs. There are two players in any duopoly game - the two firms. There are two players. The strategies available to each firm are the different prices it might charge. Naturally, prices cannot be negative, thus, each firm's strategy space can be represented as $\mathcal{S}_i = [0, \infty)$. We assume that the firm's payoffs are simply their profits. Thus, the profit to firm i when it chooses the price p_i and its rival chooses the price p_j is

$$\pi_i(p_i, p_j) = p_i q_i - C(q_i) = (p_i - c)(a - p_i - bp_j)$$

Thus, the price pair (p_i^*, p_j^*) is a Nash equilibrium if, for each firm, it solves

$$\max_{s_i \in \mathcal{S}_i} \pi_i(p_i, p_j^*) = \max_{s_i \in \mathcal{S}_i} (p_i - c)(a - p_i - bp_j^*).$$

Solving the optimisation problems by deriving the first order condition yields

$$\frac{\partial \pi_i}{\partial p_i} = 0 \implies p_i^* = \frac{1}{2}(a + bp_j^* + c)$$

Therefore, if the price pair is a Nash equilibrium, the firms' price choices must satisfy

$$p_1^* = \frac{1}{2}(a + bp_2^* + c)$$

and

$$p_2^* = \frac{1}{2}(a + bp_2^* + c)$$

Solving this pair of equation yields

$$p_1^* = p_2^* = \frac{a+c}{2-b}.$$

1.1.4 Mixed Strategies

Definition ***Pure Strategy*** *If $\mathcal{S}_i$ identifies player i's set of strategies, we refer to the members of $\mathcal{S}_i$ (i.e. s_i) as player i's pure strategies.*

Definition ***Mixed Strategy*** *In the normal-form game $G = \langle \boldsymbol{S}, \boldsymbol{U} \rangle$, suppose $\mathcal{S}_i = \{s_{i1}, \ldots s_{iK}\}$, Then a mixed strategy for player i is a probability distribution $p_i = (p_{i1}, \ldots, p_{iK})$, where $0 \leq p_{ik} \leq 1$ for $k \in \{1, 2, \ldots, K\}$, $\sum_{k=1}^{K} p_{ik} = 1$, and p_{ik} represent the probability that player i will play strategy s_{ik}.*

We shall now provide the extended definition of Nash equilibrium by restricting our attention to the two player case. Let the players 1 and 2 have strategy spaces $\mathcal{S}_1$ and $\mathcal{S}_2$, respectively. Further, let J and K denote the number of pure strategies in $\mathcal{S}_1$ and $\mathcal{S}_2$, respectively (i.e $|\mathcal{S}_1| = J$ and $|\mathcal{S}_2| = K$). We shall write $\mathcal{S}_1 = \{s_{11}, \ldots, s_{1J}\}$ and $\mathcal{S}_2 = \{s_{21}, \ldots, s_{2J}\}$, and we shall use s_{1j} and s_{2k} to denote an arbitrary pure strategy from $\mathcal{S}_1$ and $\mathcal{S}_2$, respectively.

If player 1 believes that player 2 will play the strategies $(s_{21}, \ldots, s_{2K})$ with probabilities $p_2 = (p_{21}, \ldots, p_{2K})$ then, player 1's expected payoff from playing the pure strategy s_{1j} is

$$\sum_{k=1}^{K} p_{2k} \cdot u_1(s_{1j}, s_{2k})$$

and player 1's expected payoff from playing the mixed strategy
$p_1 = (p_{1j}, \ldots, p_{1J})$ is

$$v_1(p_1, p_2) = \sum_{j=1}^{J} \sum_{k=1}^{K} p_{1j} \cdot p_{2k} \cdot u_1(s_{1j}, s_{2k})$$

where $p_{1j} \cdot p_{2k}$ is the probability that player 1 plays s_{1j} and player 2 plays s_{2k}. Hence, for the mixed strategy $p_1 = (p_{11}, \ldots, p_{1J})$ to be a best response for player 1 to 2's mixed strategy p_2, it must be that $p_{1j} > 0$ (i.e. s_{1j} has some non-zero probability of being played) only if

$$\sum_{k=1}^{K} p_{2k} \cdot u_1(s_{1j}, s_{2k}) \geq \sum_{k=1}^{K} p_{2k} \cdot u_1(s_{1j'}, s_{2k})$$

for every $s_{1j'} \in \mathcal{S}_1$. That is, for a mixed strategy to be a best response to to p_2 it must place positive probability on a given pure strategy only if the pure strategy itself is a best response to to p_2. Conversely, if player 1 has several pure strategies that are best responses to p_2 then, any mixed strategy that places all its probability on some of all of there pure strategy best responses (and zero probability on all other pure strategies) is also a best response for player 1 to p_2. Symmetrically, if player 2 believes that player 1 will play the strategies $(s_{1j}, \ldots, s_{1J})$ with probabilities $p_1 = (p_{1j}, \ldots, p_{1J})$ then, player 2's expected payoff from playing the pure strategy s_{2k} is

$$\sum_{j=1}^{J} p_{1j} \cdot u_2(s_{1j}, s_{2k})$$

and player 2's expected payoff from playing the mixed strategy
$p_2 = (p_{21}, \ldots, p_{2K})$ is

$$v_2(p_1, p_2) = \sum_{k=1}^{K} \sum_{j=1}^{J} p_{2k} \cdot p_{1j} \cdot u_2(s_{1j}, s_{2k})$$

Given $v_1(p_1, p_2)$ and $v_2(p_1, p_2)$ we can restate the requirement of Nash equilibrium to include mixed strategies (in a two-player game).

Definition ***Nash Equilibrium*** *In the two-player normal-form game* $G = \langle \boldsymbol{S}, \boldsymbol{U} \rangle$ *where* $\boldsymbol{S} = \{\mathcal{S}_1, \mathcal{S}_2\}$ *and* $\boldsymbol{U} = \{u_1, u_2\}$*, the mixed strategies* (p_1^*, p_2^*) *is a best response to to the other player's mixed strategy:*

$$v_1(p_1^*, p_2^*) \geq v_2(p_1, p_2^*)$$

for every probability distribution p_1 *over* $\mathcal{S}_1$*, and*

$$v_1(p_1^*, p_2^*) \geq v_2(p_1^*, p_2)$$

for every probability distribution over $\mathcal{S}_2$*.*

Remark: Since any pure strategy can be represented as the mixed strategy that places zero probability of all of the player's other pure strategies, this extended definition absorbs the earlier one.

Example: Game of Chicken

Let us apply this definition to the game of chicken. In normal form, the game is described by the bi-matrix below.

		Player 2	
		Tough	Swerve
Player 1	Tough	-8,-8	2,-3
	Swerve	-3,2	0,0

We start by identifying all pure strategies Nash equilibria. player 1 best response to player 2 playing Tough and Swerve is to play Swerve (as $-3 > -8$) and Tough (as $2 > 0$), respectively. Symmetrically, player 2 best response to player 1 playing Tough and Swerve is to play Swerve (as $-3 > -8$) and Tough (as $2 > 0$), respectively. Therefore, we identify two Nash equilibria (Tough,Swerve) and (Swerve,Tough). We now turn to mixed strategy. Suppose that player 1 believes that player 2 will play Tough and Swerve with probability q and $1-q$ respectively. Given this beliefs, player 1's expected payoffs are $-8q+2(1-q) = -10q+2$ from playing Tough and $-3q+0(1-q) = -3q$ from playing Swerve. Then, player 1 utility when playing Swerve will be $-8q$. Since $2-10q > -3q$ if ans only if $q < 2/7$, player 1's best pure strategy response is to play Tough if $q < 2/7$ and Swerve if $q > 2/7$, and player 1 is indifferent between the two strategies if $q = 2/7$.

Let, $(r, 1-r)$ denote the mixed strategy in which player 1 plays Tough and Swerve with probability r and $1-r$, respectively. For each value of q between 0 and 1, we can compute the value of r which makes $(r, 1-r)$ a best response for player 1 to $(q, 1-q)$ by player 2. Player 1's expected payoff from playing $(r, 1-r)$ when 2 plays $(q, 1-q)$ is

$$r(2-10q) + (1-r)(-3q) = r(2-7q) - 3q$$

which is increasing in r when $q < 2/7$ and decreasing in r if $q > 2/7$. Therefore, player 1's expected payoff will be maximised by playing Tough with probability 1 (i.e. $r = 1$) when $q < 2/7$ and Swerve with probability 1 (i.e. $r = 0$) when $q > 2/7$. And, as noted before, player 1 is indifferent between two strategies when $q = 2/7$.

If we now perform the same calculations for player 2's best response, we would find that player 2's expected payoff will be maximised by playing Tough with probability 1 (i.e. $q = 1$) when $r < 2/7$ and Swerve with probability 1 (i.e. $q = 0$) when $r > 2/7$. And, as noted before, player 2 is indifferent between two strategies when $q = 2/7$.

This information can be summarised in the following graph.

Now, the Nash equilibria are represented by the intersection points of the two players' best responses. That is, on top of the two pure strategies Nash equilibria, we have one mixed strategy Nash equilibrium: when player i plays (2/7, 3/7) then (2/7, 3/7) is a best reposne for player j.

1.1.5 Existence of Nash Equilibrium

Theorem *(Nash 1950): In the n-player normal-form game* $G = \langle \boldsymbol{S}, \boldsymbol{U} \rangle$ *where* $\boldsymbol{S} = \{\mathcal{S}_1, \ldots, \mathcal{S}_N\}$ *and* $\boldsymbol{U} = \{u_1, \ldots, u_N\}$, *if* N *is finite and* $|\mathcal{S}_i|$ *is finite for every* i *then there exist at least one Nash equilibrium, possibly involving mixed strategies.*

1.2 Dynamic Games of Complete Information

1.2.1 Dynamic Games of Complete and Perfect Information

Backwards Induction

Let us define a simple class of games of complete and perfect information as: the set of problems which fit the following descriptions:

1. Player 1 chooses an action a_1 from the feasible set $\mathcal{A}_1$.
2. Player 2 observes a_1 and chooses an action a_2 form the feasible set $\mathcal{A}_2$.
3. Payoffs are $u_1(a_1, a_2)$ and $u_2(a_1, a_2)$.

Further, the features od a dynamic game of complete and perfect information are that (i) the moves occur in sequence, (ii) all previous moves are observed before the next move is chosen, and (iii) the players' payoffs from each feasible combination of moves are common knowledge.

We solve a game from this class by backward induction, as follows. When player 2 gets to move at the second stage of the game, they will face the following maximisation problem, give the action a_1 previously chosen by player 1:

$$\max_{a_2 \in \mathcal{A}_2} u_2(a_1, a_2).$$

Assume that for each $a_1 \in \mathcal{A}_1$, player 2's optimisation problem has a unique solution denoted by $R_2(a_1)$. This, is player 2's reaction to player'1 action. Moreover, since player 1 can solve 2's optimisation problem as well 2 can (remember this is a game of complete and perfect information), player 1 should anticipate player 2's reaction to each action a_1 that 1 might take, so 1's problem at the first stage is

$$\max_{a_1 \in \mathcal{A}_1} u_2(a_1, R_2(a_1)).$$

Assume that this problem for player 1 also has a unique solution, denoted by a_1^*. We then call $(a_1^*, R_2(a_1^*))$ the *backwards-induction outcome* of this game.

Remark: The backwards-induction outcome does not involve non-credible threats.

Example: Stackelberg Model of Duopoly
Consider the dynamic model of duopoly in which a dominant (or leader firm) moves first and a subordinate (or follower) firm moves second.

The timing of the game is as follows: (1) firm 1 chooses a quantity $q_1 \geq 0$; (2) firms

2 observes q_1 and then chooses a quantity $q_2 \geq 0$; (3) the payoff to firm i is given by the profit function

$$\pi_i(q_i, q_j) = q_i(P(Q) - c),$$

where $P(Q) = a - Q$ is the market cleaning price when the aggregate quantity on the market is $Q = q_1 + q_2$, and c is the constant marginal cost of production. To solve the backwards-induction outcome of this game, we first compute firm 2's reaction to an arbitrary quantity q_1 by firm 1. Hence, for $R_2(q_1)$ to be firm 2's reaction function it must solve

$$max_{q_2 \in [0,\infty)} \pi_2(q_1, q_2) = max_{q_2 \in [0,\infty)} q_2(a - q_1 - q_2 - c),$$

which yields

$$R_2(q_1) = \frac{1}{2}(a - q_1 - c)$$

Now, since firm 1 can solve firm 2's optimisation problem as well as firm 2 can, it should choose q_1 knowing that this decision will be met with the reaction $R_2(q_1)$. Thus, firm 1's optimisation problem in the first stage is

$$max_{q_1 \in [0,\infty)} \pi_1(q_1, R_2(q_1)) = max_{q_1 \in [0,\infty)} q_1(a - q_1 - \frac{1}{2}(a - q_1 - c) - c),$$

which yields

$$q_1^* = \frac{1}{2}(a - c) \quad \text{and} \quad R_2(q_1^*) = \frac{1}{4}(a - c)$$

as the backwards-solution outcome of the Stackelberg model of duopoly.

1.2.2 Finitely Repeated Games

Let $G = \langle \boldsymbol{A}, \boldsymbol{U} \rangle$ where $\boldsymbol{A} = \{\mathcal{A}_i\}_{i=1}^N$ and $\boldsymbol{U} = \{u_i\}_{i=1}^N$ denote a static game of complete information in which players 1 through N simultaneously choose actions a_1 through a_N form the action spaces $\mathcal{A}_1$ through $\mathcal{A}_N$, respectively, na payoffs are $u_1(a_1, \ldots a_N)$ through $u_n(a_1, \ldots a_N)$. The game G will be called the stage game of the repeated game.

Definition ***Finitely Repeated Game*** *Given a stage game G, let $G(T)$ denote the finitely repeated game in which G is played T times, with the outcomes of all preceding games observed before the next play begins. The payoffs of $G(T)$ sare simply the sum of the payoffs form the T stage games.*

Proposition *If the stage game G has a unique Nash equilibrium then, for any finite T, the repeated game $G(T)$ has a unique ee-perfect outcome: the Nash equilibrium of G is played in every stage.*

Example: Twice Repeated Game

Consider a repeated game in which the stage game in the following table is played in each of two periods and there is no discounting.

		Player 1		
		L	M	R
	U	8,8	0,9	0,0
Player 2	C	9,0	0,0	3,1
	D	0,0	1,3	3,3

We want to: Identify the Nash equilibrium (equilibria) of the stage game, and fully describe a subgame perfect equilibrium in which the players select (U,L) in the first period.

1.

The Nash equilibria of the game are (C,R), (D,M), (D,R).
Now, suppose that player 1 plays U in the first period but player 2 defect by playing either M or R then the payoffs in the second stage are:

If player 2 plays M

	L	M	R
U	8,17	0,18	0,9
C	9,9	0,9	3,10
D	0,9	1,12	3,12

If player 2 plays R

	L	M	R
U	8,8	0,9	0,0
C	9,0	0,0	3,1
D	0,0	1,3	3,3

In this case, player 1 will punish player 2 for having played either M or R by playing C in the second stage. Players 2 will then best respond by playing R. Thus coordinating on (C,R).

Now, suppose that player 2 plays L in the first period but player 1 defect by playing either C or D then the payoffs in the second stage are:

If player 1 plays C

	L	M	R
U	17,8	9,9	9,0
C	18,0	9,0	12,1
D	9,0	10,3	12,3

If player 1 plays D

	L	M	R
U	8,8	0,9	0,0
C	9,0	0,0	3,1
D	0,0	1,3	3,3

In this case, player 2 will punish player 1 for having played either C or D by playing M in the second stage. Players 1 will then best respond by playing D. Thus coordinating on (D,M).

Otherwise, if (U,L) is played in the first period, players will coordinate (D,R) in the second period.

1.2.3 Infinitely Repeated Games

Definition ***Present Value*** *Given the discount factor δ, the present value of the infinite sequence of payoffs $u_1, u_2, u_3 \dots$ is*

$$u_1 + \delta u_2 + \delta^2 u_3 + \dots = \sum_{t=1}^{\infty} \delta^{t-1} u_t$$

Remark: $a + a\delta + a\delta^2 + \dots = \sum_{k=0}^{\infty} a\delta^k = 1/(1-\delta)$. Further, if the geometric series does not start from zero but from n then $a\delta^n + a\delta^{n+1} + \dots = \delta^n(a + a\delta + a\delta^2 + \dots) = \delta^n \sum_{k=0}^{\infty} a\delta^k = \delta^n/(1-\delta)$.

Definition ***Infinetely Repeated Game*** *Given a stage game G, let $G(\infty, \delta)$ denote the infinitely repeated game in which G is repeated forever and the players share the discount factor δ. For each t, the outcome of the $t-1$ preceding plays of the stage game are observed before the t^{th} stage begins. Each player's payoff in $G(\infty, \delta)$ is the present value of the player's payoff from the infinite sequence of stage games.*

1.2.4 Repeated Games

Definition ***Stratety in Repeated Games*** *In the finitely repeated game $G(T)$ or the infinitely repeated game $G(\infty, \delta)$, a player's strategy specifies the action the player will take in each stage, for each possible history of play through the previous stage.*

Definition ***Subgame in Repeated Games*** *In the finitely repeated game $G(T)$, a subgame beginning at stage $t+1$ is the repeated game in which G is played $T-1$ times, denoted $G(T-1)$. There are many subgames that begin at stage $t+1$, one for each of the possible histories of play through stage t. In the infinitely repeated game $G(\infty, \delta)$, each subgame beginning at stage $t+1$ is identical to the original game $G(\infty, \delta)$. As in the finite-horizon case, there are as many subgames beginning at stage $t+1$ of $G(\infty, \delta)$ as there are possible histories of play through stage t.*

Definition ***Subgame-Perfection in Repeated Games (Selten, 1965)*** *A Nash equilibrium is subgame-perfect if the player's strategies constitute a Nash equilibrium in every subgame.*

Example: Infinitely Repeated Prisoner Dilemma and Grim Trigger Strategies

Consider an infinitely repeated game $G(\infty, \delta)$ in which the following version of the prisoner dilemma is played at every stage.

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	1,1	-1,2
	Defect	2,-1	0,0

Further, suppose that the both players follow the grim trigger strategy: play Cooperate in the stage game t if and only if the other player has played Cooperate in all previous $t-1$ stages. If the other player plays Defect at any stage then punish him by paying Defect forever, regardless of what they play. Let us find the appropriate discount factor δ so that a cooperative equilibrium can be sustained.

The payoff form infinite cooperation is

$$1 + 1\delta + 1\delta^2 + \cdots = \frac{1}{1-\delta}.$$

Then, suppose that one player deviates to Defect. Since the other player is following the grim trigger strategy, he will play Cooperate (remember here the players move simultaneously so they are not aware of what the other player has played) when the other player has just played Defect and Defect forever after. The payoff from defecting is

$$2 + 0\delta + 0\delta^2 + \ldots \frac{0}{1-\delta} = 2.$$

Thus, cooperation can be induced as long as

$$\frac{1}{1-\delta} \geq 2 \implies \delta \geq \frac{1}{2}.$$

1.2.5 Dynamic Games of Complete but Imperfect Information

Let us define a simple class of games of complete but imperfect information as: the set of problems which fit in the following description:

1. Players 1 and 2 simultaneously choose actions a_1 and a_2 from the feasible sets $\mathcal{A}_1$ and $\mathcal{A}_2$, respectively.

2. Players 3 and 4 observe the outcome of the first stage (a_1, a_2) and then simultaneously choose actions a_3 and a_4 from $\mathcal{A}_3$ and $\mathcal{A}_4$, respectively.

3. Payoffs are $u_1(a_1, a_2, a_3, a_4)$ for $i = 1, 2, 3, 4$.

Further, the features od a dynamic game of complete but imperfect information are that (i) the moves occur either in sequence or simultaneously, (ii) all previous moves are observed before the next move is chosen, and (iii) the players' payoffs from each feasible combination of moves are common knowledge.

We solve a game from this class by using an approach similar to backwards-induction, but this time, the fist step in working backwards from the end of the game involves solving a simultaneous-move game between players 3 and 4 in stage two, given the outcome from stage one -rather than solving a single-player optimisation problem.

If player 1 and 2 anticipate that the second-stage behaviour of players 3 and 4 will be given by $(a_3^*(a_1, a_2), a_4^*(a_1, a_2))$, then the first stage interaction between players 1 and 2 consist to the following simultaneous-move game:

1. Players 1 and 2 simultaneously choose actions a_1 and a_2 from the feasible sets $\mathcal{A}_1$ and $\mathcal{A}_2$, respectively.

2. Payoffs are $u_i(a_1, a_2, a_3^*(a_1, a_2), a_4^*(a_1, a_2))$ for $i = 1, 2$

Suppose that (a_1^*, a_2^*) is the unique Nash equilibrium of this simultaneous-move game. We shall call $(a_1^*, a_2^*, a_3^*(a_1^*, a_2^*), a_4^*(a_1^*, a_2^*))$ the *subgame-perfect outcome* of this two-stage game.

1.2.6 Extensive-Form Representation of Games

Definition ***Extensive-Form Representation*** *The extensive-form representation of a game specifies: (1) the players in the game, (2a) when each player moves, (2b) what each player can do at each of their opportunities to move, (2c) what each player knows at each of their opportunities to move, and (3) the payoffs received by each player for each combination of moves that could be chosen by the players.*

Definition ***Strategy*** *A strategy for a player is a complete plan of action - it specifies a feasible action for the player in every contingency in which the player might be called to act.*

Example: Extensive-Form Representation - Game of Complete and Perfect Information

As an example fo a game in extensive form, consider the following two-stage game of complete and perfect information.

1. Player 1 chooses an action a_1 from the feasible set $\mathcal{A}_1 = \{L, R\}$

2. Player 2 observes a_1 and then chooses an action a_2 from the set $\mathcal{A}_2 = \{L', R'\}$

3. Payoffs are $u_1(a_1, a_2)$ and $u_2(a_1, a_2)$ as shown in the game tree below.

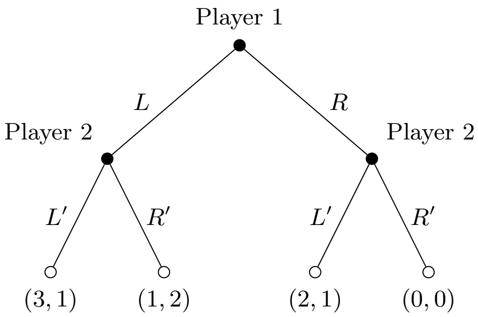

Let us now write down the strategy sets for the two players.

Strategy space for player 1

Strategy 1: play L.

Strategy 2: play R.

Strategy space for player 2

Strategy 1: if player 1 plays L then play L', if player 1 plays R then play L', denoted by (L', L').

Strategy 2: if player 1 plays L then play L', if player 1 plays R then play R', denoted by (L', R').

Strategy 3: if player 1 plays L then play R', if player 1 plays R then play L', denoted by (R', L').

Strategy 4: if player 1 plays L then play R', if player 1 plays R then play R', denoted by (R', R').

Given the strategy spaces for the two players, we can now derive the normal-form representation of this game.

		Player1			
		(L', L')	(L', R')	(R', L')	(R', R')
Player 2	L	3,1	3,1	1,2	1,2
	R	2,1	0,0	2,1	0,0

From both representation, we see can there exist only one pure strategy Nash equilibrium: (R, L').

Definition ***Information Set*** *An information set for a player is a collection of decision nodes satisfying: (i) the player has the move at every node in the information set, and (ii) when the play of the game reaches a node in the information set, the player with the move does not know which node in the information set has (or has not) been reached.*

Remark: In an extensive-form game, we shall indicate that a collection of decision nodes constitutes an information set by connecting the nodes by a dotted line.

Example: Extensive-Form Representation - Game of Complete but Imperfect Information
Here is an example of a simultaneous-move game.

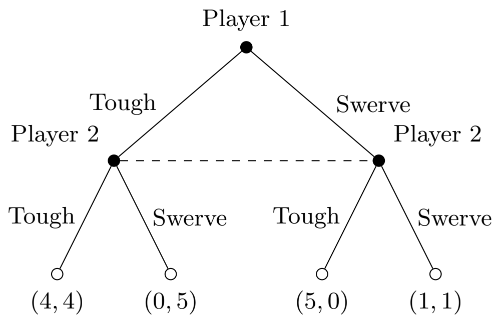

1.2.7 Subgame-Perfect Nash Equilibrium

Definition ***Subgame-Perfect Nash Equilibrium*** *A Nash equilibrium is subgame perfect if the player's strategy constitute a Nash equilibrium in every subgame.*

Definition ***Subgame-Perfect N.E. in Two-Stages Games of Complete and Perfect Information*** *In the two-stage game of complete and perfect information, the backwards induction outcome is* $(a_1^*, R_2(a_1^*))$ *but the subgame-perfect Nash equilibrium is* $(a_1^*, R_2(a_1))$.

Definition ***Subgame-Perfect N.E. in Two-Stages Games of Complete but Imperfect Information*** *In the two-stage game of complete but imperfect information, the subgame-perfect outcome is* $(a_1^*, a_2^*, a_3^*(a_1^*, a_2^*),\ a_4^*(a_1^*, a_2^*))$ *but the subgame-perfect Nash equilibrium is* $(a_1^*, a_2^*, a_3^*(a_1, a_2),\ a_4^*(a_1, a_2))$.

The above mentioned definition should make sense if the reader has well understood the difference between an action and a strategy. In fact, the reason why, for instance, $(a_1^*, R_2(a_1^*))$ does not represent the subgame perfect Nash equilibrium is because $R_2(a_1^*)$ is the action which player 2 will perform after player 1 played a_1^*. On the contrary, the best-response function $R_2(a_1)$ with $a_i \in \mathcal{A}_1$ is a strategy from player 2 because it specifies a feasible action for player 2 in every contingency in which the player might be called to act.

Example: Entry Deterrence in a Game of Complete and Perfect Information
Consider he following game:

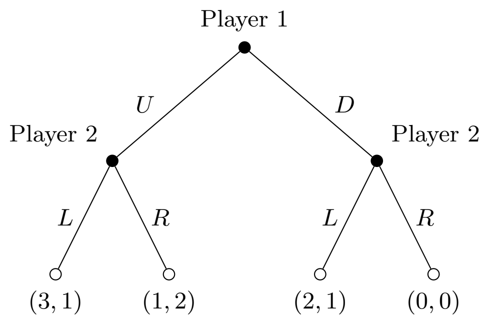

Is is evident that the unique subgame perfect Nash equilibrium in pure strategies is (D, RL) that is, player 1 plays D and player 2 responds by playing R if player 1 has played U and by plying L if player 1 has played D. This follows from the back-wards induction approach.

1.3 Static Games of Incomplete Information - Static Bayesian Games

Definition ***Normal-Form Representation of Static Bayesian Games*** *The normal-form representation of an n-player static Bayesian game specifies the player action spaces* $\boldsymbol{A} = \{A_i\}_{i=1}^N$*, their type spaces* $\boldsymbol{T}\{T_i\}_{i=1}^N$*, their beliefs* $\boldsymbol{P} = \{p_i\}_{i=1}^N$*, and their payoffs functions* $\boldsymbol{U} = \{u_i\}_{i=1}^N$*. Player i's type,* t_i*, is privately known to player i, determines player i's payoff function,* $u_i(a_1, \ldots, a_N s; t_i)$*, and is a member of the set of possible types,* $\mathcal{T}_i$*. Player i's belief* $p_i(t_{\neg i}|t_i)$ *describes i's uncertainty about the* $n-1$ *other players' possible types,* $t_{\neg i}$*, given i's own type. We denote this game by* $G = \langle \boldsymbol{A}, \boldsymbol{T}, \boldsymbol{P}, \boldsymbol{U} \rangle$

Definition ***Bayesian Nash Equilibrium*** *in the static Bayesian game* $G = \langle \boldsymbol{A}, \boldsymbol{T}, \boldsymbol{P}, \boldsymbol{U} \rangle$*, the strategies* $(s_1^*, \ldots, s_N^*)$ *are a (pure strategy) Bayesian Nash equilibrium if for each player i and for each of i's types* t_i *in* $(T)_i$*,* $s_i^*(t_i)$ *solves*

$$\max_{a_i \in \boldsymbol{A}_i} \sum_{t_{\neg i} \in \boldsymbol{T}_{\neg i}} u_i(s_1^*(t_1), \ldots, s_{i-1}^*(t_{i-1}), a_i, s_{i+1}^*(t_{i+1}), \ldots, s_N^*(t_N); t) p_i(t_{\neg i}|t_i)$$

That is, no player wants to change their strategy, even if the change involves only one action by one type.

Example: Battle of the Sexes with Incomplete and Imperfect Information
Consider the following version of the Battle of the Sexes game, where a man and a woman simultaneously choose whether to go to football or to the opera. Suppose that each person in the couple can be in love, or not in love. The probability that they are both in love is 0.1, the probability that they are both not in love is 0.1, the probability that she is in love and he is not is 0.4, the same as the probability that he is in love and she is not. The payoffs are as follows. For the woman: If she is in love, her payoff is 3 if she is at the opera with him; 2 if she is at football with him; and 0 if she is without him (wherever she is). If she is not in love, her payoff is 3 if she is at the opera (with or without him); 0 if she is at football without him; and 1 if she is at football with him. For the man: If he is in love, his payoff is 3 if he is at football with her; 2 if he is at the opera with her; and 0 if he is without her (wherever he is). If he is not in love, his payoff is 3 if he is at football (with or without her); 0 if he is at the opera without her; and 1 if he is at the opera with her.

			Husband			
		t_W	Loving		Not Loving	
	t_H	u_H, u_W	Opera	Football	Opera	Football
Wife	Loving	Opera	3,2	0,0	3,-1	0,3
		Football	0,0	2,3	0,0	2,3
	Not Loving	Opera	3,2	3,0	3,-1	3,3
		Football	0,0	-1,3	0,0	-1,3

We seek to determine whether the following pair of strategies is a Bayes-Nash equilibrium: if she is in love she goes to football, and if she is not in love she goes to opera, that is $(t_W = L \to F, t_W = NL \to O)$; while if he is in love he goes to opera, and if he is not in love he goes to football, that is $(t_H = L \to O, t_H = NL \to F)$.

Firstly, we need to calculate the following conditional probabilities:

$$P(t_H = L|t_W = L) = \frac{P(t_H = L, t_W = L)}{P(t_W = L)}$$
$$= \frac{P(t_H = L, t_W = L)}{P(t_W = L, t_H = L) + P(t_W = L, t_H = NL)} = 0.2$$
$$P(t_H = L|t_W = NL) = \frac{P(t_H = L, t_W = NL)}{P(t_W = NL)}$$
$$= \frac{P(t_H = L, t_W = L)}{P(t_W = NL, t_H = L) + P(t_W = NL, t_H = NL)} = 0.8$$
$$P(t_W = L|t_H = NL) = \frac{P(t_H = NL, t_W = L)}{P(t_H = NL)}$$
$$= \frac{P(t_H = NL, t_W = L)}{P(t_H = NL, t_W = L) + P(t_H = NL, t_W = NL)} = 0.8$$
$$P(t_H = NL|t_W = NL) = \frac{P(t_H = NL, t_W = NL)}{P(t_W = NL)}$$
$$= \frac{P(t_H = NL, t_W = NL)}{P(t_W = NL, t_H = L) + P(t_W = NL, t_H = NL)} = 0.2$$

Hence, this game can be described as $G = \langle \boldsymbol{A}, \boldsymbol{T}, \boldsymbol{P}, \boldsymbol{U} \rangle$ where $\boldsymbol{A} = \{\text{Opera}, \text{Football}\}$, $\boldsymbol{T} = \{\text{Loving}, \text{Not Loving}\}$, and $\boldsymbol{P}$ and $\boldsymbol{U}$ are as calculated above.

We can use this information to construct a table of payoffs.

Suppose that the Husband plays the proposed strategy, that is $(t_H = L \to O, t_H = NL \to F)$ then, the Wife payoffs are:

s_W	$t_H = L$	$t_H = NL$	$u_W(s_W, s_H, t_w = L)$
$(t_W = L \to O, t_W = NL \to O)$	3	0	$0.2 \times 3 + 0.8 \times 0 = 0.6$
$(t_W = L \to O, t_W = NL \to F)$	3	0	$0.2 \times 3 + 0.8 \times 0 = 0.6$
$(t_W = L \to F, t_W = NL \to O)$	0	2	$0.2 \times 0 + 0.8 \times 2 = 1.6$
$(t_W = L \to F, t_W = NL \to F)$	0	2	$0.2 \times 0 + 0.8 \times 2 = 1.6$

if the Wife is Loving and

s_W	$t_H = L$	$t_H = NL$	$u_W(s_W, s_H, t_w = NL)$
$(t_W = L \to O, t_W = NL \to O)$	3	3	$0.2 \times 3 + 0.8 \times 3 = 3$
$(t_W = L \to O, t_W = NL \to F)$	0	-1	$0.2 \times 0 + 0.8 \times -1 = -0.8$
$(t_W = L \to F, t_W = NL \to O)$	3	3	$0.2 \times 3 + 0.8 \times 3 = 3$
$(t_W = L \to F, t_W = NL \to F)$	0	-1	$0.2 \times 0 + 0.8 \times -1 = -0.8$

if the Wife is Not Loving. Now, if the wife is to best respond then she will play the strategy which maximises her utility that is she will play Football when she is Loving and Opera when she is Not Loving. Since the play is symmetric, we do not need to do the same for the husband and we are ready to conclude that the proposed equilibrium is indeed a Bayes-Nash equilibrium.

1.3.1 Mixed Strategies

Having introduced static Bayesian games, we can now give a more precise statement of the idea behind mixed strategies: 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 little bit of incomplete information. This, the crucial feature of a mixed strategy Nash equilibrium is not that player i chooses a strategy randomly, but rather that player i is uncertain about player j's choice.

1.4 Dynamic Games of Incomplete Information - Dynamic Bayesian Games

Definition ***Information Set On/Off the Equilibrium Path*** *For a given equilibrium in a given extensive-form game, an information set is on the equilibrium path if it will be reached with positive probability if the game is played according to the equilibrium strategies, and if off the equilibrium path if it is certain not to be reached if the game is played according to the equilibrium strategies (where equilibrium can mean Nash, subgame-perfect, Bayesian or perfect Bayesian equilibrium)*

Definition ***Perfect Bayesian Equilibrium*** *A perfect Bayesian equilibrium consists of strategies and beliefs satisfying:*

Requirement 1 *At each information set, the player with the move must have a belief about which node in the information set has been reached by the play of the game. For a non-singleton information set, a belief is a probability distribution over the nodes in the information set; for a singleton information set, the player's belief puts probability one on the single decision node.*

Requirement 2 *Given their beliefs, the players' strategies must be sequentially rational. That is, at each information set the action taken by the player with the move (and the player's subsequent strategy) must be optimal given the player's belief at that information set and the other player's subsequent strategies (where a subsequent strategy is a complete plan of action covering every contingency that might arise after the given information set has been reached.)*

Requirement 3 *At information sets on the equilibrium path, beliefs are determined by Bayes' rule and the players' equilibrium strategies.*

Requirement 4 *At information sets off the equilibrium path, beliefs are determined by Bayes' rule and the players' equilibrium strategies where possible.*

1.4.1 Perfect Bayesian Equilibrium in Signalling Games

A signalling game is a dynamic game of incomplete information involving two players: a Sender (S) and a Receiver (R). The timing of the game is as follows:

1. Nature draws a type t_i for the Sender from a set of feasible types $\mathcal{T} = \{t_1, \ldots, t_N\}$ according to a probability distribution $p(t_i)$, where $p(t_i) > 0\ \forall i$ and $\sum_{i=1}^{N} p(t_i) = 1$

2. The Sender observes t_i and then chooses a message m_j from a set of feasible messages $\mathcal{M} = \{m_1 \ldots m_N\}$.

3. The receiver observes m_j but not t_i and then chooses an action a_k from a set of feasible actions $\mathcal{A} = \{a_1, \ldots, a_N\}$.

4. Payoffs are given by $u_S(t_i, m_j, a_k)$ and $u_R(t_i, m_j, a_k)$.

Definition ***Pooling Strategy*** *A strategy from the feasible set of strategies of the Sender is called pooling if each the Sender sends the same message regardless of their type.*

Definition ***Separating Strategy*** *A strategy from the feasible set of strategies of the Sender is called separating if each Sender type sends a different message.*

Definition ***Pooling/Separating Equilibrium*** *If the Sender's strategy is pooling or separating then we call the equilibrium pooling or separating, respectively.*

Definition ***Pure-strategy Perfect Bayesian Equilibrium in a Signalling Game*** *A pure-strategy perfect Bayesian equilibrium in a signalling game is a pair of strategies $m^*(t_i)$ and $a^*(m_j)$ and a belief $\mu(t_i|m_j)$ satisfying:*

Requirement 1 *After observing any message $m_j \in \mathcal{M}$,the Receiver must have a belief about which types could have sent m_j. Denote this belief by the probability distribution $\mu(t_i|m_j)$, where $\mu(t_i|m_j)$ $\forall t_i \in \mathcal{T}$ and $\sum_{t_i \in \mathcal{T}} \mu(t_i|m_j) = 1$.*

Requirement 2 *For each $m_j \in \mathcal{M}$, the Receiver's action $a^*(m_j)$ must maximize the Receiver's expected utility,given the belief $\mu(t_i|m_j)$ about which types could have sent m_j. That is, $a^*(m_j)$ solves*

$$\max_{a_k \in \mathcal{A}} \sum_{t_i \in \mathcal{T}} \mu(t_i|m_j) u_R(t_i, m_j, a_k).$$

Requirement 3 *For each $t_i \in \mathcal{T}$ the Sender's message $m^*(t_i)$ must maximize the Sender's utility, given the Receiver's strategy $a^*(m_j)$. That is, $m^*(t_i)$ solves*

$$\max_{m_j \in \mathcal{M}} u_S(t_i, m_j, a^*(m_j)).$$

Requirement 4 *For each $m_j \in \mathcal{M}$, if there exist $t_i in \mathcal{T}$ such that $m^*(t_i) = m_j$, then the Receiver's belief at the information set corresponding to m_j must follow from Bayes' rule and the Sender's strategy:*

$$\mu(t_i|m_j) = \frac{p(t_i)}{\sum_{t_i \in \mathcal{T}} p(t_i)}.$$

1.5 Evolutionary Game Theory

Suppose that we have a population whose individuals have different phenotypes. As a consequence, they can have different behaviours which might make them more or less fit for a given environment. Let $u(s_1, s_2)$ be the fitness of an individual following a strategy s_i against an opponent following a strategy s_j with $s_j, s_i \in \mathcal{S}$, where $\mathcal{S}$ is the set of feasible strategies that an agent can follow.

Definition ***Evolutionarily Stable Strategy (ESS)*** *A strategy s_i^* from a set of feasible strategies $\mathcal{S}$ is evolutionary stable if there exist $\hat{\epsilon} > 0$ such that for any $s_i \in \mathcal{S}$, $s_i \neq s_i^*$ and for any $\epsilon < \hat{\epsilon}$, we have*

$$u(s_i^*, \epsilon s_i + (1-\epsilon)s_i^*) > u(s_i, \epsilon s_i + (1-\epsilon)s_i^*),$$

Remark: recalling that $u(\cdot, \cdot)$, being an expected utility, is linear in its arguments, we can rewrite

$$u(s_i^*, \epsilon s_i + (1-\epsilon)s_i^*) > u(s_i, \epsilon s_i + (1-\epsilon)s_i^*)$$

as

$$\epsilon u(s_i^*, s_i) + (1-\epsilon)u(s_i^*, s_i^*) > \epsilon u(s_i, s_i) + (1-\epsilon)u(s_i, s_i^*)$$

which can be interpreted as: the strategy s_i^* is evolutionary stable if it cannot be invaded by any $s_i \neq s_i^*$, that is to say, if starting with a population playing s_i^*, a small proportion $\epsilon < \hat{\epsilon}$ of individuals playing s_i, have a lower fitness than those playing s_i^*.

Definition ***Polymorpihc/Monomorpihc ESS*** *If an evolutionary stable strategy allows for more than one strategy to coexist then we say that the ESS is polymorphic. Otherwise, we say that the ESS is mono-morphic.*