1 Games in Strategic Form

A game in **strategic form** or **normal form** is a triple $\Gamma \equiv (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ in which

- $N = \{1, 2, \ldots, n\}$ is a finite set of players,
- S_i is the set of strategies of player i, for every player $i \in N$ - the set of strategy profiles is denoted as $S \equiv S_1 \times \ldots \times S_n$,
- $u_i : S \to \mathbb{R}$ is a utility function that associates with each profile of strategies $s \equiv (s_1, \ldots, s_n)$, a payoff $u_i(s)$ for every player $i \in N$.

Here, the set of strategies can be finite or infinite. The assumption is that players choose these strategies simultaneously in the game, i.e., no player observes the strategies played by other players before playing his own strategy. Here, simultaneous only means they choose their strategies independently without observing each others strategies - one can think of a situation where each player writes down the possible course of action for every possible contingencies in the future and submit it to the game. A strategy profile of all the players will be denoted as $s \equiv (s_1, \ldots, s_n) \in S$. A strategy profile of all the players excluding a Player i will be denoted by s_{-i}. The set of all strategy profiles of players other than a Player i will be denoted by S_{-i}.

We give two examples to illustrate games in strategic form.

1. The first game is the game of Prisoner's Dilemma. Suppose $N = \{1, 2\}$. These players are prisoners. Because of lack of evidence, they have been questioned in separate rooms and made to confess their crimes. If they both confess, then they each achieve a payoff of 1. If both of them do not confess, then they can achieve higher payoffs of 2 each. However, if one of them confesses, but the other one does not confess, then the confessed player gets a payoff of 3 but the player who does not confess gets a payoff of 0.

   What are the strategies in this game? For both the players, the set of strategies is $\{$Confess (C), Do not confess (D)$\}$. The payoffs from the four strategy profiles can be written in a matrix form. It is shown in Table 1.

2. Now, consider an example of an auction. There are two bidders in an auction. Each bidder $i \in \{1, 2\}$ has a value v_i for the object being sold. Each bidder reports a bid in the auction. The highest bidder wins and pays an amount equal to his bid - in case of ties, both win the object with equal probability. The payoff of the bidder from winning

	c	d
C	$(1,1)$	$(3,0)$
D	$(0,3)$	$(2,2)$

Table 1: The Prisoner's Dilemma

is his value minus his bid - in case of ties, $\frac{1}{2}$ times this value. The payoff of the bidder from losing is zero. The strategy of each player in this game in any non-negative real number. If strategy b_i is used by player $i \in \{1,2\}$, then $f_i(b_i, b_{-i})$ be the probability of winning for bidder i - this is 1 if $b_i > b_{-i}$ and zero if $b_i < b_{-i}$ and $\frac{1}{2}$ otherwise. The utility of the bidder i is just $(v_i - b_i)f_i(b_i, b_{-i})$ at a strategy profile (b_i, b_{-i}). Notice that the set of strategies for each bidder is the set of all non-negative real numbers in this case - an infinite set.

The strategy of a game is a powerful tool for representation. It can potentially represent many situations. It provides a complete description of actions that need to be taken in all possible contingencies. As an example, suppose two individuals work every day together on some project for 2 days. Based on the effort put by the individuals on these days, they realize payoffs at the end of two days. Here, a strategy is an effort level in Day 1 and an effort level in Day 2. Players choose such strategies (a combination of effort levels for two days) and that results in payoffs. Later, we will show that many strategic interactions can be reduced to such strategic form by specifying the strategies appropriately.

2 Beliefs of Players

The objective of game theory is to provide predictions of games. To arrive at reasonable predictions for normal form games, let us think how agents will behave in these games. One plausible idea is each agent forms a belief about how other agents will play the game and play his own strategy accordingly. For instance, in the Prisoner's Dilemma game in Table 1, Player 1 may believe that Player 2 will play c with probability $\frac{3}{4}$ and play d with probability $\frac{1}{4}$. In that case, he can compute his payoffs (using expected utility) from both the strategies:

- from playing C: $\frac{3}{4}1 + \frac{1}{4}3 = \frac{6}{4}$,
- from playing D: $\frac{3}{4}0 + \frac{1}{4}2 = \frac{2}{4}$.

Clearly, playing C is better under this belief. Hence, Player 1 will play D given his belief.

Formally, each player i forms a belief $\mu_i \in \Delta S_{-i}$, where ΔS_{-i} is the set of all probability distributions over S_{-i}. Given these beliefs, it computes his utility given his beliefs as:

$$\mathcal{U}_i(s_i, \mu_i) := \sum_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i})\mu_i(s_{-i}) \ \forall \ s_i \in S_i.$$

Then it chooses a strategy s_i^* such that $\mathcal{U}_i(s_i^*, \mu_i) \geq \mathcal{U}_i(s_i, \mu_i)$ for all $s_i \in S_i$.

There are two reasons why this may not work. First, beliefs may not be formed, i.e., where do beliefs come from? Second, beliefs may be incorrect. Even if agent i believes certain strategies will be played by others, other agents may not play them. In game theory, there are two kinds of solution concepts to tackle these issues: (a) solution concepts that work independent of beliefs and (b) solution concepts that assume correct beliefs. The former is sometimes referred to as a *non-equilibrium* solution concept, while the latter is referred to as an *equilibrium* solution concept.

3 Domination

The idea of domination is probably the strongest possible prediction of a game. Dominance is a concept that uses strategies whose performance is good irrespective of the beliefs.

Definition **1** *A strategy $s_i \in S_i$ for Player i is* **strictly dominant** *if for every $s_{-i} \in S_{-i}$, we have*

$$u_i(s_i, s_{-i}) > u_i(s_i', s_{-i}) \ \forall \ s_i' \in S_i \setminus \{s_i\}.$$

Similarly, a strategy $s_i \in S_i$ for Player i is **weakly dominant** *if for every $s_{-i} \in S_{-i}$, we have*

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \ \forall \ s_i' \in S_i \setminus \{s_i\}.$$

It is fairly clear that the idea of domination requires a strategy to be optimal for a player irrespective of what he believes other players are doing. The following lemma formalizes it.

Lemma **1** *A strategy s_i for Player i is strictly dominant if and only if for* all *beliefs μ_i*

$$\mathcal{U}_i(s_i, \mu_i) > \mathcal{U}_i(s_i', \mu_i) \ \forall \ s_i' \in S_i \setminus \{s_i\}.$$

A strategy s_i for Player i is strictly dominant if and only if for all *beliefs μ_i*

$$\mathcal{U}_i(s_i, \mu_i) \geq \mathcal{U}_i(s_i', \mu_i) \ \forall \ s_i' \in S_i \setminus \{s_i\}.$$

Proof: We do the proof for strictly dominant strategies - the weak dominance part follows similarly. Suppose s_i is a strictly dominant strategy for Player i. Fix a belief μ_i. Now, note the following:

$$\begin{aligned}
\mathcal{U}_i(s_i, \mu_i) &= \sum_{s_{-i}} u_i(s_i, s_{-i})\mu_i(s_{-i}) \\
&> \sum_{s_{-i}} u_i(s_i', s_{-i})\mu_i(s_{-i}) \qquad \text{(By definition of strict dominance)} \\
&= \mathcal{U}_i(s_i', \mu_i).
\end{aligned}$$

For the other direction, suppose s_i is an optimal strategy for Player i for all beliefs μ_i. Now, choose some s_{-i} and consider the belief that $\mu_i(s_{-i}) = 1$. Then, it follows that

$$u_i(s_i, s_{-i}) = \mathcal{U}_i(s_i, \mu_i) > \mathcal{U}_i(s_i', \mu_i) = u_i(s_i', s_{-i}).$$

■

In the Prisoner's Dilemma game in Table 1, the strategy C (or c) is a strictly dominant strategy for each player.

If we assume a modest amount of *rationality* in players, we must believe that players must play strictly dominant strategies (whenever they exist). Here, rationality requires that every player plays a strategy that maximizes his utility given his belief about other players' strategies. However, many games do not have a strictly dominant strategy for both the players. For instance, in the game in Table 2, there is no strictly dominant strategy for either of the players.

	L	C	R
T	$(2,2)$	$(6,1)$	$(1,1)$
M	$(1,3)$	$(5,5)$	$(9,2)$
B	$(0,0)$	$(4,2)$	$(8,8)$

Table 2: Domination

However, irrespective of the strategy played by Player 2, Player 1 always gets less payoff in B than in M. In such a case, we will say that Strategy B is strictly dominated.

Definition **2** *A strategy $s_i \in S_i$ for Player i is* **strictly dominanted** *if there exists $s_i' \in S_i$ such that for every $s_{-i} \in S_{-i}$, we have*

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

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

A belief based definition is also possible: irrespective of beliefs of Player i, playing s_i is worse than playing s'_i.

A rational player will never play a strictly dominated strategy. But does that imply we can forget about a strictly dominated strategy? The main issue is removing a strategy of Player i influences the support of the belief of other players. So, unless we assume something about the knowledge level of other players, it is not clear whether we can remove a strategy from Player i.

To see this, consider the example in Table 2. Strategy B is strictly dominated by Strategy M for Player 1. Hence, if Player 1 is rational, then he will not play B. Suppose **Player 2 knows that Player 1 is rational**. Then, he can conclude that Player 1 will not play B ever. As a result, his belief on what Player 1 can play must put probability zero on B. In that case, his Strategy R is strictly dominated by Strategy L. So, he will not play R. Now, if **Player 1 knows that Player 2 is rational and Player 1 knows that Player 2 knows that Player 1 is rational**, then he will not play M because it is now strictly dominated by T. Continuing in this manner, we will get that Player 2 does not play C. Hence, the only strategy profile surviving such elimination is (T, L).

The process we just described is called *iterated elimination of strictly dominated strategies*. It requires more than rationality.

Definition 3 *A* fact *is* **common knowledge** *among players in a game if for any finite chain of player* $(i_1, \ldots, i_k)$ *the following holds:*

Player i_1 *knows that Player* i_2 *knows that Player* i_3 *knows that* $\ldots$ *Player* i_k *knows the fact.*

Iterated elimination of strictly dominated strategies require the following assumption. We will provide a more formal treatment later in this course.

Definition 4 **Common Knowledge of Rationality (CKR)***: The fact that all players are rational is common knowledge.*

Let us consider another example in Table 3. Strategy R is strictly dominated by Strategy M for Player 2. If Player 2 is rational, he does not play R. If Player 1 knows that Player 2 is rational and he himself is rational, then he will assume that R is not played, and T strictly dominates B after removing R. So, he will not play B. If Player 2 knows that Player 1 is rational and Player 2 knows that Player 1 knows Player 2 is rational, then he will not play L. So, iteratively deleting all strictly dominated strategies lead to a unique prediction of (T, M).

	L	M	R
T	$(1,0)$	$(1,2)$	$(0,1)$
B	$(0,3)$	$(0,1)$	$(2,0)$

Table 3: Domination

In many games, iterated elimination of strictly dominated strategies lead to a unique outcome of the game. In those case, we call it a **solution** of the game. However, absence of strictly dominated strategies will imply that no strategies can be eliminated. In such case, iterated elimination of strictly dominated strategies result in no solution. However, the order in which we eliminate strictly dominated strategies does not matter. A formal proof of this fact will be presented later.

In some games, there may not exist any strictly dominated strategy. In such a case, the following weaker notion of weak domination is considered.

DEFINITION **5** *Strategy s_i of Player i is* **weakly dominated** *if there exists another strategy t_i of Player i such that for all $s_{-i} \in S_{-i}$, we have*

$$u_i(s_i, s_{-i}) \leq u_i(t_i, s_{-i}),$$

with strict inequality holding for at least one $s_{-i} \in S_{-i}$. In this case, we say that t_i weakly dominates s_i.

There is no foundation for eliminating (iteratively or otherwise) weakly dominated strategies. Indeed, if we remove weakly dominated strategies iteratively, then the order of elimination matters. This is illustrated in the following example in Table 4.

	L	C	R
T	$(1,2)$	$(2,3)$	$(0,3)$
M	$(2,2)$	$(2,1)$	$(3,2)$
B	$(2,1)$	$(0,0)$	$(1,0)$

Table 4: Order of elimination of weakly dominated strategies

The game in Table 4, there are two weakly dominated strategies for Player 1: $\{T, B\}$. Suppose Player 1 eliminates T first. Then, strategies in $\{C, R\}$ are weakly dominated for Player 2. Suppose Player 2 eliminates R. Then, Player 1 eliminates the weakly dominated strategy B. Finally, Player 2 eliminates Strategy C to leave us with (M, L).

Now, suppose Player 1 eliminates B first. Then, both L and C are weakly dominated. Suppose Player 2 eliminates L first. Then, T is weakly dominated for Player 1. Eliminating T, we see that C is weakly dominated for Player 2. So, we are left with (M, R).

3.1 An Auction Example

However, in some games, weakly dominant strategies give striking prediction. One such example is given below.

The Vickrey Auction. An indivisible object is being sold. There are n buyers (players). Each buyer i has a value v_i for the object, which is completely known to the buyer. Each buyer is asked to report or bid a non-negative real number - denote the bid of buyer i as b_i. The highest bidder wins the object but asked to pay an amount equal to the second highest bid. In case of a tie, all the highest bidders get the object with equal probability and pay the second highest bid, which is also their bid amount in this case. Any buyer who does not win the object pays zero. If a buyer i wins the object and pays a price p_i, then his utility is $v_i - p_i$.

Lemma **2** *In the Vickrey auction, it is a weakly dominant strategy for every buyer to bid his value.*

Proof: Suppose for all $j \in N \setminus \{i\}$, buyer j bids an amount b_j. If buyer i bids v_i, then there are two cases to consider.

Case 1. $v_i > \max_{j \neq i} b_j$. In this case, the payoff of buyer i from bidding v_i is $v_i - \max_{j \neq i} b_j > 0$. By bidding something else, if he is not the unique highest bidder, then he either does not get the object or he gets the object with lower probability and pays the same amount. In the first case, his payoff is zero and in the second case, his payoff is strictly less than $v_i - \max_{j \neq i} b_j$. Hence, bidding v_i is a weakly dominant strategy.

Case 2. $v_i \leq \max_{j \neq i} b_j$. In this case, the payoff of buyer i from bidding v_i is zero. If he bids an amount smaller than v_i, then he does not get the object and his payoff is zero. If he bids an amount larger than v_i, then he gets the object with probability one and pays $\max_{j \neq i} b_j$, and hence, his payoff is $v_i - \max_{j \neq i} b_j \leq 0$. Hence, bidding v_i is a weakly dominant strategy for buyer i. ■

4 Mixed Strategies

We now consider a game with a finite set of strategies. Sometimes, it is natural to assume that players play different strategies with different probabilities - the idea of belief was already reflecting this.

Formally, for any finite set A, we denote by ΔA, the set of all probability distributions over A: $\Delta A := \{p : A \to [0,1] : \sum_{a \in A} p(a) = 1\}$. For any finite strategy set S_i of Player i, every $\sigma_i \in \Delta S_i$ is a **mixed strategy** of Player i. In this case S_i is called the set of **pure strategies** of Player i. A mixed strategy profile is $\sigma \equiv (\sigma_1, \ldots, \sigma_n) \in \prod_{i \in N} \Delta S_i$. Under mixed strategy, players are assumed to randomize independently, i.e., how a player randomizes does not depend on how others randomize.

Often, a finite normal form game $\Gamma \equiv (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ may be given. The **mixed extension** of Γ is given by $(N, \{\Delta S_i\}_{i \in N}, \{U_i\}_{i \in N})$, where for all $i \in N$, for all $\sigma \in \prod_{i \in N} \Delta S_i$, we have

$$U_i(\sigma) = \sum_{s \equiv (s_1, \ldots, s_n) \in S} u_i(s) \sigma_1(s_1) \ldots \sigma_n(s_n).$$

Note that the mixed extension of a game is an infinite game - it includes all possible lotteries over pure strategies of a player. Further, the utility function is a linear extension of the utility function of the original pure strategy game.

Consider the following game in Table 5. Suppose Player 1 plays the mixed strategy A with probability $\frac{3}{4}$ and B with probability $\frac{1}{4}$. Suppose Player 2 plays a with probability $\frac{1}{4}$ and b with probability $\frac{3}{4}$. Then, the mixed strategy profile is

$$\sigma \equiv (\sigma_1, \sigma_2) = \Big((\sigma_1(A), \sigma_1(B)), (\sigma_2(a), \sigma_2(b)) \Big) = \Big((\frac{3}{4}, \frac{1}{4}), (\frac{1}{4}, \frac{3}{4}) \Big).$$

	a	b
A	$(3,1)$	$(0,0)$
B	$(0,0)$	$(1,3)$

Table 5: Mixed strategies

From this, the probability with which each pure strategy profile is played can be computed (using independence). These probabilities are shown in Table 6. A player computes the utility from a mixed strategy profile using expected utility. The mixed strategy profile σ

gives players payoffs as follows:

$$\begin{aligned}
U_1(\sigma) &= u_1(A,a)\sigma_1(A)\sigma_2(a) + u_1(A,b)\sigma_1(A)\sigma_2(b) + u_1(B,a)\sigma_1(B)\sigma_2(a) + u_1(B,b)\sigma_1(B)\sigma_2(b) \\
&= 3\frac{3}{16} + 0 + 0 + 1\frac{3}{16} \\
&= \frac{3}{4} \\
U_2(\sigma) &= u_2(A,a)\sigma_1(A)\sigma_2(a) + u_2(A,b)\sigma_1(A)\sigma_2(b) + u_2(B,a)\sigma_1(B)\sigma_2(a) + u_2(B,b)\sigma_1(B)\sigma_2(b) \\
&= 1\frac{3}{16} + 0 + 0 + 3\frac{3}{16} \\
&= \frac{3}{4}.
\end{aligned}$$

	a	b
A	$\frac{3}{16}$	$\frac{9}{16}$
B	$\frac{1}{16}$	$\frac{3}{16}$

Table 6: Mixed strategies - probability of all pure strategy profiles

4.1 Domination

Nothing changes in strict dominance if we consider mixed strategies. We make the following observations.

- A mixed strategy that puts positive probability on more than one pure strategies cannot be strictly dominant. To see this, suppose it puts positive probability on s_i and t_i. But then, the utility from such a mixed strategy cannot exceed $\max(u_i(s_i), u_i(t_i))$. This contradicts the fact that it is a strictly dominant strategy.

- If a pure strategy is a strictly dominant strategy in a finite normal game with pure strategies, then it is also a strictly dominant strategy in its mixed extension. This is because if a pure strategy dominates all other pure strategies, it must dominate any lottery involving those pure strategies and itself.

- A pure strategy that is not dominated by any pure strategy may be dominated by a mixed strategy. To see this, consider the example in Table 7. Strategy C is not dominated by any pure strategy for Player 1. However, the mixed strategy $\frac{1}{2}A$ and $\frac{1}{2}B$

	a	b
A	$(3,1)$	$(0,4)$
B	$(0,2)$	$(3,1)$
C	$(1,0)$	$(1,2)$

Table 7: Mixed strategies may dominate pure strategies

strictly dominates the pure strategy C. Hence, C is a strictly dominated strategy for Player 1 in the mixed extension of the game described in Table 7.

- If a pure strategy is strictly dominated, then any mixed strategy which has this pure strategy in its support is also strictly dominated. This is because if a pure strategy s_i is strictly dominated by σ_i. Then, in any mixed strategy with s_i in its support, we can transfer the probability on s_i to σ_i to increase its utility, and this will dominate the mixed strategy. For instance, in the example in Table 7, C is strictly dominated by the mixed strategy $\frac{1}{2}A + \frac{1}{2}B$. Hence, the mixed strategy $\frac{2}{3}B + \frac{1}{3}C$ is strictly dominated by the mixed strategy $\frac{2}{3}B + \frac{1}{3}\Big(\frac{1}{2}A + \frac{1}{2}B\Big) \equiv \frac{1}{6}A + \frac{5}{6}B$.

- Even if a group of pure strategies are not strictly dominated, a mixed strategy with only these strategies in its support may be strictly dominated. To see this, consider the game in Table 8. The pure strategies A and B are not strictly dominated. But the mixed strategy $\frac{1}{2}A + \frac{1}{2}B$ is strictly dominated by pure strategy C.

	a	b
A	$(3,1)$	$(0,4)$
B	$(0,2)$	$(3,1)$
C	$(2,0)$	$(2,2)$

Table 8: Mixed strategies may be dominated

5 Nash Equilibrium

One of the problems with the idea of domination is that often there are no dominated strategies. Hence, it fails to provide any prediction about many games. For instance, consider the game in Table 9. No pure strategy in this game is dominated.

We may now revisit the strong requirement of domination that a strategy is best irrespective of the beliefs we have about what others are playing. In many cases, games are

	a	b
A	$(3,1)$	$(0,4)$
B	$(0,2)$	$(3,1)$

Table 9: No dominated strategies

results of repeated outcomes. For instance, if two firms are interacting in a market, they have a good idea about each other's cost and technology. As a result, they can form accurate beliefs about what other player is playing. The idea of Nash equilibrium takes this accuracy to the limit - it assumes that each player has **correct** belief about what others are playing and responds optimally given his (correct) beliefs.

DEFINITION **6** *A strategy profile* $(s_1^*, \ldots, s_n^*)$ *in a strategic form game* $\Gamma \equiv (N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N})$ *is a* **Nash equilibrium** *of* Γ *if for all* $i \in N$

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*) \; \forall \; s_i \in S_i.$$

The game Γ can be a mixed extension of another game. In that case, the strategy profile under consideration in the above definition may be a mixed strategy profile. Similarly, the game Γ in the above definition may be a finite or an infinite game.

The idea of a Nash equilibrium is that of a *steady state*, where each player is responding optimally given the strategies of the other players - no unilateral deviation is possible. It does not argue how this steady state is reached. It has a notion of stability - if a player finds certain unilateral deviation profitable, then such a steady state cannot be sustained.

An alternate definition using the idea of *best response* is often useful. A strategy s_i of Player i is a **best response** to the strategy s_{-i} of other players if

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \; \forall \; s_i' \in S_i.$$

The set of all best response strategies of Player i given the strategies of other players is denoted by $B_i(s_{-i})$. This definition can be written in terms of beliefs as well - s_{-i} is a belief over the strategies of other players.

Now, a strategy profile $(s_1^*, \ldots, s_n^*)$ is a **Nash equilibrium** if for all $i \in N$,

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

The following observation is immediate.

CLAIM **1** *If* s_i *is a strictly dominant strategy of Player* i*, then* $\{s_i\} = B_i(s_{-i})$ *for all* $s_{-i} \in S_{-i}$*. Hence, if* $(s_1, \ldots, s_n)$ *is a unqiue Nash equilibrium.*

It is extremely important to remember that Nash equilibrium assumes correct beliefs and best responding with respect to these correct beliefs of other players. There are other interpretations of Nash equilibrium. Consider a mediator who offers the players a strategy profile to play. A player agrees with the mediator if (a) he believes that others will agree with the mediator and (b) strategy proposed to him by the mediator is a best response to the strategy proposed to others. This is precisely the idea behind a Nash equilibrium.

5.1 Examples (Pure Strategies)

We give various examples of games where a Nash equilibrium (in pure strategies) exist. In Table 10, we consider the Prisoner's Dilemma game. We claim that (A, a) is a Nash equilibrium of this game - if Player 1 plays A, the best response of Player 2 consists of only strategy a and if Player 2 plays a, the best response of Player 1 consists of only strategy A. Note that this is also the outcome in strictly dominant strategies.

	a	b
A	$(1,1)$	$(5,0)$
B	$(0,5)$	$(4,4)$

Table 10: Nash equilibrium in Prisoner's Dilemma

Consider now the game (called the *coordination game*) in Table 11. The game is called coordination game since if players do not coordinate in this game they both get zero payoff. If they coordinate, then they get the same payoff but (A, a) is worse than (B, b) for both the players. If Player 2 plays a, then $B_1(a) = \{A\}$ and if Player 1 plays A, then $B_2(A) = \{a\}$. So, (A, a) is a Nash equilibrium. Now, if Player 2 plays b, then $B_1(b) = \{B\}$ and if Player 1 plays B, then $B_2(B) = \{b\}$. Hence, (B, b) is another Nash equilibrium. This example shows you that there may be more than one Nash equilibrium in a game.

	a	b
A	$(1,1)$	$(0,0)$
B	$(0,0)$	$(3,3)$

Table 11: Nash equilibrium in the Coordination game

Another game that has more than one Nash equilibrium is the *Battle of the sexes*. A man and a woman are deciding which movie to go between two movies $\{X, Y\}$. Man wants to see movie X and woman wants to see movie Y. However, if both of them go to separate movies, then they get zero payoff. Their preferences are reflected in Table 12. If Woman plays x, then Man's best response is $\{X\}$ and if Man plays X, then Woman's best response is $\{x\}$. Hence, (X, x) is a Nash equilibrium. Using a similar logic, we can compute (Y, y) to be a Nash equilibrium. These are the only Nash equilibria of the game.

	x	y
X	$(2, 1)$	$(0, 0)$
Y	$(0, 0)$	$(1, 2)$

Table 12: Nash equilibrium in the Battle of the Sexes game

Now, we discuss a game with infinite number of strategies. This game is called the the Cournot Duopoly game. Two firms $\{1, 2\}$ produce the same product in a market where there is a common price for the product. They simultaneously decide how much to produce - denote by q_1 and q_2 respectively the quantities produced by firms 1 and 2. If the total quantity produced by both the firms is $q_1 + q_2$, then the product price is assumed to be $2 - q_1 - q_2$. Suppose the per unit cost of productions are: $c_1 > 0$ for firm 1 and $c_2 > 0$ for firm 2. We will assume that $q_1, q_2, c_1, c_2 \in [0, 1]$. We will now compute the Nash equilibrium of this game.

This is a two player game. Each player's strategy is the quantity it produces. If firms 1 and 2 produce q_1 and q_2 respectively, then their payoffs are given by

$$u_1(q_1, q_2) = q_1(2 - q_1 - q_2) - c_1 q_1$$
$$u_2(q_1, q_2) = q_2(2 - q_1 - q_2) - c_2 q_2.$$

Given q_2, firm 1 can maximize its payoff my maximizing u_1 over all q_1. To do so, we take the first order condition for u_1 to get

$$2 - 2q_1 - q_2 - c_1 = 0.$$

This simplifies to

$$q_1 = \frac{1}{2}(2 - c_1 - q_2).$$

Similarly, we get

$$q_2 = \frac{1}{2}(2 - c_2 - q_1).$$

Solving these two equations we get

$$q_1^* = \frac{2 - 2c_1 + c_2}{3}, q_2^* = \frac{2 - 2c_2 + c_1}{3}.$$

These are necessary conditions for optimality. We need to verify that it is a Nash equilibrium. For this, first note that

$$\begin{aligned} u_1(q_1^*, q_2^*) &= (q_1^*)^2 \\ u_2(q_1^*, q_2^*) &= (q_2^*)^2 \end{aligned}$$

Now, given firm 2 sets q_2^*, let us find the utility when firm 1 sets q_1:

$$\begin{aligned} u_1(q_1, q_2^*) &= \frac{q_1}{3}\big[4 + 2c_2 - 4c_1 - 3q_1\big]. \\ &= 2q_1 q_1^* - (q_1)^2 \\ &\leq (q_1^*)^2 \\ &= u_1(q_1^*, q_2^*). \end{aligned}$$

A similar calculation suggests

$$u_2(q_1^*, q_2) \leq u_2(q_1^*, q_2^*).$$

Hence, (q_1^*, q_2^*) is a Nash equilibrium. This is also a unique Nash equilibrium (why?).

We now consider an example of a two-player game where payoffs of both the players add up to zero. This particular game is called the *matching pennies*. Two players toss two coins. If they both turn Heads or Tails, then Player 1 is paid by Player 2 Rs. 1. Else, Player 1 pays Player 2 Rs. 1. The payoff of each player is the money he receives (or the negative of the money he pays). The payoffs are shown in Table 13. For the moment assume that, what turns up in the coin is in the control of the players - for instance, a player may choose to show Heads in his coin.

The Matching Pennies game has no Nash equilibrium. To see this, note that when Player 2 plays h, then the unique best response of Player 1 is H. But when Player 1 plays H, the unique best response of Player 2 is t. Also, when Player 2 plays t the unique best response of Player 1 is T, but when Player 1 plays T the unique best response of Player 2 is h.

	h	t
H	$(1, -1)$	$(-1, 1)$
T	$(-1, 1)$	$(1, -1)$

Table 13: The Matching Pennies game

6 The Maxmin Value

Consider a game shown in Table 14. There is a unique Nash equilibrium of this game: (B, R) - verify this. But, will Player 1 play strategy B? What if Player 2 makes a mistake in his belief and plays L? Then, Player 1 will get -100 by playing B. Thinking this, Player 1 may like to play safe, and play a strategy like T that guarantees him a payoff of 2. For Player 2 also, strategy R may be bad if Player 1 decides to play T. On the other hand, strategy L can guarantee him a payoff of 0.

	L	R
T	$(2, 1)$	$(2, -20)$
M	$(3, 0)$	$(-10, 1)$
B	$(-100, 2)$	$(3, 3)$

Table 14: The Maxmin idea

The main message of the example is that sometimes players may choose to play strategy to guarantee themselves some *safe* level of payoff without assuming anything about the rationality level of other players. In particular, we consider the case where every player believes that the other players are *adversaries* and are here to punish him - this is a very pessimistic view of the opponents. In such a case, what can a player guarantee for himself?

If Player i chooses a strategy $s_i \in S_i$ in a game, then the worst payoff he can get is

$$\min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}).$$

Of course, we are assuming here that the strategy sets and the utility functions are such that a minimum exists - else, we can define an infimum.

Definition **7** *The* **maxmin** *value for Player i in a strategic form game $(N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ is given by*

$$\underline{v}_i := \max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}).$$

Any strategy that guarantees Player i a value of $\underline{v}_i$ is called a **maxmin** *strategy.*

Note that if s_i is a maxmin strategy for Player i, then it satisfies

$$\min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) \geq \min_{s_{-i} \in S_{-i}} u_i(s'_i, s_{-i}) \ \ \forall \ s'_i \in S_i.$$

This also means that $u_i(s_i, s_{-i}) \geq \underline{v}_i$ for all $s_{-i} \in S_{-i}$.

In the example in Table 14, we see that $\underline{v}_1 = 2$ and $\underline{v}_2 = 0$. Strategy T is a maxmin strategy for Player 1 and strategy L is a maximin strategy for Player 2. Hence, when players play their maxmin strategy, the outcome of the game is $(2, 1)$. However, there can be more than one maxmin strategies in a game, in which case no unique outcome can be predicted. Consider the example in Table 15. The maxmin strategy for Player 1 is B. But Player 2 has two maxmin strategies $\{L, R\}$, both giving a payoff of 1. Depending on which maxmin strategy Player 2 plays the outcome can be $(2, 3)$ or $(1, 1)$.

	L	R
T	$(3, 1)$	$(0, 4)$
B	$(2, 3)$	$(1, 1)$

Table 15: More than one maxmin strategy

It is clear that if a player has a weakly dominant strategy, then it is a maxmin strategy - it guarantees him the best possible payoff irrespective of what other agents are playing. Hence, if every player has a weakly dominant strategy, then the vector of weakly dominant strategies constitute a vector of maxmin strategies. This was true, for instance, in the example involving the second-price sealed-bid auction. Further, if there are strictly dominant strategies for each player (note such strategy must be unique for each player), then the vector of strictly dominant strategies constitute a unique vector of maxmin strategies.

The following theorem shows that a Nash equilibrium of a game guarantees the maxmin value for every player.

Theorem 1 *Every Nash equilibrium s^* of a strategic form game satisfies*

$$u_i(s^*) \geq \underline{v}_i \quad \forall \; i \in N.$$

Proof: For any Player i and for every $s_i \in S_i$, we know that

$$u_i(s_i, s^*_{-i}) \geq \min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}).$$

By definition, $u_i(s^*_i, s^*_{-i}) = \max_{s_i \in S_i} u_i(s_i, s^*_{-i})$. Combining with the above inequality, we get

$$u_i(s^*_i, s^*_{-i}) = \max_{s_i \in S_i} u_i(s_i, s^*_{-i}) \geq \max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}) = \underline{v}_i.$$

■

6.1 Elimination of Dominated Strategies

We now describe the effect on maxmin value by eliminating dominated strategies. Though elimination of dominated strategies require extreme assumptions on rationality compared to maxmin strategies, the relation between the outcomes in both the cases is interesting. As a consequence, we will see that the relationship between the set of Nash equilibria of a game and the set of Nash equilibria of a game that survives iterated elimination of dominated strategies.

Theorem **2** *Let $\Gamma = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ be a finite game in strategic form and Γ' be the game generated by removing a weakly dominated strategy s'_j of Player j from Γ. Then, the maxmin value of Player j in Γ' is equal to his maxmin value in Γ.*

Proof: Let s_j be a strategy that weakly dominates s'_j for Player j in Γ. Then, $u_j(s_j, s_{-j}) \geq u_j(s'_j, s_{-j})$ for all s_{-j}. Hence,

$$\min_{s_{-j} \in S_{-j}} u_j(s_j, s_{-j}) \geq \min_{s_{-j} \in S_{-j}} u_j(s'_j, s_{-j})$$

Now, note that

$$\max_{t_j \in S_j, t_j \neq s'_j} \min_{s_{-j} \in S_{-j}} u_j(t_j, s_{-j}) \geq \min_{s_{-j} \in S_{-j}} u_j(s_j, s_{-j}) \geq \min_{s_{-j} \in S_{-j}} u_j(s'_j, s_{-j}).$$

This implies that

$$\begin{aligned}
\underline{v}_j &= \max_{t_j \in S_j} \min_{s_{-j} \in S_{-j}} u_j(t_j, s_{-j}) \\
&= \max\big(\max_{t_j \in S_j, t_j \neq s'_j} \min_{s_{-j} \in S_{-j}} u_j(t_j, s_{-j}), \min_{s_{-j} \in S_{-j}} u_j(s'_j, s_{-j})\big) \\
&= \max_{t_j \in S_j, t_j \neq s'_j} \min_{s_{-j} \in S_{-j}} u_j(t_j, s_{-j}) \\
&= \underline{v}'_j,
\end{aligned}$$

where $\underline{v}_j$ and $\underline{v}'_j$ are the maxmin values of Player j in games Γ and Γ' respectively. ■

Note that elimination of weakly or strictly dominated strategy of Player j does not have any effect on the maxmin value of Player j but it may increase (though never decrease) the maxmin value of other players - this follows from the fact that eliminating strategies of other players only increases your worst payoff for every strategy, and hence, increases your maxmin value.

The next result states that if we eliminate some strategies (dominated or not) of a player, then every Nash equilibrium of the original game that survived this elimination continues to be a Nash equilibrium of the new game.

Theorem 3 *Let Γ be a finite game in strategic form and Γ' be a game derived from Γ by eliminating some of the strategies of each player. If s^* is a Nash equilibrium of Γ and s^* is available in Γ', then s^* is a Nash equilibrium in Γ'.*

Proof: Let S'_i be the set of strategies remaining for each player i in Γ' and S_i be the set of original strategies in Γ for each player i. By definition,

$$u_i(s^*) \geq u_i(s_i, s^*_{-i}) \ \ \forall \ s_i \in S_i.$$

But $S'_i \subseteq S_i$ implies that $u_i(s^*) \geq u_i(s_i, s^*_{-i}) \ \ \forall \ s_i \in S'_i$. Hence, s^* is also a Nash equilibrium of Γ'. ∎

Note that eliminating arbitrary strategies though will not eliminate original Nash equilibria, it may introduce new Nash equilibria. The following theorem shows that this is not possible if weakly dominated strategies are eliminated.

Theorem 4 *Let Γ be a finite game in strategic form and s_j be a weakly dominated strategy for Player j in this game. Denote by Γ' the game derived by eliminating strategy s_j from Γ. Then, every Nash equilibrium of Γ' is also a Nash equilibrium of Γ.*

Proof: Let s^* be a Nash equilibrium of Γ'. Consider a player $i \neq j$. By definition, $u_i(s^*) = \max_{s_i \in S_i} u_i(s_i, s^*_{-i})$. Since the set of strategies of i is the same in both the games, this establishes that i cannot unilaterally deviate. For Player j, we note that s_j is weakly dominated, say by strategy t_j. Then,

$$u_j(s_j, s^*_{-j}) \leq u_j(t_j, s^*_{-j}) \leq \max_{s'_j \in S_j : s'_j \neq s_j} u_j(s'_j, s^*_{-j}) = u_j(s^*_j, s^*_{-j}),$$

where the last equality follows since s^* is a Nash equilibrium of Γ'. This shows that $u_j(s^*_j, s^*_{-j}) \geq u_j(s'_j, s^*_{-j})$ for all $s'_j \in S_j$. Hence, s^* is also a Nash equilibrium of Γ. ∎

The above theorem implies that if we iteratively eliminate weakly dominated strategies and look at the Nash equilibria of the resulting game, they will also be Nash equilibria of the original game. However, we may lose some of the Nash equilibria of the original game. Consider the game in Table 16. Suppose Player 2 eliminates L and then Player 1 eliminates B. We are then left with (T, R). However, (B, L) is a Nash equilibrium of the original game. Note that (T, R) is also a Nash equilibrium of the original game (implied by Theorem 4).

However, this cannot happen if we eliminate strictly dominated strategies.

Theorem 5 *Let Γ be a finite game in strategic form and s_j be a strictly dominated strategy for Player j in this game. Denote by Γ' the game derived by eliminating strategy s_j from Γ. Then, the set of Nash equilibria in Γ and Γ' are the same.*

	L	R
T	$(0,0)$	$(2,1)$
B	$(3,2)$	$(1,2)$

Table 16: Elimination may lose equilibria

Proof: By Theorem 4, we need to show that if s^* is a Nash equilibrium of Γ, then s^* is also a Nash equilibrium of Γ'. Note that the strategy profile s^* is still available to all the agents in Γ' since only a strictly dominated strategy is eliminated for Player j. Formally, for Player j, there exists a strategy t_j such that $u_j(t_j, s^*_{-j}) > u_j(s_j, s^*_{-j})$. Hence, $u_j(s^*_j, s^*_{-j}) \geq u_j(t_j, s^*_{-j}) > u_j(s_j, s^*_{-j})$. So, $s^*_j \neq s_j$. Since s^* is available in Γ', by Theorem 3, s^* is a Nash equilibrium of game Γ'. ■

This theorem leads to some interesting corollaries. First, a strictly dominated strategy cannot be part of a Nash equilibrium. Second, if elimination of strictly dominated strategies lead to a unique outcome, then that outcome is the unique Nash equilibrium of the original game. In other words, to compute the Nash equilibrium or maxmin value, we can iteratively eliminate all strictly dominated strategies of the players.

7 Existence of Nash Equilibrium in Finite Games

As we have seen that not all games have a Nash equilibrium. This section is devoted to results that describe sufficient conditions on games for a Nash equilibrium to exist. We start from the celebrated theorem of Nash and end with some theorems on existence of pure strategy Nash equilibrium. All the theorems have one theme in common - existence of Nash equilibrium is equivalent to establishing existence of a fixed point of an appropriate map.

In this section, instead of talking about mixed extension of a game, we will refer to the mixed strategies of a player in a game explicitly. Before establishing the main theorem, we provide a useful lemma.

Lemma **3 (Indifference Principle)** *Suppose $\sigma_i \in B_i(\sigma_{-i})$ and $\sigma_i(s_i) > 0$. Then, $s_i \in B_i(\sigma_{-i})$.*

Proof: Suppose $\sigma_i \in B_i(\sigma_{-i})$. Let $S_i(\sigma_i) := \{s_i \in S_i : \sigma_i(s_i) > 0\}$. If $|S_i(\sigma_i)| = 1$, then the claim is obviously true. Else, pick $s_i, s'_i \in S_i(\sigma_i)$. We argue that $u_i(s_i, \sigma_{-i}) = u_i(s'_i, \sigma_{-i})$. First note that the net utility from playing σ_i is given by

$$\sum_{s''_i \in S_i(\sigma_i)} u_i(s''_i, \sigma_{-i})\sigma_i(s''_i).$$

Suppose $u_i(s_i, \sigma_{-i}) > u_i(s'_i, \sigma_{-i})$. Then, transferring the probability on s'_i to s_i in σ_i increases the net utility of agent i, contradicting the fact that σ_i is best response to σ_{-i}. This shows that

$$u_i(s_i, \sigma_{-i}) = u_i(s'_i, \sigma_{-i}) \;\; \forall \; s_i, s'_i \in S_i(\sigma_i).$$

This also means that $U_i(\sigma_i, \sigma_{-i}) = u_i(s_i, \sigma_{-i})$ for all $s_i \in S_i(\sigma_i)$. Hence, $s_i \in B_i(\sigma_{-i})$ for all $s_i \in S_i(\sigma_i)$. ■

Now, we prove Nash's seminal theorem.

THEOREM **6 (Nash)** *Every finite game has a Nash equilibrium in mixed strategies.*

Proof: We do the proof in several steps.

STEP 1. For each profile of mixed strategy σ, for each player $i \in N$, and for each pure strategy $s_i^j \in S_i$, we define

$$g_i^j(\sigma) := \max\Big(0, U_i(s_i^j, \sigma_{-i}) - U_i(\sigma)\Big),$$

where U_i is the net payoff function agent i from playing a mixed strategy, which is derived using the von-Neumann-Morgenstern expected utility.

The interpretation of $g_i^j(\sigma)$ is that it is zero if Player i does not find deviating to s_i^j from σ profitable. Else, it captures the increase in payoff of Player i from (σ) to (s_i^j, σ_{-i}). Note that Player i can profitably deviate from σ if and only if it can profitably deviate from σ using a pure strategy - Lemma 3. This implies that σ is a Nash equilibrium if and only if $g_i^j(\sigma) = 0$ for all $i \in N$ and for all $j \in \{1, \ldots, |S_i|\}$.

STEP 2. Now, we show that for each i and each j, g_i^j is a continuous function. To see this note that U_i is continuous in σ and σ_{-i}. As a result, $U_i(s_i^j, \sigma_{-i}) - U_i(\sigma)$ is a continuous function. The max of two continuous functions is continuous. Hence, g_i^j is continuous.

STEP 3. Using g_i^j, we define another map f_i^j in this step. For every $i \in N$, for every $s_i^j \in S_i$, and for every σ, define

$$f_i^j(\sigma) := \frac{\sigma_i(s_i^j) + g_i^j(\sigma)}{1 + \sum_k g_i^k(\sigma)}.$$

The amount $f_i^j(\sigma)$ is supposed to hint that if σ_i is not a better response to σ_{-i}, then how much probability on s_i^j should be increased - thus, it gives another improved mixed strategy.

Also, it is easy to see that for each i and each j, $f_i^j(\sigma) \geq 0$. Further,

$$\begin{aligned}
\sum_{j=1}^{|S_i|} f_i^j(\sigma) &= \sum_{j=1}^{|S_i|} \frac{\sigma_i(s_i^j) + g_i^j(\sigma)}{1 + \sum_k g_i^k(\sigma)} \\
&= \frac{\sum_{j=1}^{|S_i|} \sigma_i(s_i^j) + g_i^j(\sigma)}{1 + \sum_k g_i^k(\sigma)} \\
&= \frac{1 + \sum_{j=1}^{|S_i|} g_i^j(\sigma)}{1 + \sum_k g_i^k(\sigma)} \\
&= 1.
\end{aligned}$$

Hence, $f_i(\sigma) \equiv (f_i^1(\sigma), \ldots, f_i^{|S_i|}(\sigma))$ is another mixed strategy of Player i. Further, f_i^j is a continuous function since both numerator and denominator are non-negative continuous functions. Hence, $f(\sigma) \equiv (f_1(\sigma), \ldots, f_n(\sigma))$ is also a continuous function.

Step 4. In this step, we introduce the idea of a fixed point of a function and use it to show a result.

Definition **8** *Let $F : X \rightarrow X$ be a function defined on X. If $F(x) = x$ for some $x \in X$, then x is called a fixed point of F.*

We show that if $f(\sigma) = \sigma$, i.e., σ is a fixed point of f, then for all $i \in N$ and for all j,

$$g_i^j(\sigma) = \sigma_i(s_i^j) \sum_k g_i^k(\sigma).$$

To see this, using the fixed point property and the definition of f_i^j, we see that

$$\begin{aligned}
f_i^j(\sigma) &= \sigma_i(s_i^j) \\
&= \frac{\sigma_i(s_i^j) + g_i^j(\sigma)}{1 + \sum_k g_i^k(\sigma)}.
\end{aligned}$$

Rearranging, we get the desired equality.

Step 5. In this step of the proof, we show that if σ is a fixed point of f, then σ is a Nash equilibrium. Suppose σ is not a Nash equilibrium. Then, for some Player i, there is a strategy s_i^j such that $g_i^j(\sigma) > 0$. As a result $\sum_k g_i^k(\sigma) > 0$. From the previous step, we know that $\sigma_i(s_i^k) > 0$ if and only if $g_i^k(\sigma) > 0$.

Now, note that $U_i(\sigma) = \sum_k \sigma_i(s_i^k) U_i(s_i^k, \sigma_{-i})$. Hence,

$$\begin{aligned}
0 &= \sum_k \sigma_i(s_i^k)\big(U_i(s_i^k, \sigma_{-i}) - U_i(\sigma)\big) \\
&= \sum_{k:\sigma_i(s_i^k)>0} \sigma_i(s_i^k)\big(U_i(s_i^k, \sigma_{-i}) - U_i(\sigma)\big) \\
&= \sum_{k:\sigma_i(s_i^k)>0} \sigma_i(s_i^k) g_i^k(\sigma) \\
&> 0,
\end{aligned}$$

where the last equality and the strict inequality follows from our earlier observation that $g_i^k(\sigma) > 0$ if and only if $\sigma_i(s_i^k) > 0$.

STEP 6. This leads to the last step of the theorem. In this step, we show that a fixed point of f exists. For this, we use the following fixed point theorem due to Brouwer.

THEOREM **7 (Brouwer's fixed point theorem)** *Let X be a convex and compact set in $\mathbb{R}^k$ and let $F : X \to X$ be a continuous function. Then, there exists a fixed point of F.*

Now, we have already argued that f is a continuous function. The domain of f is the set of all strategy profiles. Since this is the set of all mixed strategies of a finite set of pure strategies, it is a compact and convex set. Finally, the range of f belongs to the set of strategy profiles. Hence, by Brouwer's fixed point theorem, there exists a fixed point of f. By the previous step, such a fixed point corresponds to the Nash equilibrium of the finite game. ■

Some comments about the proof of Nash's theorem. Simpler proofs are possible using a stronger version of fixed point theorem - due to Kakutani. This proof is the original proof of Nash, where he uses the Brouwer's fixed point theorem. The Brouwer's fixed point theorem is not simple to proof, but you are encouraged to look at its proof. In one-dimension, the Brouwer's fixed point theorem is the *intermediate value theorem*.

7.1 Computing Mixed Strategy Equilibrium - Examples

In general, computing mixed strategy equilibrium of a game is computationally difficult. However, couple of thumb-rules make it easier for finding the set of all Nash equilibria. First, we should iteratively eliminate all strictly dominated strategies. As we have learnt, the set of Nash equilibria remains the same after iteratively eliminating strictly dominated strategies. The second is a crucial property that we have already established - the indifference principle in Lemma 3.

We start off by a simple example on how to compute all Nash equilibria of a game. Consider the game in Table 17.

	L	R
T	$(8,8)$	$(8,0)$
B	$(0,8)$	$(9,9)$

Table 17: Nash equilibria computation

First, note that no strategies can be eliminated as strictly dominated. It is easy to verify that (T, L) and (B, R) are two pure strategy Nash equilibria of the game. To compute mixed strategy Nash equilibria, suppose Player 1 plays T with probability p and B with probability $(1-p)$, where $p \in (0, 1)$. Then, by playing L, Player 2 gets

$$8p + 8(1-p) = 8.$$

By playing R, Player 2 gets

$$9(1-p).$$

L is best response to $pT + (1-p)B$ if and only if $8 \geq 9(1-p)$ or $p \geq \frac{1}{9}$. Else, R is a best response. Note that Player 2 is indifferent between L and R when $p = \frac{1}{9}$ - this follows from the indifference lemma that we have proved. Hence, if Player 2 mixes, then Player 1 must play $\frac{1}{9}T + \frac{8}{9}B$. But, when Player 2 plays $qL + (1-q)R$, then Player 1 gets 8 by playing T and $9(1-q)$ by playing B. For Player 1 to mix, Player 2 must make him indifferent between playing T and B, which happens at $q = \frac{1}{9}$. Thus, $(\frac{1}{9}T + \frac{8}{9}B, \frac{1}{9}L + \frac{8}{9}R)$ is also a Nash equilibrium of this game. Note that the payoff achieved by both the players by playing this strategy profile is 8.

There are some strategies of a player which are not strictly dominated, but which can still be eliminated before computing the Nash equilibrium. These are strategies which are *never* best responses.

DEFINITION **9** *A strategy $\sigma_i \in \Delta S_i$ is* **never a best response** *for Player i if for every $\sigma_{-i} \in \Delta S_{-i}$,*

$$\sigma_i \notin B_i(\sigma_{-i}).$$

The following claim is a straightforward observation.

CLAIM **2** *If a strategy is strictly dominated, then it is never a best response.*

The next claim says that we can remove all pure strategies that are not best responses to compute Nash equilibrium.

LEMMA **4** *If a pure strategy $s_i \in S_i$ is never a best response, then any mixed strategy σ_i with $\sigma_i(s_i) > 0$ is not a Nash equilibrium strategy.*

Proof: Suppose $s_i \in S_i$ is never a best response but there is a mixed strategy Nash equilibrium σ with $\sigma_i(s_i) > 0$. By the Indifference Lemma (Lemma 3), s_i is also a best response to σ_{-i}, contradicting the fact s_i is never a best response. ■

The connection between never best response strategies and strictly dominated strategies is deeper. Indeed, in two-player games, a strategy is strictly dominated if and only if it is never a best response. We will come back to this once we discuss zero-sum games. We will now use Lemma 4 to compute Nash equilibria efficiently.

Consider the two player game in Table 18. Computing Nash equilibria of such a game can be quite tedious. However, we can be smart in avoiding certain computations.

	L	C	R
T	$(3,3)$	$(0,0)$	$(0,2)$
M	$(0,0)$	$(3,3)$	$(0,2)$
B	$(2,2)$	$(2,2)$	$(2,0)$

Table 18: Nash equilibria computation

In two player 3-strategy games, we can draw the best response correspondences in a 2-d simplex - Figure 1 represents the simplex of Player 1's strategy space for the game in Table 18. Any point inside the simplex represents a probability distribution over the three strategies of Player 1, and these probabilities are given by the lengths of perpendiculars to the three sides. To see this suppose we pick a point in the simplex with lengths of perpendiculars to sides $(T,B),(T,M),(M,B)$ as p_m, p_b, p_t respectively. The following fact from Geometry is useful.

FACT **1** *For every point inside an equilateral triangle with lengths of perpendiculars (p_m, p_b, p_t), the sum of $p_m + p_b + p_t$ equals to $\sqrt{3}a/2$, where a is the length of sides of the equilateral triangle.*

This fact can be proved easily by using the fact the sum of three triangles generated by any point is the same - $\sqrt{3}a^2/4 = \frac{1}{2}a(p_m + p_t + p_b)$. Hence, without loss of generality, we will scale the lengths of the sides of the simplex to $\frac{2}{\sqrt{3}}$. As a result, $p_m + p_t + p_b = 1$ and the numbers p_m, p_t, p_b reflect a probability distribution. We will follow this term to represent strategies in two player 3-strategy games.

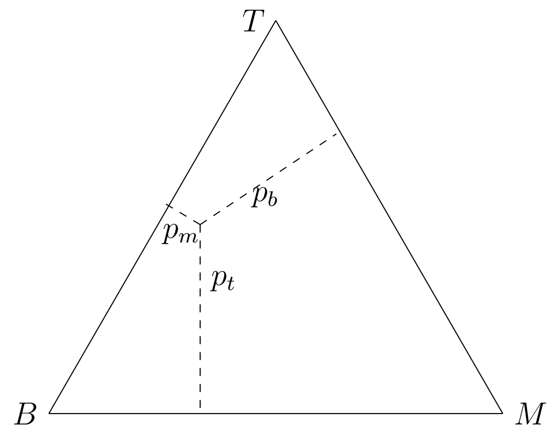

Figure 1: Representing probabilities on a 2d-simplex

Now, let us draw the best response correspondence of Player 1 for various strategies of Player 2: $B_1(\sigma_2)$ will be drawn on the simplex of strategies of Player 2 - see Figure 2. For this, we fix a strategy $\sigma_2 = (\alpha L + \beta C + (1-\alpha-\beta)R)$ of Player 2. We now identify conditions on α and β to identify pure strategy best responses of Player 1. By the Indifference Lemma, the mixed strategy best responses happen at the intersection of these pure strategy best response regions. We consider three cases:

Case 1- T. $T \in B_1(\sigma_2)$ if

$$3\alpha \geq 3\beta$$
$$3\alpha \geq 2.$$

Combining these conditions together, we get $\alpha \geq \frac{2}{3}$ and $\alpha \geq \beta$. The second condition holds if $\alpha \geq \frac{2}{3}$. So, we deduce that the best response region of T are all mixed strategies where L is played with at least $\frac{2}{3}$ probability. This is shown in Figure 2.

Case 2 - M. $M \in B_1(\sigma_2)$ if

$$3\beta \geq 3\alpha$$
$$3\beta \geq 2.$$

This gives us a similar condition to Case 1: $\beta \geq \frac{2}{3}$. The best response region of M is shown in the simplex of Player 2's strategies in Figure 2.

Case 3 - B. Clearly $B \in B_1(\sigma_2)$ in the remaining regions and at all the boundary points where B and T are indifferent and B and M are indifferent. This is shown in Figure 2 in the simplex of Player 2's strategy.

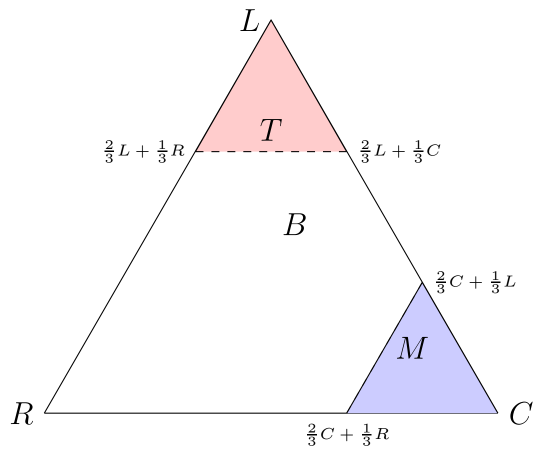

Figure 2: Best response map of Player 1

Once the best response map of Player 1 is drawn, we conclude that no best response involves mixing T and M together. So, every mixed strategy best response involves mixing B.

We now draw the best response map of Player 2. For this we consider a mixed strategy $\alpha T + \beta M + (1 - \alpha - \beta)B$ of Player 1. For L to be a best response of Player 2 against this strategy, we must have

$$3\alpha + 2(1 - \alpha - \beta) \geq 3\beta + 2(1 - \alpha - \beta)$$
$$3\alpha + 2(1 - \alpha - \beta) \geq 2(\alpha + \beta).$$

This gives us

$$\alpha \geq \beta$$
$$2 \geq \alpha + 4\beta.$$

The line $\alpha = \beta$ is shown in Figure 2. To draw $2 = \alpha + 4\beta$, we pick two points: (i) $\alpha = 0$ and $\beta = \frac{1}{2}$ and (ii) $\alpha + \beta = 1$ and $\beta = \frac{2}{3}$. The line joining these two points depict $2 = \alpha + 4\beta$. Now, the entire best response region of L is shown in Figure 2.

An analogous argument shows that for C to be a best response we must have

$$\beta \geq \alpha$$
$$2 \geq \beta + 4\alpha.$$

The best response region of strategy C is shown in Figure 3. The remaining area is the best response region of strategy R (including the borders with L and C).

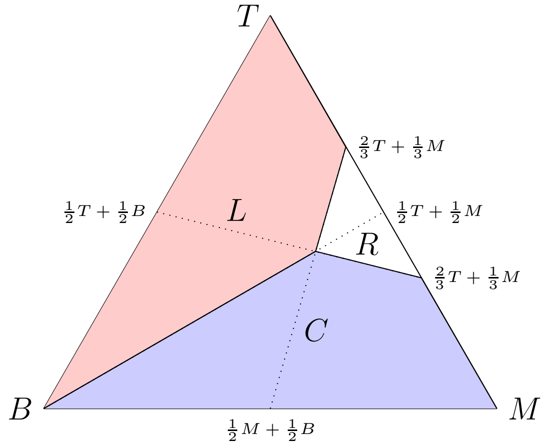

Figure 3: Best response map of Player 2

Computing Nash equilibria. To compute Nash equilibria, we see that there is no best response of Player 1 where T and M are mixed. Further, R is a best response of Player 2 when T and M are mixed. Hence, there cannot be a Nash equilibrium (σ_1, σ_2) such that $\sigma_2(R) > 0$. So, in any Nash equilibrium, Player 2 either plays L or C or mixed L and C but puts zero probability on R.

Since no mixing of T and M is possible for Player 1 in Nash equilibrium, we must look at the best response map of Player 2 when mix of T and B and mix of M and B is played. That corresponds to the two edges of the simplex corresponding to (T, B) and (M, B) in Figure 3. In that region, mixture of L and C is a best response when B is played with probability 1. So, in any Nash equilibrium where L and C is mixed Player 1 plays B for sure. But then looking into the best response map of Player 1 in Figure 2, we see that Player 1 best responds B for sure if Player 2 mixes $\alpha L + (1-\alpha)C$ with $\alpha \in [\frac{1}{3}, \frac{2}{3}]$. The other pure strategy Nash equilibria are (T, L) and (M, C).

So, we can enumerate all the Nash equilibria of the game in Table 18 now:

$$(T, L), (M, C), (B, \alpha L + (1-\alpha)C),$$

where $\alpha \in [\frac{1}{3}, \frac{2}{3}]$.

7.2 Two Player Zero-Sum Games

The two player zero-sum games occupy a central role in game theory because of variety of reasons. First, they were the first set of games to be theoretically analyzed by von-Neumann and Morgenstern when they came up with the theory of games. Second, the zero-sum games are ubiquitous - examples include any real game where one player's loss is another player's gain. Formally, a zero-sum game is defined as follows.

Definition **10** *A finite* **zero-sum game of two players** *is defined as* $N = \{1, 2\}$ *and* (S_1, S_2), (u_1, u_2) *with the restriction that for all* $(s_1, s_2) \in S_1 \times S_2$*, we have*

$$u_1(s_1, s_2) + u_2(s_1, s_2) = 0.$$

Because of this restriction, we can define a zero-sum two player game by a single utility function $u : S_1 \times S_2 \to \mathbb{R}$, where $u(s_1, s_2)$ represents utility of Player 1 and $-u(s_1, s_2)$ represents the utility of Player 2.

	h	t
H	$(1, -1)$	$(-1, 1)$
T	$(-1, 1)$	$(1, -1)$

Table 19: Matching pennies

Consider the two player zero-sum game in Table 19. It is called the *matching pennies* game - the strategies are sides of a coin, if the sides match then Player 1 wins and pays Player 2 Rs. 1, else Player 2 wins and pays Player 1 Rs. 1. There is no pure strategy Nash equilibrium of this game. To compute mixed strategy Nash equilibrium, suppose Player 2 plays $\alpha h + (1 - \alpha)t$. To make Player 1 indifferent between H and T, we see that

$$\alpha + (-1)(1 - \alpha) = -\alpha + (1 - \alpha).$$

This gives us $\alpha = \frac{1}{2}$. A similar calculation suggests that if Player 2 has to mix in best response, Player 1 must play $\frac{1}{2}H + \frac{1}{2}T$. Hence, $(\frac{1}{2}H + \frac{1}{2}T, \frac{1}{2}h + \frac{1}{2}t)$ is the unique Nash equilibrium of this game. Note that the payoff achieved by both the players in this Nash equilibrium is zero.

Now, suppose Player 1 plays $\frac{1}{2}H + \frac{1}{2}T$, the worst payoff that he can get from Player 2's strategies can be computed as follows. If Player 2 plays h or t Player 1 gets a payoff of 0. Hence, his worst payoff is 0. As a result, the maxmin value of Player 1 is at least zero. We know (by Theorem 1) that the Nash equilibrium payoff is at least the maxmin value. [1]

[1] Theorem 1 continues to hold even we allow for mixed strategies.

Hence, the maxmin value is also zero. A similar calculation suggests that the maxmin value of Player 2 is also zero. We show that this is true for *any* finite two player zero-sum game.

The maxmin value of Player 1 in a zero sum game is denoted by

$$\underline{v}_1 := \max_{\sigma_1 \in \Delta S_1} \min_{\sigma_2 \in \Delta S_2} u(\sigma_1, \sigma_2).$$

The maxmin value of Player 2 in a zero sum game is denoted by

$$\underline{v}_2 := \max_{\sigma_2 \in \Delta S_2} \min_{\sigma_1 \in \Delta S_1} -u(\sigma_1, \sigma_2) = -\min_{\sigma_2 \in \Delta S_2} \max_{\sigma_1 \in \Delta S_1} u(\sigma_1, \sigma_2).$$

We denote by $\underline{v} := \max_{\sigma_1 \in \Delta S_1} \min_{\sigma_2 \in \Delta S_2} u(\sigma_1, \sigma_2)$ and $\bar{v} := \min_{\sigma_2 \in \Delta S_2} \max_{\sigma_1 \in \Delta S_1} u(\sigma_1, \sigma_2)$. Note that $\underline{v}_1 = \underline{v}$ and $\underline{v}_2 = -\bar{v}$.

DEFINITION **11** *A finite two player zero-sum game* **has a value** *if $\underline{v} = \bar{v}$. In that case, $\underline{v} = \bar{v}$ is called the value of the game, and is denoted by v. Any maxmin and minmax strategies of Player 1 and Player 2 respectively are called* **optimal** *strategies.*

The main result for two person zero-sum game is the following.

THEOREM **8** *If a finite two player zero-sum game has a value v and if σ_1^* and σ_2^* are optimal strategies of the two players, then $\sigma^* \equiv (\sigma_1^*, \sigma_2^*)$ is a Nash equilibrium with payoff $(v, -v)$. Conversely, if $\sigma^* \equiv (\sigma_1^*, \sigma_2^*)$ is a Nash equilibrium of a finite two player zero-sum game, then the game has a value $v = u(\sigma_1^*, \sigma_2^*)$, and strategies σ_1^* and σ_2^* are optimal strategies.*

Proof: Suppose a two player zero-sum game has a value v and if σ_1^* and σ_2^* are optimal strategies of the two players. Then, since σ_1^* is optimal for Player 1, we get

$$u(\sigma_1^*, \sigma_2^*) = v = \min_{\sigma_2 \in \Delta S_2} u(\sigma_1^*, \sigma_2).$$

Hence, for all $\sigma_2 \in \Delta S_2$,

$$u(\sigma_1^*, \sigma_2^*) \leq u(\sigma_1^*, \sigma_2).$$

This gives us for all $\sigma_2 \in \Delta S_2$, $u_2(\sigma_1^*, \sigma_2^*) \geq u_2(\sigma_1^*, \sigma_2)$. Further, since σ_2^* is optimal for Player 2, we get

$$u(\sigma_1^*, \sigma_2^*) = v = \max_{\sigma_1 \in \Delta S_1} u(\sigma_1, \sigma_2^*).$$

Hence, for all $\sigma_1 \in \Delta S_1$,

$$u_1(\sigma_1^*, \sigma_2^*) \geq u_1(\sigma_1, \sigma_2^*).$$

This establishes that (σ_1^*, σ_2^*) is a Nash equilibrium. Clearly, the payoffs are $(v, -v)$.

For the other direction, suppose (σ_1^*, σ_2^*) is a Nash equilibrium. Then, for all $\sigma_1 \in \Delta S_1$, we have $u(\sigma_1^*, \sigma_2^*) \geq u(\sigma_1, \sigma_2^*)$. Hence,

$$u(\sigma_1^*, \sigma_2^*) = \max_{\sigma_1 \in \Delta S_1} u(\sigma_1, \sigma_2^*) \geq \min_{\sigma_2 \in \Delta S_2} \max_{\sigma_1 \in \Delta S_1} u(\sigma_1, \sigma_2) = \bar{v}.$$

Note that by Theorem 1, $-u(\sigma_1^*, \sigma_2^*) \geq -\bar{v}$ or $\bar{v} \geq u(\sigma_1^*, \sigma_2^*)$. Hence, we have

$$u(\sigma_1^*, \sigma_2^*) = \bar{v}.$$

Next, for all $\sigma_2 \in \Delta S_2$, we have $-u(\sigma_1^*, \sigma_2^*) \geq -u(\sigma_1^*, \sigma_2)$. Hence,

$$u(\sigma_1^*, \sigma_2^*) = \min_{\sigma_2 \in \Delta S_2} u(\sigma_1^*, \sigma_2) \leq \max_{\sigma_1 \in \Delta S_1} \min_{\sigma_2 \in \Delta S_2} u(\sigma_1, \sigma_2) = \underline{v}.$$

By Theorem 1, $u(\sigma_1^*, \sigma_2^*) \geq \underline{v}$. Hence, we get

$$\bar{v} = u(\sigma_1^*, \sigma_2^*) = \underline{v}.$$

Hence, the game has a value $v = u(\sigma_1^*, \sigma_2^*)$ and σ_1^* and σ_2^* are optimal strategies. ■

An immediate corollary using Nash theorem is the following.

Corollary **1** *Every two player zero-sum game has a value v. The payoff from any Nash equilibrium correspond to $(v, -v)$.*

Proof: Every finite game has a Nash equilibrium. By Theorem 8, a value of a two player zero sum game exists and the value corresponds to the payoff of Player 1 and negative of payoff of Player 2. ■

7.3 Interpretations of Mixed Strategy Equilibrium

Considering mixed strategies guarantee existence of Nash equilibrium in finite games. However, it is not clear why a player will randomize in the precise way prescribed by a mixed strategy Nash equilibrium, specially given the fact he is indifferent between the pure strategies in the support of such a Nash equilibrium. There are no clear answers to this question. However, following are some arguments to validate that mixed strategies can be part of Nash equilibrium play.

- Players some times randomize deliberately. For instance, in zero-sum games with two players, players randomize to play their max min strategies. In games like Poker, players have been shown to randomize.

- Mixed strategy equilibrium can be thought to be a belief system - if σ^* is a Nash equilibrium, then σ_i^* describes the belief that opponents of Player i have on Player i's behavior. This means that Player i may not actually randomize but his opponents collectively believe that σ_i^* is the strategy he will play. Hence, a mixed strategy equilibrium is just a steady state of beliefs.

- One can think of a strategic form game being played over time repeatedly (payoffs and actions across periods do not interact). Suppose players choose a best response in each period assuming time average of plays of past (with some initial conditions on how to choose strategies). In particular, they observe that opponents have been playing a strategy A for $\frac{3}{4}$ times and another strategy B for the remaining time. So, they optimally respond by forming this as their belief. It has been shown that such plays eventually converge to a steady state where the average play of each player is some mixed strategy.

- Another interpretation that is provided by Nash himself interprets Nash equilibrium as population play. There are two pools of large population. We draw a player at random from each pool and pair them against each other. The strategy of that player will reflect the expected strategy played by the population and will represent a mixed strategy. So, Nash equilibrium represents some kind of stationary distribution of pure strategies in such these population.

8 Existence of Pure Strategy Nash Equilibrium

In many games pure strategy Nash equilibria exist. Whenever it exists, it provides more compelling prediction of a game than a mixed strategy Nash equilibrium. Further, the mixed strategy Nash equilibrium existence theorem of Nash only applies to finite games. In this section, we investigate settings under which a pure strategy Nash equilibrium exists in a game (finite or infinite). Typically, in these games we need to assume certain topological and geometrical properties about the strategy sets of players and their utility functions.

8.1 Continuity and Convexity Assumptions

The first such existence theorem is a generalization of the ideas found in Nash's theorem.

Theorem **9** *Suppose* $\Gamma \equiv (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ *is a game in strategic form such that for each* $i \in N$

1. *S_i is a compact and convex subset of $\mathbb{R}^{K_i}$ for some integer K_i.*

2. *$u_i(s_i, s_{-i})$ is continuous in s_{-i}.*

3. *$u_i(s_i, s_{-i})$ is continuous and concave in s_i.* [2]

Then, Γ has a pure strategy Nash equilibrium.

Proof: The proof of this theorem is done using Kakutani's fixed point theorem.

THEOREM **10 (Kakutani's Fixed Point Theorem)** *Let A be a non-empty subset of a finite dimensional Euclidean space. Let $f : A \to 2^A$ be a map which satisfies the following properties.*

1. *A is compact and convex.*

2. *$f(x)$ is a non-empty subset of A for each $x \in A$.*

3. *$f(x)$ is a convex subset of A for each $x \in A$.*

4. *$f(x)$ has a closed graph for each $x \in A$, i.e., if $\{x^k, y^k\} \to \{x, y\}$ with $y^k \in f(x^k)$ for each k, then $y \in f(x)$.*

Then, there exists $x \in A$ such that $x \in f(x)$.

We use Theorem 10 in a straightforward manner to establish existence of Nash equilibrium. For every strategy profile s, define

$$B(s) = \{s' : s'_i \in B_i(s_{-i}) \ \forall \ i \in N\}.$$

Note that s is a Nash equilibrium if and only if $s \in B(s)$. We show that B satisfies all the conditions of Theorem 10.

1. Since each S_i is compact and convex, the set of strategy profiles is also compact and convex.

2. For every s and for every $i \in N$,

   $$B_i(s_{-i}) = \{s'_i \in S_i : u_i(s'_i, s_{-i}) = \max_{s''_i \in S_i} u_i(s''_i, s_{-i})\}.$$

   This set is non-empty because of u_i is continuous in s''_i and S_i is compact - so, by Weirstrass theorem a maximum of the function exists. As a result $B(s)$ is also non-empty.

[2] A concave function is continuous in the interior of the domain. Requiring continuity here makes it continuous even at the boundary points.

3. Next, we show that $B(s)$ is convex. Pick, $t, t' \in B(s)$ and $\lambda \in (0,1)$. Define $t'' \equiv \lambda t + (1-\lambda)t'$. We show that for every $i \in N$, $t''_i \in B_i(s_{-i})$. Since $t_i, t'_i \in B_i(s_{-i})$, we get

$$u_i(t_i, s_{-i}) = u_i(t'_i, s_{-i}) = \max_{s'_i} u_i(s'_i, s_{-i}).$$

But then concavity of u_i implies that

$$u_i(t''_i, s_{-i}) \geq \lambda u_i(t_i, s_{-i}) + (1-\lambda)u_i(t'_i, s_{-i}) = \max_{s'_i} u_i(s'_i, s_{-i}).$$

Hence, $t''_i \in B_i(s_{-i})$, and this implies that $B(s)$ is convex.

4. Finally, we show that B has a closed graph. To see this, assume for contradiction that B does not have a closed graph. Then, for some sequence $(t^k, \bar{t}^k) \to (t, \bar{t})$ with $\bar{t}^k \in B(t^k)$, we have $\bar{t} \notin B(t)$. This means, for some $i \in N$, $\bar{t}_i \notin B_i(t_{-i})$. This implies that there exists some $s'_i \in S_i$ and $\epsilon > 0$ such that

$$u_i(s'_i, t_{-i}) > u_i(\bar{t}_i, t_{-i}) + \epsilon.$$

By the continuity of u_i, we get that there is some k such that t^k and t are close enough such that

$$u_i(s'_i, t^k_{-i}) \geq u_i(s'_i, t_{-i}) - \frac{\epsilon}{2}.$$

Combining these two inequalities we get

$$u_i(s'_i, t^k_{-i}) > u_i(\bar{t}_i, t_{-i}) + \frac{\epsilon}{2} \geq u_i(\bar{t}^k_i, t^k_{-i}) + \frac{\epsilon}{4},$$

where we used continuity of u_i in the second inequality. This is a contradiction because $\bar{t}^k_i \in B_i(t^k_{-i})$ implies $u_i(\bar{t}^k_i, t^k_{-i}) \geq u_i(s'_i, t^k_{-i})$.

Now, we apply Kakutani's fixed point theorem (Theorem 10) to conclude that there exists s such that $s \in B(s)$. This implies that s is a pure strategy Nash equilibrium. ■

Notice that the Nash's theorem is an immediate corollary of Theorem 9. To see how Theorem 9 can and cannot be applied, consider the following location game. Two shops (players) are locating on the line segment $[0,1]$ which has a uniform distribution of customers. Once the shops are located, customers go to the nearest shop with tie broken with equal probability. The utility of a shop is the mass of customers that go there. So, strategy sets of both the players are $S_1 = S_2 = [0,1]$, a convex and compact set. If the shops locate themselves at (s_1, s_2) with $s_1 \leq s_2$, then the utilities of the shops are

$$u_1(s_1, s_2) = \frac{s_1 + s_2}{2}, u_2(s_1, s_2) = 1 - \frac{s_1 + s_2}{2}.$$

Hence, fixing s_2 as s_1 approaches s_2, we see that $u_1(s_1, s_2)$ approaches s_2 but as s_1 crosses s_2 for values arbitrarily close to s_2 it has a value of $1 - s_2$. Hence, u_1 is not continuous in s_1. So, Theorem 9 cannot be applied here. But we know that pure strategy Nash equilibrium exists in such games.

Second, consider the Cournot duopoly game with two firms. When firms produce q_1 and q_2, the price in the market is $2 - q_1 - q_2$ and unit costs of the firms are c_1 and c_2 respectively. Then, the utility function of each firm i is

$$u_i(q_1, q_2) = q_i(2 - q_1 - q_2) - c_i q_i.$$

This is continuous in both q_i and q_{-i}. Further, it is concave in q_i. Hence, it satisfies all the conditions of Theorem 9. Further, if we assume that the allowable quantities are some closed interval in the non-negative real line, then the strategy set of each firm is compact and convex. Theorem 9 guarantees that a pure strategy Nash equilibrium exists.

8.2 Supermodular Games

The concavity assumption made in Theorem 9 does not hold in many games. We now discuss a class of games where we provide a different set of sufficient conditions that guarantee existence of pure strategy Nash equilibrium. These are called *supermodular* games. Supermodular games capture the idea that strategies of players are complements of each other. The main idea of a supermodular game is that the marginal utility of one player's utility is non-decreasing in the strategies of the other players.

To define supermodularity, we need to introduce the mathematical structure of lattice. For each $x, y \in \mathbb{R}^K$, define $x \wedge y \in \mathbb{R}^K$ (meet) and $x \vee y \in \mathbb{R}^K$ (join) as: $(x \wedge y)_i = \min(x_i, y_i)$ for all i and $(x \vee y)_i = \max(x_i, y_i)$ for all i.

Definition **12** *A set $X \subseteq \mathbb{R}^K$ is a* **lattice** *if $(x \vee y), (x \wedge y) \in X$ for all $x, y \in X$.*

Lattice is more general than what this definition suggests. This definition is a particular type of lattice which is a subset of a finite dimensional Euclidian space.

Consider the following two examples of lattice: (a) an interval $[0, 1]$ and (b) $\{x \in \mathbb{R}^K : x_i \geq x_{i+1} \ \forall\ i \in \{1, \ldots, K-1\}\}$. The set of points inside a circle on $\mathbb{R}^2$ is *not* a lattice but inside a square is a lattice.

We say a point x^* in a lattice X is the **greatest** element of X if $x^* \geq x$ for all $x \in X$. Similarly, $x_* \in X$ is the **least** element of X if $x_* \leq x$ for all $x \in X$. [3] It is an easy exercise to verify the following fact.

[3]Here, when we say $y \geq z$, we mean $y_i \geq z_i$ for all i.

FACT **2** *If $X \subseteq \mathbb{R}^K$ is a compact set and a lattice, then it has a greatest and a least element.*

This fact is easily observed by minimizing or maximizing over all $x \in X$ the sum $\sum_{i=1}^{K} x_i$. Since this is a continuous function and X is compact, a minimum and a maximum exists. Further, such a minimum (and maximum) must be unique because of the lattice property of X.

One easy fixed point theorem to prove is by Tarski. We state a weak variant of this fixed point theorem.

THEOREM **11 (Tarski's Fixed Point Theorem)** *Let $X \subseteq \mathbb{R}^K$ be a compact lattice and $f : X \to X$ be a monotone function (i.e., $f(x) \geq f(y)$ for all $x \geq y$). Then, f has a fixed point.*

The Tarski's fixed point theorem works for arbitrary compact lattices (not necessary that the lattice is a subset of a Euclidean space). We give a proof below assuming continuity of f.

Proof: Since X is a compact lattice, there is a least element of X, denote this by x_*. Start from $x^1 := x_*$. Let $y^1 := f(x^1)$. If $y^1 = x^1$, then we are done. Since x_* is the least element of X, $y^1 \geq x_* = x^1$. Then, set $x^2 := y^1 > x^1$. By monotonicity of f, $f(x^2) \geq f(x^1)$. By setting $y^2 := f(x^2)$, we see that $y^2 \geq y^1 = x^2$. We now iterate this process - at every stage $x^{k+1} := y^k \geq x^k$. By monotonicity, $y^{k+1} := f(x^{k+1}) \geq f(x^k) = y^k = x^{k+1}$. If $y^{k+1} = x^{k+1}$, then we are done. Else, $y^{k+1} > x^{k+1}$. Since $x^{k+2} = y^{k+1}$, we see that $x^{k+2} \geq x^{k+1}$. Hence, the sequences $\{x_k\}_k$ and $\{y_k\}_k$ are monotone sequences in X with the property that $x^{k+1} = y^k$ and $y^k \geq x^k$ for every k. In particular we get sequences $\{x_k\}_k \equiv (x^1, x^2, x^3, \ldots, x^k, \ldots)$ and $\{y^k\}_k \equiv (x^2, x^3, \ldots, x^k, \ldots)$. Since X has a greatest element, these sequences converge. Further since the sequence $\{y_k\}_k$ is just $\{x_{k+1}\}_k$, both the sequences converge to the same point - denote this limit as x (note that the proof does not end here since we still have to show that $f(x) = x$).

Since f is continuous $\lim_k f(x^k) = f(x)$. But $f(x^k) = x^{k+1}$ for each k. Hence, $\lim_k f(x^k) = \lim_k x^{k+1} = x$. This implies that $x = f(x)$. ■

Now, we define the increasing differences property. Suppose $X \subseteq \mathbb{R}^K$ and $Y \subseteq \mathbb{R}^L$ be two compact lattices. It is easy to verify that $X \times Y$ is also a compact lattice $\mathbb{R}^{K+L}$. Consider a function $f : X \times Y \to \mathbb{R}$, where X and Y are compact lattices as mentioned above. Then, f will be said to have the increasing differences property if marginal value from elements in X is increasing with increase in elements from Y.

Definition **13** *Let $X \subseteq \mathbb{R}^K$ and $Y \subseteq \mathbb{R}^L$ be two lattices. A function $f : X \times Y \to \mathbb{R}$ satisfies increasing differences in (x, y) if for all $x, x' \in X$ with $x' \geq x$ and for all $y, y' \in Y$ with $y' \geq y$, we have*

$$f(x', y') - f(x, y') \geq f(x', y) - f(x, y).$$

To understand increasing differences, consider a function $f : \mathbb{R}^2 \to \mathbb{R}$ and note that $\mathbb{R}^2$ is a lattice. Suppose $f(x, y) = x(1 - y)$. Now, $f(1, 1) - f(0, 1) = 0$ and $f(1, 0) - f(0, 0) = 1$. Hence, such a function does not satisfy increasing differences - increasing y decreases the marginal value of x. However, $f(x, y) = x(1 + y)$ satisfies increasing differences.

A closely related concept is supermodularity.

Definition **14** *Let $X \subseteq \mathbb{R}^K$ be a lattice. A function $f : X \to \mathbb{R}$ is* **supermodular** *if for all $x, x' \in X$, we have*

$$f(x \vee x') + f(x \wedge x') \geq f(x) + f(x').$$

We state (without proof) some elementary facts about supermodularity and increasing differences. We assume X and Y are two lattices below.

1. A function $f : X \to \mathbb{R}$ is supermodular if and only if for every $i, j \in \{1, \ldots, K\}$, and every x_{-ij} $f(x_i, x_j, x_{-ij})$ satisfies increasing differences for all x_i, x_j.

2. A function $f : X \times Y$ satisfies increasing differences in (x, y) if and only if f satisfies increasing differences for any pair (x_i, y_j) given any (x_{-i}, y_{-j}).

3. If f is twice continuously differentiable on $X = \mathbb{R}^K$, f is supermodular if and only if $\frac{\partial^2 f}{\partial x_i \partial x_j} \geq 0$ for all x_i, x_j.

The following is an important result regarding monotone comparative statics on lattices.

Theorem **12 (Topkis Monotone Comparative Statics)** *Let $X \subseteq \mathbb{R}^K$ be a compact lattice and $T \subseteq \mathbb{R}^L$ be a lattice. Suppose $f : X \times T \to \mathbb{R}$ is supermodular and continuous on X for every $t \in T$ and satisfies increasing differences in (x, t). Define for every $t \in T$,*

$$x^*(t) := \{x \in X : f(x, t) \geq f(x', t) \ \forall \ x' \in X\}.$$

Then, the following are true:

1. *for every $t \in T$, $x^*(t) \subseteq X$ is a non-empty compact lattice.*

2. *for every $t, t' \in T$ with $t' > t$ and for every $x \in x^*(t)$ and $x' \in x^*(t')$, we have*

$$x \vee x' \in x^*(t') \quad \text{and} \quad x \wedge x^* \in x^*(t).$$

3. *for every $t, t' \in T$ with $t' > t$ we have*

$$\bar{x}^*(t') \geq \bar{x}^*(t) \quad \text{and} \quad \underline{x}^*(t') \geq \underline{x}^*(t),$$

where for every $t'' \in T$, $\bar{x}^*(t'')$ and $\underline{x}^*(t'')$ are the greatest and least elements of the lattice $x^*(t'')$ respectively.

Proof: For every $t \in T$, $x^*(t)$ is non-empty and compact by the Weierstrass theorem. For any $x, x' \in x^*(t)$, we know that

$$f(x \vee x', t) + f(x \wedge x', t) \geq f(x, t) + f(x', t).$$

Either $f(x \vee x', t) \geq f(x \wedge x', t)$ or $f(x \vee x', t) \leq f(x \wedge x', t)$. Suppose $f(x \vee x', t) \geq f(x \wedge x', t)$. Since $x, x' \in x^*(t)$, we get $f(x, t) = f(x', t)$, and hence, $f(x \vee x', t) \geq f(x, t)$. Since $x \in x^*(t)$, $x \vee x' \in x^*(t)$. This implies that $f(x \vee x', t) = f(x, t) = f(x', t)$. But then, $f(x \wedge x', t) \geq f(x, t)$, implying that $x \wedge x' \in x^*(t)$. A similar proof works if $f(x \vee x', t) \leq f(x \wedge x', t)$. This shows that $x^*(t)$ is a compact lattice.

Now pick $t, t' \in T$ with $t' > t$ and $x, x' \in X$ with $x' \in x^*(t')$ and $x \in x^*(t)$. We know that $f(x, t) - f(x \wedge x', t) \geq 0$. By increasing differences, we get $f(x, t') - f(x \wedge x', t') \geq 0$. By supermodularity, we get $f(x' \vee x, t') - f(x', t') \geq 0$. Hence, $x' \vee x \in x^*(t')$. Hence, for any $x \in x^*(t)$ and $x' \in x^*(t')$, we have $x \leq x' \vee x \leq \bar{x}^*(t')$. Hence, $\bar{x}^*(t) \leq \bar{x}^*(t')$.

Also, $f(x \vee x', t') - f(x', t') \leq 0$. By increasing differences, $f(x \vee x', t) - f(x', t) \leq 0$. By supermodularity, $f(x, t) - f(x \wedge x', t) \leq 0$. Since $x \in x^*(t)$, we see that $x \wedge x' \in x^*(t)$. Hence, for any $x \in x^*(t)$ and $x' \in x^*(t')$, we have $x' \geq x \wedge x' \geq \underline{x}^*(t)$. Hence, $\underline{x}^*(t) \leq \underline{x}^*(t')$. ■

This leads us to the definition of the supermodular game.

Definition 15 *A game $(N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ is* **supermodular** *if for every $i \in N$,*

- *$S_i \subseteq \mathbb{R}^{K_i}$ is a compact lattice,*
- *u_i is continuous and supermodular in s_i for every s_{-i},*
- *u_i satisfies increasing differences in (s_i, s_{-i}).*

Note that a supermodular game does not assume continuity of u_i with respect to other players strategies s_{-i}. It also does not assume concavity of u_i with respect to s_i. For instance, if $S_i \subseteq \mathbb{R}$ for every i, then u_i is vacuously supermodular in s_i for every s_{-i}. Hence, we will only need continuity of u_i in s_i (contrast this to concavity requirement in Theorem 9). Another **important** point: all the lattice-theoretic results we proved for lattices in $\mathbb{R}^K$ can also be proved for finite lattices with a greatest element and a least element - this is a general definition of a compact lattice. Hence, supermodular games can also be defined when S_i for each i is finite and a compact lattice. The result below will apply to such a case also.

Now, we state the main result of this section.

THEOREM **13** *Every supermodular game has a pure strategy Nash equilibrium.*

Proof: Pick any strategy profile s. For every $i \in N$ and for every $s_{-i} \in S_{-i}$, define

$$B_i(s_{-i}) = \{s_i : u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}) \; \forall \; s'_i\}.$$

Since S_i and S_{-i} are compact lattices, by Theorem 12, $B_i(s_{-i})$ is a non-empty compact lattice. Now, we define $\bar{B}_i(s_{-i})$ as the greatest element of $B_i(s_{-i})$ - note that this is a strategy in S_i. Now, we can define for every strategy profile s,

$$\bar{B}(s) := (\bar{B}_1(s_{-1}), \ldots, \bar{B}_n(s_{-n})).$$

Hence, $\bar{B} : S_1 \times \ldots \times S_n \to S_1 \times \ldots \times S_n$. By Theorem 12, if $s' \geq s$, then $\bar{B}_i(s'_{-i}) \geq \bar{B}_i(s_{-i})$ for all $i \in N$. Hence, $\bar{B}$ is a monotone function defined on a compact lattice. By Theorem 11, a fixed point of $\bar{B}$ exists. But such a fixed point is a Nash equilibrium, which completes the proof. ■

We now do an example to illustrate the usefulness of supermodular games. Consider the classic Bertrand game, where two firms are producing the same good. Each firm chooses a price: say p_1 for firm 1 and p_2 for firm 2. Suppose the prices lie in $[0, M]$ for some positive real number M. The demand for firm i for a pair of prices p_i, p_j is given by

$$D_i(p_i, p_j) = g_i(p_i) + p_j,$$

where g_i some continuous and decreasing function of p_i. If the marginal cost of production is c for both the firms, the utility of firm i is

$$u_i(p_i, p_j) = (p_i - c)(g(p_i) + p_j).$$

Note that u_i is continuous and supermodular in p_i for every p_j (supermodularity is vacuously satisfied). For increasing differences, we pick $p_i' > p_i$ and $p_j' = p_j + \delta$ for $\delta > 0$. So, we have

$$\begin{aligned} u_i(p_i', p_j') - u_i(p_i, p_j') &= (p_i' - c)(g_i(p_i') + p_j') - (p_i - c)(g_i(p_i) + p_j') \\ &= (p_i' - c)(g_i(p_i') + p_j + \delta) - (p_i - c)(g_i(p_i) + p_j + \delta) \\ &= (p_i' - c)\delta - (p_i - c)\delta + u_i(p_i', p_j) - u_i(p_i, p_j) \\ &= (p_i' - p_i)\delta + u_i(p_i', p_j) - u_i(p_i, p_j) \\ &\geq u_i(p_i', p_j) - u_i(p_i, p_j). \end{aligned}$$

By Theorem 13, a pure strategy Nash equilibrium exists in this Bertrand game.

The existence of pure strategy equilibrium in supermodular game is an interesting result because it does not require some concavity and continuity assumptions of Theorem 9. However, there are even more striking results one can establish for supermodular games. Below, we show how we can compute a pure strategy Nash equilibria of a supermodular game.

Theorem **14** *Suppose $\Gamma \equiv (N, \{S_i\}_i, \{u_i\}_i)$ is a supermodular game and u_i is continuous in s_{-i} for every $i \in N$. Suppose $\bar{s}$ and $\underline{s}$ are the greatest and least pure strategy profiles that survive iterated strict dominance in pure strategies. Then, $\bar{s}$ and $\underline{s}$ are the greatest and the least Nash equilibrium profiles of Γ.*

Proof: We iterate through the best response map by successively eliminating strictly dominated strategies. Initially, we set $S_i^0 = S_i$ for all $i \in N$. Let $S^0 \equiv (S_1^0, \ldots, S_n^0)$. Denote by $s^0 \equiv (s_1^0, \ldots, s_n^0)$ the greatest element of the lattice S.

Now, for every $i \in N$, choose

$$s_i^1 = \bar{B}_i(s_{-i}^0) \quad \text{and} \quad S_i^1 = \{s_i \in S_i^0 : s_i \leq s_i^1\}.$$

The first claim is that any $s_i > s_i^1$ (i.e., $s_i \notin S_i^1$) is strictly dominated by s_i^1. To see this, for all $s_{-i} \in S_{-i}$, we have

$$\begin{aligned} u_i(s_i, s_{-i}) - u_i(s_i^1, s_{-i}) &\leq u_i(s_i, s_{-i}^0) - u_i(s_i^1, s_{-i}^0) \\ &< 0, \end{aligned}$$

where the first inequality followed from increasing differences and the second strict inequality from the fact that $s_i^1 = \bar{B}_i(s_{-i}^0)$ and $s_i \notin B_i(s_{-i}^0)$.

Note that $s_i^1 \leq s_i^0$. We now inductively define a sequence. Having defined S_i^{k-1} and s_i^{k-1} for all $i \in N$, we define

$$s_i^k = \bar{B}_i(s_{-i}^{k-1}) \quad \text{and} \quad S_i^k = \{s_i \in S_i^{k-1} : s_i \leq s_i^k\}.$$

As before, we note that for all $s_i \in S_i^{k-1} \setminus S_i^k$, s_i is strictly dominated by s_i^k for all strategies $s_{-i} \in S_{-i}^{k-1}$. To see this, pick $s_i \in S_i^{k-1}$ and $s_{-i} \in S_{-i}^{k-1} \setminus S_{-i}^k$, and note that

$$\begin{aligned} u_i(s_i, s_{-i}) - u_i(s_i^k, s_{-i}) &\leq u_i(s_i, s_{-i}^{k-1}) - u_i(s_i^k, s_{-i}^{k-1}) \\ &< 0, \end{aligned}$$

where the first inequality followed from increasing differences and the second strict inequality from the fact that $s_i^k = \bar{B}_i(s_{-i}^{k-1})$ and $s_i \notin B_i(s_{-i}^{k-1})$. Thus, $\{S_i^k\}_i$ defines a new game where players eliminate strictly dominated strategies from the previous stage game with strategies $\{S_i^{k-1}\}_i$.

Further, note that if $s^k \leq s^{k-1}$, then for every $i \in N$,

$$s_i^{k+1} = \bar{B}_i(s_{-i}^k) \leq \bar{B}_i(s_{-i}^{k-1}) = s_i^k,$$

where the inequality followed from the monotone comparative statics result of Topkis. This implies that the sequence $\{s^k\}_k$ is a non-increasing sequence which is bounded from below. Hence, it has a limit point - denote this limit as $\bar{s}$.

We now show that $\bar{s}$ is a Nash equilibrium. To see this, we show that for all $i \in N$ and for all $s_i \in S_i$, we have

$$u_i(s_i^{k+1}, s_{-i}^k) \geq u_i(s_i, s_{-i}^k).$$

First $u_i(s_i^1, s_{-i}^0) \geq u_i(s_i, s_{-i}^0)$ for all $s_i \in S_i$. Now assume that $u_i(s_i^k, s_{-i}^{k-1}) \geq u_i(s_i, s_{-i}^{k-1})$ for all $s_i \in S_i$. Now, choose $s_i \in S_i \setminus S_i^k$. By definition $s_i^k \leq s_i$. Since $s_{-i}^k \leq s_{-i}^{k-1}$ increasing differences imply that

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

But, $s_i^{k+1} = \bar{B}_i(s_{-i}^k)$. Hence, $u_i(s_i^{k+1}, s_{-i}^k) \geq u_i(s_i^k, s_{-i}^k) \geq u_i(s_i, s_{-i}^k)$. This shows that for all $s_i \in S_i \setminus S_i^k$,

$$u_i(s_i^{k+1}, s_{-i}^k) \geq u_i(s_i, s_{-i}^k).$$

Since $s_i^{k+1} = \bar{B}_i(s_{-i}^k)$, we know that for all $s_i \in S_i^k$,

$$u_i(s_i^{k+1}, s_{-i}^k) \geq u_i(s_i, s_{-i}^k).$$

This completes the argument that for all $s_i \in S_i$, we have

$$u_i(s_i^{k+1}, s_{-i}^k) \geq u_i(s_i, s_{-i}^k).$$

Taking limit, and using the fact that u_i is continuous, we get

$$u_i(\bar{s}_i, \bar{s}_{-i}) \geq u_i(s_i, \bar{s}_{-i}).$$

Hence, $\bar{s}$ is a Nash equilibrium of the original game.

Suppose there is another Nash equilibrium $\bar{s}'$ such that $\bar{s}'_i > \bar{s}_i$ for some i. Then, there is a stage k of iterated elimination with s^k as the greatest strategy profile. An s^k can be chosen such that $\bar{s}'_i > s^k_i > \bar{s}_i$. We know that a Nash equilibrium of the original game is also a Nash equilibrium of this game (strict iterated elimination preserves the set of Nash equilibrium - Theorem 5). But $\bar{s}'_i$ is strictly dominated in this game. Hence, it cannot be part of a Nash equilibrium. This is a contradiction.

Similarly, we can start with $s^0 \equiv (s^0_1, \ldots, s^0_n)$ as the least element in S and identify the limit point of an non-decreasing sequence as $\underline{s}$. Using a similar proof technique, we can show that $\underline{s}$ is also a Nash equilibrium. This will correspond to the least Nash equilibrium. ■

We now apply the idea of Theorem 14 to a Bertrand game. Suppose there are two firms producing the same good. Both the firms choose prices in $[0,1]$. Depending on prices p_1, p_2, the demand of firm 1 is

$$D_i(p_1, p_2) = 1 - 2p_i + p_j.$$

Suppose the marginal cost is zero for both the firms. Then, utility of firm i is

$$u_i(p_1, p_2) = p_i(1 - 2p_i + p_j).$$

Set $S^0_i = [0,1]$. The greatest element strategy profile is $(1,1)$. If one firm sets price equal to 1, then $u_i(p_i, 1) = 2p_i(1 - p_i)$. There is a unique best response to it - $p_i = \frac{1}{2}$. Now, we set $S^1_i = [0, \frac{1}{2}]$ and $s^1_i = \frac{1}{2}$ for each i. Then, $u_i(p_i, \frac{1}{2}) = p_i(\frac{3}{2} - 2p_i)$. This gives a unique best response of $\frac{3}{8}$. So, we set $S^2_i = [0, \frac{3}{8}]$ and $s^2_i = \frac{3}{8}$. So, we get a sequence $(1, \frac{1}{2}, \frac{3}{8}, \frac{11}{32}, \ldots)$. Note that this sequence is $(1, \frac{1}{2}, \frac{1}{4} + \frac{1}{4}\frac{1}{2}, \frac{1}{4} + \frac{1}{4}\frac{3}{8}, \ldots)$. Hence, the k-th term is

$$\frac{1}{4} + \frac{1}{16} + \ldots + \frac{1}{4^k} + \frac{s^0_i}{4^k}$$

As k tends to infinity, this becomes $\frac{1}{3}$. Hence, the greatest Nash equilibrium is $(\frac{1}{3}, \frac{1}{3})$.

Now, we start from the least strategy profile $(0,0)$. Then, $u_i(p_i, 0) = p_i(1 - 2p_i)$. Hence, the unique best response is $p_i = \frac{1}{4}$. So, $S^1_i = [\frac{1}{4}, 1]$ and $s^1_i = \frac{1}{4}$ for each i. Then, $u_i(p_i, \frac{1}{4}) = p_i(\frac{5}{4} - 2p_i)$. Unique best response is $\frac{5}{16}$. Hence, we get a sequence $(0, \frac{1}{4}, \frac{1}{4} + \frac{1}{4}\frac{1}{4}, \ldots)$. Hence, the k-th term is

$$\frac{1}{4} + \frac{1}{16} + \ldots + \frac{1}{4^k} + \frac{s^0_i}{4^k},$$

whose limit is the same $\frac{1}{3}$. Hence, the least Nash equilibrium is also $(\frac{1}{3}, \frac{1}{3})$. So, $(\frac{1}{3}, \frac{1}{3})$ is the only Nash equilibrium.

Important Note: As we saw in this example, the strategy space of players in many games is a subset of $\mathbb{R}$. In that case, the every compact subset of $\mathbb{R}$ will be a compact lattice. Hence, the lattice requirement is vacuously satisfied. Further, supermodularity is also vacuously satisfied. The only restriction that supermodular games impose is increasing differences in (s_i, s_{-i}) and continuity with respect to s_i. Theorem 14 also makes use of continuity in s_{-i}.

9 Correlated Equilibrium

Consider the following game - usually called the game of "chicken". There are two players - $N = \{1, 2\}$. Player 1 has two pure strategies $S_1 = \{T, B\}$ and Player 2 has two pure strategies $S_2 = \{L, R\}$. The payoffs are shown in Table 20. The story that accompanies this game is that two drivers are racing towards each other on a single lane. Each driver can either stay on or move away from the road. If both move away, then they get a payoff of 6 each. If both stay on, then they get a payoff of 0. If one of them stays on but the other moves away, then the one who stays on gets a payoff of 7 but the other one gets a payoff of 2.

	L	R
T	$(6,6)$	$(2,7)$
B	$(7,2)$	$(0,0)$

Table 20: Game of chicken

There are three Nash equilibria of this game: $(T, R), (B, L), \left(\frac{2}{3}T + \frac{1}{3}B, \frac{2}{3}L + \frac{1}{3}R\right)$. Now, consider the following "extended" game. There is an outside observer. The observer recommends each player *privately* a pure strategy to play. Note that no player observes the recommendation of the other player. Given his own recommended strategy, a player forms belief about the recommended strategy of the other player, assuming that the other player follows the recommendation. He follows his recommended strategy if and only if it is a best response given his belief about other player's recommended strategy.

Two natural confusions arise - (a) How does the observer recommend? and (b) How do the players form beliefs? It is assumed that the observer has access to a randomization device which is public, i.e., players know the distribution from which the recommendations are derived. Given the distribution of recommendation, players form beliefs by using Bayes' rule - they compute conditional probabilities.

In the game in Table 20, suppose the observer recommends pure strategy profiles in Nash equilibrium: (T, R) and (B, L) with probability p and $(1-p)$. Then, given his recommended strategy each player can uniquely infer the recommended strategy of the other player. Player 1 gets a recommendation of T means, Player 2 must have received a recommendation of R. So, Player 1 forms a belief that Player 2 plays R with probability 1. But (T, R) is a Nash equilibrium means, T is a best response to R. A similar logic shows that Player 1 will also accept B if it is recommended. Same argument applies to Player 2. Hence, *any* convex combination of pure strategy Nash equilibrium can be sustained as a *correlated* equilibrium of this extended game. In particular $p(T, R)+(1-p)(B, L)$ for any p is an equilibrium of this game. The set of payoffs that can be obtained are convex combination of $(7, 2)$ and $(2, 7)$.

Can we get other equilibrium? Suppose the observer recommends $(T, R), (B, L)$, and (T, L) with probability $\frac{1}{3}$ each. Then, if Player 1 observes T as a recommendation, then he can infer that Player 2 will have R as recommendation with probability $\frac{1}{2}$ and L as recommendation with probability $\frac{1}{2}$. Hence, he forms belief that Player 2 plays $\frac{1}{2}R + \frac{1}{2}L$. Is T a best response of Player 1 to this strategy? Playing T gives him 4 and playing B gives him 3.5. So, T is a best response, and Player 1 accepts the recommendation. If Player 1 receives B as a recommendation, then he forms a belief that Player 2 must receive L as recommendation. Since (B, L) is a Nash equilibrium, B is a best response to L. For Player 2, if he receives R as a recommendation, then he infers Player 1 must have received T and that being a Nash equilibrium, he accepts the recommendation. If Player 2 receives L as a recommendation, then he believes Player 1 must have received T as recommendation with probability $\frac{1}{2}$ and B as recommendation with probability $\frac{1}{2}$. Indeed, L is a best response to this strategy. Hence, both the players agree to accept the recommendations of the observer using this randomization device. The equilibrium payoff of both players from this is $(5, 5)$ which could not be obtained if we just randomize over Nash equilibria. Hence, an observer using a public randomizing device allows players to get payoff outside the convex hull of Nash equilibrium payoffs.

As the previous example illustrated, using public randomization allowed the players to avoid the worst payoff $(0, 0)$ by putting zero probability on that profile. This is impossible in a mixed strategy - independent randomization. To be able to play strategy profile (T, R), Player 2 must play R with some probability and that will mean playing (B, R) with some probability.

9.1 Correlated Strategies

A crucial assumption in mixed strategies is that players randomize independently. Each of them have access to a randomizing device (say, a coin to toss or a random number generating computer program) and these devices are independent. In some circumstances, players may have access to the same randomizing device. For instance, players observe some common event in the nature and decide to play their strategies based on this common event - say weather in a particular area.

Consider the same example in Table 5. Suppose Player 1 plays A and Player 2 plays a if it rains and Player 1 plays B and Player 2 plays b if it does not rain. Suppose the probability of rain is $\frac{1}{2}$. This means that the strategy profiles (A, a) and (B, b) is played with probability $\frac{1}{2}$ each but other strategy profiles are played with zero probability. There is strong correlation between the strategies played by both the players. Formally, a correlated strategy ρ is a map $\rho : S \to [0, 1]$ with $\sum_{s \in S} \rho(s) = 1$. The correlated strategy discussed above is shown in Table 21.

	a	b
A	$\frac{1}{2}$	0
B	0	$\frac{1}{2}$

Table 21: Correlated strategies - probability of all pure strategy profiles

An important fact to note is that a correlated strategy may not be obtained from a mixed strategy. For instance, consider the correlated strategy in Table 21. If Player 1 and Player 2 play mixed strategies that generates the same distribution over strategy profile as in Table 21, then either 1 must put zero weight on A or 2 must put zero weight on b. This implies that we cannot get the distribution in Table 21.

In general, the correlated strategy $\rho \in \Delta\Big(\prod_{i \in N} S_i\Big)$ and a mixed strategy $\sigma \in \prod_{i \in N} \Delta S_i$. Every mixed strategy generates a correlated strategy. Hence, the set of distributions over strategy profiles that can be obtained by correlated strategy is larger than the set of distributions generated by mixed strategies. Player i evaluates a correlated strategy ρ using expected utility:

$$U_i(\rho) = \sum_{s \in S} u_i(s)\rho(s).$$

9.2 Formal Definition

We will now define a correlated equilibrium based on the notion of correlated strategies. Let $\Gamma \equiv (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ be a finite strategic form game. To avoid confusion, we will refer to strategies in S_i for each i as **actions** of Player i.

For every probability vector (correlated strategy) p over $S \equiv S_1 \times \ldots \times S_n$, an **extended game of** Γ is defined as:

- An outside observer chooses a profile of pure actions $s \in S$ using the correlated strategy p.

- It reveals to each player i, his recommendation s_i but not s_{-i}.

- Each player i chooses an action $s'_i \in S_i$ after receiving his recommendation.

We denote this extended game as $\Gamma(p)$. Hence, formally a strategy in this extended game is a different object compared to the strategy in a strategic form game.

Definition **16** *A strategy of Player i in the extended game $\Gamma(p)$ is a map $\psi_i : S_i \to S_i$, i.e., giving an action for every possible recommended action.*

One strategy is the **obedient** strategy map - for every $s_i \in S_i$, $\psi^*_i(s_i) = s_i$ for each i. Below, we show the mathematical implication of the fact that ψ^* is a Nash equilibrium of $\Gamma(p)$. What does it mean to say that ψ^* is a Nash equilibrium of $\Gamma(p)$? It means that given that everyone else is playing the strategy ψ^*, payoff of an agent i is maximized by playing ψ^*. Since there is uncertainty about the recommendation of other players, payoff of agent i has to be computed by taking expectation over all possible recommendations.

Theorem **15** *The strategy profile ψ^* is a Nash equilibrium of $\Gamma(p)$ if and only if for every $i \in N$, for every $s_i, s'_i \in S_i$, we have*

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

Proof: The strategy profile ψ^* is a Nash equilibrium if and only if no player i can unilaterally deviate from his recommended action. If Player i receives recommendation s_i, then his conditional belief that other players received recommendation s_{-i} is

$$\frac{p(s_i, s_{-i})}{\sum_{t_{-i}} p(s_i, t_{-i})},$$

where the denominator is positive from the fact that $p(s_i, s_{-i}) > 0$. Then, his expected payoff from following $\psi^*_i(s_i) = s_i$ (given others are following recommendation) is

$$\sum_{s_{-i} \in S_{-i}} \frac{p(s_i, s_{-i})}{\sum_{t_{-i}} p(s_i, t_{-i})} u_i(s_i, s_{-i}).$$

His expected payoff from playing s'_i (given others are following recommendation) is

$$\sum_{s_{-i} \in S_{-i}} \frac{p(s_i, s_{-i})}{\sum_{t_{-i}} p(s_i, t_{-i})} u_i(s'_i, s_{-i}).$$

Since the denominator is positive, we can say that s_i is best response if and only if

$$\sum_{s_{-i}\in S_{-i}} p(s_i, s_{-i})u_i(s_i, s_{-i}) \geq \sum_{s_{-i}\in S_{-i}} p(s_i, s_{-i})u_i(s_i', s_{-i}).$$

■

This leads to the definition of a correlated equilibrium.

Definition **17** *A correlated strategy p over S is a* **correlated equilibrium** *if for every $i \in N$, for every $s_i, s_i' \in S_i$,*

$$\sum_{s_{-i}\in S_{-i}} p(s_i, s_{-i})u_i(s_i, s_{-i}) \geq \sum_{s_{-i}\in S_{-i}} p(s_i, s_{-i})u_i(s_i', s_{-i}).$$

In other words, a correlated strategy p over S is a **correlated equilibrium** *if the strategy profile ψ^* is a Nash equilibrium of the extended game $\Gamma(p)$.*

This shows that the set of correlated equilibria are solutions to a finite set of inequalities in a finite game. As result, they form a convex and compact set (in particular, a *polytope*, defined by a system of linear inequalities).

Every Nash equilibrium σ^* of Γ induces a probability distribution p_{σ^*}, where for every $(s_1, \ldots, s_n)$,

$$p_{\sigma^*}(s_1, \ldots, s_n) = \sigma_1^*(s_1) \times \ldots \times \sigma_n^*(s_n).$$

Below, we formally show that every Nash equilibrium induces a distribution over strategy profiles that is a correlated equilibrium.

Theorem **16** *For every Nash equilibrium σ^* of Γ, the induced correlated strategy p_{σ^*} is a correlated equilibrium of $\Gamma(p_{\sigma^*})$.*

Proof: Note that $p_{\sigma^*}(s) > 0$ if and only if for every $i \in N$, s_i is in the support of σ^*. Pick agent i, $s_i, s_i' \in S_i$. We see that

$$\sum_{s_{-i}\in S_{-i}} p_{\sigma^*}(s_i, s_{-i})u_i(s_i, s_{-i}) = \sum_{s_{-i}\in S_{-i}} \sigma_1^*(s_1) \times \ldots \times \sigma_n^*(s_n)u_i(s_i, s_{-i}) = \sigma_i^*(s_i)\mathcal{U}_i(s_i, \sigma_{-i}^*).$$

Further,

$$\sum_{s_{-i}\in S_{-i}} p_{\sigma^*}(s_i, s_{-i})u_i(s_i', s_{-i}) = \sum_{s_{-i}\in S_{-i}} \sigma_1^*(s_1) \times \ldots \times \sigma_n^*(s_n)u_i(s_i', s_{-i}) = \sigma_i^*(s_i)\mathcal{U}_i(s_i', \sigma_{-i}^*).$$

Since s_i is in the support of Nash equilibrium at σ^*, it implies that $\sigma_i^*(s_i) > 0$. Further, by the indifference lemma, s_i is a best response to σ_{-i}^*, and hence,

$$u_i(s_i, \sigma_{-i}^*) \geq u_i(s_i', \sigma_{-i}^*).$$

This gives us that

$$\sum_{s_{-i} \in S_{-i}} p_{\sigma^*}(s_i, s_{-i}) u_i(s_i, s_{-i}) \geq \sum_{s_{-i} \in S_{-i}} p_{\sigma^*}(s_i, s_{-i}) u_i(s_i', s_{-i}),$$

as required. ■

10 Correlated Rationalizability

In this section, we introduce the idea of correlated rationalizability. Here, we entertain beliefs of players that allow strategies of other players to be correlated. Formally, belief of player i in a game is a probability distribution over S_{-i} - it specifies a probability of each of the strategy profile s_{-i} being played. Such probabilities need not be computed using independence of strategies of other players. So, belief of player i is a map $\mu_i : S_{-i} \rightarrow [0, 1]$, with $\sum_{s_{-i}} \mu_i(s_{-i}) = 1$. Note that a mixed strategy profile σ induces a belief for every player i: $\mu_i(s_{-i}) := \times_{j \neq i} \sigma_j(s_j)$ for all s_{-i}. These beliefs are generated by independent probabilities of each player $j \neq i$. In general, beliefs may allow correlations.

A strategy $s_i \in S_i$ is a **best response with respect to** a belief μ_i if

$$\sum_{s_{-i}} u_i(s_i, s_{-i}) \mu_i(s_{-i}) \geq \sum_{s_{-i}} u_i(s_i', s_{-i}) \mu_i(s_{-i}) \ \forall \ s_i' \in S_i.$$

Definition 18 *A strategy $s_i \in S_i$ is* **rationalizable** *in the strategic form game $(N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ if for every $j \in N$ there is a strategy set $Z_j \subseteq S_j$ such that*

- $s_i \in Z_i$
- *every $s_j \in Z_j$ is a best response (over all strategies in S_j) with respect to some belief μ_j of Player j whose support is a subset of Z_{-j}.*

The set of all strategies that are rationalizable for a player are called his **rationalizable** *strategies.*

Note that the strategies in Z_j for each j are only used to form beliefs - strategy profiles involving strategies outside them get zero probability. The best response is with respect to all the strategies.

Consider the example in Table 22. $(\{A\}, \{a\})$ is not a set of rationalizable strategies. This is because here there is only one degenerate belief: Player 1 must believe Player 2 plays a and Player 2 must believe that Player 1 plays A. But a is not a best response if Player 1 plays A. On the other hand, $(\{A, C\}, \{a, b\})$ is a set of rationalizable strategies. How do we verify this? A is a best response if a is played and C is a best response if b is played. Similarly, for Player 2, a is a best response if C is played and b is a best response if A is played.

	a	b	c
A	$(6,2)$	$(0,6)$	$(4,4)$
B	$(2,12)$	$(4,3)$	$(2,5)$
C	$(0,6)$	$(10,0)$	$(2,2)$

Table 22: Two Player Game

Also, note that if a set of strategies $S_i' \subseteq S_i$ is rationalizable with respect to $\{Z_j\}_j$ and another set of strategies $S_i'' \subseteq S_i$ is rationalizable with respect to $\{Z_j'\}_j$, then $S_i' \cup S_i''$ is also rationalizable with respect to $\{Z_j \cup Z_j'\}_j$. Hence, the set of rationalizable strategies is the largest collection of $\{Z_j\}_j$ that can be rationalized.

An immediate claim is the following.

Lemma **5** *Every strategy in the support of a Nash equilibrium is rationalizable.*

Proof: Suppose s_i is a strategy of Player i in the support Nash equilibrium σ^*. Now for every j, Z_j are all the strategies in the support of the Nash equilibrium σ^* and the belief μ_j is the product

$$\times_{k \neq j} \sigma_k^*(s_k) \ \forall \ s_{-j}.$$

By the definition of Nash equilibrium and the indifference lemma, each s_j in the support of σ_j^* is a best response of j with respect to the belief μ_j. ∎

One can also show that strategies used with positive probability in a correlated equilibrium are also rationalizable. In general, finding the set of rationizable strategies can be quite cumbersome. Below, we provide an easy method with the help of a cute result.

10.1 Never Best Responses and Strict Domination

We define the notion of never best response using correlated beliefs now.

Definition 19 *A strategy $s_i \in S_i$ is a* **never-best response** *in the strategic form game $(N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N})$ if it is not a best response with respect to any belief μ_i.*

Clearly, if a strategy is a never-best response, it cannot be rationalized. We show that for finite games, the set of rationalizable strategies can be found by iteratively eliminating strictly dominated strategies.

We remind the definition of a strictly dominated strategies.

Definition 20 *A strategy $s_i \in S_i$ is a* **strictly dominated strategy** *in the strategic form game $(N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N})$ if there is a mixed strategy of $\sigma_i \in \Delta S_i$ such that*

$$u_i(s_i, s_{-i}) < u_i(\sigma_i, s_{-i}) \;\; \forall \; s_{-i}.$$

We prove the following.

Theorem 17 *A strategy of a player in a strategic game is a never-best response if and only if it is strictly dominated.*

Proof: Clearly, every strictly dominated strategy is a never-best response strategy. For the other direction, fix a player j in a strategic form game $\Gamma \equiv (N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N})$ and a strategy $s_j^* \in S_j$. Consider the reduced game in which there are two players j and $-j$. The set of strategies available to Player j is $S'_j := S_j \setminus \{s_j^*\}$ and to Player $-j$ is S_{-j}. The utility of Player j is:

$$v_j(s_j, s_{-j}) = u_j(s_j, s_{-j}) - u_j(s_j^*, s_{-j}).$$

The payoff to Player $-j$ is negative of payoff to Player j - hence, it is a zero-sum game. Denote this game as Γ'.

Now, notice that a mixed strategy of Player $-j$ corresponds to a correlated belief of Player j in the original game. Hence, strategy s_j^* is a never-best response implies for every mixed strategy σ_{-j} of Player $-j$, there exists a strategy s_j such that $v_j(s_j, \sigma_{-j}) > 0$. This implies that

$$\min_{\sigma_{-j}} \max_{s_j} v_j(s_j, \sigma_{-j}) > 0.$$

The maximization in the second part can also be done with respect to all the mixed strategies, and by linearity of expected utility the value of maximum will not change. Hence, s_j^* is a never-best response implies

$$\min_{\sigma_{-j}} \max_{\sigma_j} v_j(\sigma_j, \sigma_{-j}) > 0.$$

But by the max-min value theorem for zero-sum games, this is equivalent to

$$\max_{\sigma_j} \min_{\sigma_{-j}} v_j(\sigma_j, \sigma_{-j}) > 0.$$

This implies there is a mixed strategy σ_j^* such that $v_j(\sigma_j^*, \sigma_{-j}) > 0$ for all σ_{-j}. This implies that $u_j(\sigma_j^*, \sigma_{-j}) > u_j(s_j^*, \sigma_{-j})$ for all σ_{-j}. Since σ_{-j} is a belief over the strategies of players other than j, this means that σ_j^* dominates s_j^* for all possible beliefs of Player j. Hence, s_j^* is strictly dominated. ■

Notice that this equivalence is only valid if we allow for correlated beliefs - of course, for two-player games these correlated belief is same as independent belief. But, for more than two player games, using only independent beliefs does not lead to Theorem 17.

10.2 A Formal Definition of Iterated Elimination Procedure

We now formally introduce the notion of iterated elimination of strictly dominated strategies.

Definition **21** *The set $X \subseteq S$ of strategy profiles survives* **iterated elimination of strictly dominated strategies** *if $X \equiv \times_{j \in N} X_j$ and there is a collection $(\{X_j^t\}_{j \in N})_0^T$ of sets that satisfy for each $j \in N$ the following:*

- *$X_j^0 = S_j$ and $X_j^T = X_j$,*
- *$X_j^{t+1} \subseteq X_j^t$ for each $t < T$,*
- *for each $t < T$, every strategy in $X_j^t \backslash X_j^{t+1}$ is strictly dominated in the game $(N, \{X_i^t\}_i, \{u_i^t\}_i)$, where u_i^t is the restriction of u_i to strategy profiles in this game.*
- *No strategy in X_j^T is strictly dominated.*

Note that the definition does not require you to eliminate **all** the strictly dominated strategies in a stage of elimination.

Theorem **18** *The set of rationalizable strategies of a player is the set of strategies available after iterated elimination of strictly dominated strategies.*

Proof: Suppose s_i is rationalizable for Player i. Let $\{Z_j\}_j$ be the profile of strategies that supports s_i. When we run the iterated elimination of strictly dominated strategies, in Stage t, each strategy in Z_j is a best response to some belief over Z_{-j}, and by Theorem 17, it is not strictly dominated. Hence, $Z_j \subseteq X_j^t$ for each t. So, $s_i \in X_i$.

Now, for the converse, pick $s_i \in X_i$. By definition every strategy in X_i is not strictly dominated in the game with strategy sets X_i. So, by Theorem 17, every strategy in X_i is a best response among strategies in X_i to some belief over X_{-i}. We need to argue that every strategy in X_i is a best response among strategies in S_i to some belief over X_{-i}. Suppose strategy $s_i \in X_i$ is not a best response among strategies in S_i to some belief over X_{-i}. Then, there must exist some stage t where s_i is a best response among X_i^t to some belief over X_{-i} but it is not a best response among X_i^{t+1} to some belief over X_{-i}. Then, there is a strategy $s'_i \in X_i^{t+1} \setminus X_i^t$ that is a best response among X_i^{t+1} to some belief over X_{-i}. By Theorem 17, such a strategy is not strictly dominated. Hence, s'_i cannot be eliminated in this stage, which is a contradiction. ■

Since the procedure we defined for iterated elimination did not specify any order of elimination, this also implies that order of elimination of strictly dominated strategies does not matter.

11 Bayesian Games

Often, the strategic form game depends on some external factor. These factors may be known to some agents with varying certainty. To make ideas clear, consider a situation in which two agents are deciding where to meet. Each agent privately observes the weather in his city but does not know the weather of the other agent's city. Based on the weather in the city, an agent has a set of *actions* available to him, and his utility will depend on the weather in both the cities and the actions chosen by both the agents. Here, the weather in each city is a *signal* that is privately observed by the player. The signal determines the action set of the strategic game. The utility in the strategic form game is determined by the signals realized by all the agents and the actions taken.

The kind of uncertainty in this example is about the weather in the cities. Each agent uses a *common prior* to evaluate uncertainty using expected utility. In this example, there is a probability distribution about the weather in both the cities. Note that since an agent only observes weather in his own city, he can use Bayes rule to update the conditional probabilities.

Note that the strategy of a player and his payoff functions are complicated objects in this environment because (a) it depends on the signals players receive and (b) there is uncertainty about the signals of other players. Harasanyi was the first to formally define an analogue of a strategic game in this uncertain environment.

Definition **22** *A* **game of incomplete information** *is defined by* $(N, \{T_i\}_{i\in N}, p, \{\Gamma_t\}_{t\in \times_{i\in N} T_i})$, *where*

- *N is a finite set of players,*

- *T_i is a set of* **types** *(signals) for player i, and $T = \times_{i \in N} T_i$ is the set of type vectors,*

- *p is a probability distribution (***belief or prior***) over T with the restriction that $\sum_{t_{-i} \in T_{-i}} p(t_i, t_{-i}) > 0$ for each $t_i \in T_i$ and for each $i \in N$,*

- *For each $t \in T$, a strategic form game $\Gamma_t \in S \equiv (N, \{A_i(t_i)\}_{i \in N}, \{u_i(t)\}_{i \in N})$.*

A game of incomplete information proceeds in a sequence where some of the associated uncertainties are resolved.

- The type vector $t \in T$ is chosen (by nature) using the probability distribution p.

- Each player $i \in N$ observes his own type t_i but does not know the types of other agents.

- After observing their types, each player i plays an action $a_i \in A_i(t_i)$.

- Each player i receives an utility equal to $u_i(t, (a_1, \ldots, a_n))$ when the type vector realized is t and the action vector is $(a_1, \ldots, a_n)$.

Because of uncertainty, the players do not even know the action set available to other players. So, they do not know which strategic form game is being played. Note that the action set depends on the type of the player. Further, the utility depends on the type vector realized and the actions taken by all the players.

Strategies in such games are complicated objects. To remind, a strategy must describe the action to be taken for every possible contingency. Hence, here also, a strategy must describe what action to take for every signal/type that the player receives.

A **pure strategy** of Player i in a Bayesian game is a map $s_i : T_i \to \cup_{t_i \in T_i} A_i(t_i)$ such that $s_i(t_i) \in A_i(t_i)$ for all $t_i \in T_i$. Thus, a pure strategy prescribes one action for every type. However, one can also chose a probability distribution over the set of actions - this will be the analogue of the mixed strategy. [4] This is called a *behavior strategy*. Formally, a mixed strategy of Player i is a map $\sigma_i : T_i \to \cup_{t_i \in T_i} \Delta A_i(t_i)$ such that for every $t_i \in T_i$, $\sigma_i(t_i) \in \Delta A_i(t_i)$.

What is the payoff of Player i from a strategy profile σ? There are two ways to think about it: ex-ante payoff, which is computed before realization of the type, and interim payoff, which is computed after realization of the type. Ex-ante payoff from strategy profile σ is

$$U_i(\sigma) := \sum_{t \in T} p(t) U_i(t; \sigma) = \sum_{t \in T} p(t) \sum_{a \in A(t)} u_i(t, a) \sigma_1(t_1; a_1) \times \ldots \times \sigma_n(t_n; a_n) \equiv \sum_{t \in T} p(t) U_i(t; \sigma).$$

[4] We will come back to this point later.

The interim payoffs are computed by updating beliefs after realizing the types. In particular, once Player i knows his type to be $t_i \in T_i$, he computes his conditional probabilities as follows. For every $t_{-i} \in T_{-i}$,

$$p(t_{-i}|t_i) := \frac{p(t_i, t_{-i})}{\sum_{t'_{-i} \in T_{-i}} p(t_i, t'_{-i})},$$

where we will denote $p(t_i) \equiv \sum_{t'_{-i} \in T_{-i}} p(t_i, t'_{-i})$ and note that it is positive by our assumption. The interim payoff of Player i with type t_i from a strategy profile σ is thus

$$U_i(\sigma|t_i) := \sum_{t_{-i} \in T_{-i}} p(t_{-i}|t_i) U_i(t; \sigma).$$

If the beliefs are independent, then observing own type gives no extra information to the players. Hence, no updating of prior belief is required by the players.

Note: The above expressions are for finite type spaces, but similar expressions (using integrals) can also be written with infinite type spaces. We will discuss them when we do particular applications with infinite type spaces later.

11.1 An Example

We give an informal description of a Bayesian game before describing the equilibrium concepts. This Bayesian game is in the context of an auction - a popular subfield of economic theory where game theory has been applied successfully in practice and theory. There is an indivisible object for sale to a set of buyers (players). Each buyer has a value v_i for the object. The value is the type of the object, and hence, every buyer only knows his own value but not the value of others. The values are drawn using a distribution p over the set of all value profiles.

The set of actions available to a player in this game is the set of all non-negative real numbers. Such actions are called **bids** in auction literature. A bid specifies the amount a buyer is willing to pay. The buyer with the highest bid (ties broken with equal probability) wins the object. If a buyer i with value v_i wins the object by bidding b_i, then his utility is $v_i - b_i$ times the probability of winning. A losing buyer gets zero utility. Note that the amount a biddder bids may depend on his type. Whether a buyer wins or not depends on the bids of all the players. The utility of a player depends on this probability of winning and his own type.

To be a little more specific, let us study strategies which are commonly referred to as **symmetric monotone bidding strategies**. Assume that type space $T_i = \mathbb{R}_+$. A symmetric monotone strategy is a map $b : T_i \to \mathbb{R}_+$. Note that every bidder is using the same bid function (strategy). We further assume that b is strictly increasing and differentiable. Suppose each bidder draws his type independently from T_i using a distribution F (identical distribution for all bidders). If types of all the buyers are $(v_1, \ldots, v_n)$, the the probability of this type vector is $F(v_1) \times \ldots \times F(v_n)$. Since bids are monotone functions, a bidder with type v_i wins when everyone follows this strategy if $v_i > \max_{j \neq i} v_j$. The probability of this event is $F(v_i)^{n-1}$. The interim payoff of Player i with type v_i from this strategy is

$$F(v_i)^{n-1}(v_i - b_i(v_i)).$$

Hence, the ex-ante payoff from of Player i with type v_i from this strategy is

$$\int_{v_i} F(v_i)^{n-1}(v_i - b_i(v_i))f(v_i)dv_i,$$

where f is the density function. Note that this expression is independent of the uncertainty about other players' types. This is because of the particular strategies (symmetric and monotone) strategies that we are considering.

11.2 Bayesian Equilibrium

As we saw, there are two points at which a player may evaluate his utility: ex-ante or interim. Depending on that the notion of equilibrium can be defined. The ex-ante notion coincides with the idea of a Nash equilibrium.

Definition **23** *A strategy profile σ^* is a* **Nash equilibrium** *in a Bayesian game if for each player i and each strategy σ_i,*

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

There is also an interim way of defining the equilibrium. This is called the Bayesian equilibrium, and is the common way of defining equilibrium in Bayesian games.

Definition **24** *A strategy profile σ^* is a* **Bayesian equilibrium** *in a Bayesian game if for each player i, each type $t_i \in T_i$, and each action $a_i \in A_i(t_i)$,*

$$U_i((\sigma_i^*, \sigma_{-i}^*)|t_i) \geq U_i((a_i, \sigma_{-i}^*)|t_i) \quad \forall\ t_i \in T_i.$$

Informally, it says that a player i of type t_i maximizes his expected/interim payoff by following σ_i^* given that all other players follow σ_{-i}^*.

The restriction to pure actions on the RHS of the inequality in the above definition is without loss of generality since it automatically implies the inequality over mixed actions also. In games with finite type spaces, Bayesian equilibrium and Nash equilibrium are equivalent concepts - we show this below.

We will say a Bayesian game is finite if the set of types of each player is finite *and* the set of actions available to each player and each type is finite. Finite Bayesian games useful properties (some of which extend to infinite games).

The first property that we show is that in a finite game, a strategy profile is a Nash equilibrium if and only if it is a Bayesian equilibrium. In other words, a player has a profitable deviation in Bayesian game before he learns his type if and only if he has a profitable deviation after he learns his type. This result will use the fact that probability of every type occurring is positive.

THEOREM **19** *In a finite Bayesian game, a strategy profile is a Bayesian equilibrium if and only if it is a Nash equilibrium.*

Proof: Consider a strategy profile σ^*. Suppose σ^* is a Bayesian equilibrium. Then, for every $i \in N$, for every $t_i \in T_i$, and every $a_i \in A_i(t_i)$, we have

$$U_i(\sigma_i^*, \sigma_{-i}^*|t_i) \geq U_i(a_i, \sigma_{-i}^*|t_i).$$

For any pure strategy s_i, we know that

$$U_i(s_i, \sigma_{-i}^*) = \sum_{t_i \in T_i} p(t_i) U_i(s_i(t_i), \sigma_{-i}^*|t_i) \leq \sum_{t_i \in T_i} p(t_i) U_i(\sigma_i^*, \sigma_{-i}^*|t_i) = U_i(\sigma_i^*, \sigma_{-i}^*).$$

Since this holds for every pure strategy s_i, it holds for any mixed strategy σ_i by the indifference principle. Hence, σ^* is a Nash equilibrium.

Now, suppose that σ^* is a Nash equilibrium. Assume for contradiction that σ^* is not a Bayesian equilibrium. Then, there is some $i \in N$ and some $t_i \in T_i$ with $a_i \in A_i(t_i)$ such that

$$U_i(a_i, \sigma_{-i}^*|t_i) > U_i(\sigma_i^*, \sigma_{-i}^*|t_i).$$

Then, we consider a new strategy $\hat{\sigma}_i$ of Player i such that $\hat{\sigma}_i(t_i') = \sigma_i(t_i')$ if $t_i' \neq t_i$ and

$\hat{\sigma}_i(t_i) = a_i$. Now, observe the following:

$$\begin{aligned}
U_i(\hat{\sigma}_i, \sigma^*_{-i}) &= \sum_{t'_i \in T_i} p(t'_i) U_i(\hat{\sigma}_i, \sigma^*_{-i} | t'_i) \\
&= \sum_{t'_i \in T_i \setminus \{t_i\}} p(t'_i) U_i(\hat{\sigma}_i, \sigma^*_{-i} | t'_i) + p(t_i) U_i(\hat{\sigma}_i, \sigma^*_{-i} | t_i) \\
&= \sum_{t'_i \in T_i \setminus \{t_i\}} p(t'_i) U_i(\sigma^*_i, \sigma^*_{-i} | t'_i) + p(t_i) U_i(a_i, \sigma^*_{-i} | t_i) \\
&> \sum_{t'_i \in T_i \setminus \{t_i\}} p(t'_i) U_i(\sigma^*_i, \sigma^*_{-i} | t'_i) + p(t_i) U_i(\sigma^*_i, \sigma^*_{-i} | t_i) \\
&= U_i(\sigma^*).
\end{aligned}$$

This contradicts the fact that σ^* is a Nash equilibrium. ■

The next property we show is that finite Bayesian games have a Bayesian equilibrium.

Theorem **20** *Every finite Bayesian game has a Bayesian equilibrium.*

Proof: Consider a new strategic form game where the set of players is the set $\cup_{i \in N} T_i$, i.e., every type of each player is a player in this new game. The set of pure strategy of a player t_i is $A_i(t_i)$. Note that an action profile a will consist of an action for every type. Hence $a \equiv \{a(t_i)\}_{t_i \in \cup_{i \in N} T_i}$. The payoff of Player t_i in this game in a pure action profile a is

$$u_{t_i}(a) = \sum_{t_{-i} \in T_{-i}} p(t_{-i} | t_i) U_i((t_i, t_{-i}); \{a_j(t_j)\}_{j \in N}).$$

Now, a mixed strategy of player t_i in the new game is a probability distribution over $\Delta A_i(t_i)$, which is exactly a strategy in the Bayesian game. Similarly, a strategy $\sigma_i(t_i) \in \Delta A_i(t_i)$ in the original game corresponds to a mixed strategy of Player t_i in the new game. Further, note that a Nash equilibrium of the new game is a Bayesian equilibrium of the original game.

Since the new game is a finite strategic game, a Nash equilibrium in mixed strategies exist, and hence, a Bayesian equilibrium of the original game exists. ■

Though this existence result was for finite Bayesian games, many non-finite Bayesian games have a Bayesian equilibrium. In fact, games with infinite type spaces and infinite action spaces are common in economics - one of the reason is that one can make use of tools and techniques from analysis easily.

12 Analysis of First-price Auction

We will study a model of selling a single indivisible object. Each agent derives some utility by acquiring the object - we will refer to this as his **valuation**. In the terminology of the Bayesian games, the valuation is the type of the agent.

We will study auction formats to sell the object. This will involve payments. A central assumption in auction theory is that utility from monetary payments is **quasi-linear**, i.e., if an agent gets utility v from the object and pays an amount p, then his net utility is

$$v - p.$$

Implicitly, this assumes risk neutral bidders - the net utility of a bidder is his net payoff.

Another fundamental assumption that is commonly made is that of **no externality**, i.e., if an agent does not win the object then he gets zero utility. All the auctions that we will study will involve zero payments by the agent who does not win the object.

We will now study the equilibrium, revenue, and welfare properties of two auction formats for the sale of a single object.

12.1 The Second-Price (Vickrey) Auction

Suppose each buyer $j \in N$ bids an amount b_j. Then the highest buyer wins the object. We assume that in case of a tie for the highest bid, each bidder gets the good with equal probability. We denote the probability of winning at a profile of bids $b \equiv (b_1, \dots, b_n)$ as $\phi_j(b)$ for each buyer $j \in N$. Note that $\phi_j(b) = 1$ if $b_j > \max_{k \neq j} b_k$ and $\phi_j(b) = 0$ if $b_j < \max_{k \neq j} b_k$. Then the payoff of buyer $j \in N$ with value x_j is given by

$$\pi_j(b) = \phi_j(b)\big[x_j - \max_{k \neq j} b_k\big]$$

The following theorem has already been proved (see Lemma 2)

Theorem **21** *A weakly dominant strategy in the second-price auction (Vickrey auction) is to bid your true value.*

Note that the statement of this theorem did not refer to the Bayesian game (in particular to the priors) because the solution concept used here is (weak) dominant strategy, which is a prior-free solution concept.

12.2 The Symmetric Model

We will now analyze another auction format that does not have a dominant strategy equilibrium but has a nice Bayesian equilibrium. We will assume that all the bidders draw their value from some interval $[0, w]$ using a distribution F (same for all the bidders). We also assume that F admits a density function f such that $f(x) > 0$ for all $x \in [0, w]$. It is possible that the interval is the whole non-negative real line, in which case, we will abuse notation to let $w = \infty$. But the mean of this distribution will be finite.

12.3 Payment in the Vickrey Auction

Consider any arbitary bidder, say 1. Let the random variable of the highest value of the remaining $n-1$ bidders be Y_1 (it is the random variable of maximum of $n-1$ random variables). Let G be the cumulative distribution function of Y_1. Notice that for all y, $G(y) = F(y)^{n-1}$. Also, if any bidder has true value x_1, then his probability of winning in the Vickrey auction is $G(x_1)$. If he wins, his expected payment is $E(Y_1|Y_1 < x_1)$.

Hence, the expected payment of a bidder in the Vickrey auction when a bidder has true value x is

$$\begin{aligned}\pi^{II}(x) &= G(x)E(Y_1|Y_1 < x) \\ &= G(x)\frac{\int_0^x yg(y)dy}{G(x)} \\ &= \int_0^x yg(y)dy.\end{aligned}$$

12.4 The First-Price Auction

Like in the Vickrey auction, the highest buyer wins the object in the first-price auction too. We assume that in case of a tie for the highest bid, each bidder gets the good with equal probability. We denote the probability of winning at a profile of bids $b \equiv (b_1, \ldots, b_n)$ as $\phi_j(b)$ for each buyer $j \in N$. Note that $\phi_j(b) = 1$ if $b_j > \max_{k \neq j} b_k$ and $\phi_j(b) = 0$ if $b_j < \max_{k \neq j} b_k$.

Given a profile of bids $b \equiv (b_1, \ldots, b_n)$ of bidders, the payoff to bidder j with value x_j is given by

$$\pi_j(b) = \phi_j(b)\big[x_j - b_j\big]$$

12.5 Symmetric Equilibrium

Unlike the Vickrey auction, the first-price auction has no weakly dominant strategy (verify). Obviously, bidding your true value guarantees a payoff of zero, and there are obvious ways to generate positive expected payoff. Hence, we adopt the weaker solution concept of Bayesian equilibrium. In fact, we will restrict ourselves to equilibria where bidders use the same *bidding function* which are technically well behaved.

In particular, for any bidder $j \in N$, a strategy $\beta_j : [0, w] \to \mathbb{R}_+$ is his bidding function. The focus in our study will be **monotone symmetric equilibria**, where every bidder uses the same bidding function. So, we will denote the bidding function (strategy in the Bayesian game) by simply $\beta : [0, w] \to \mathbb{R}_+$. We assume $\beta(\cdot)$ to be strictly increasing and differentiable.

Bayesian equilibrium requires that if every bidder except bidder i follows $\beta(\cdot)$ strategy, then the expected payoff maximizing strategy for bidder i must be $\beta(x)$ when his value is x. Note that if bidder i with value x bids $\beta(x)$, and since everyone else is using $\beta(\cdot)$ strategy, increasingness of β ensures that the probability of winning for bidder i is equal to the probability that x is the highest value, which in turn is equal to $G(x)$. Thus, we can define the notion of a symmetric (Bayesian) equilibrium in this case as follows.

Definition **25** *A strategy profile $\beta : [0, w] \to \mathbb{R}_+$ for all agents is a* **symmetric Bayesian equilibrium** *if for every bidder i and every type $x \in [0, w]$*

$$G(x)(x - \beta(x)) \geq \textit{Probability of winning by bidding } b(x - b) \qquad \forall\, b \in \mathbb{R}_+,$$

where the probability of winning is calculated by assuming bidders other than bidder i are following $\beta(\cdot)$ strategy.

Remember that due to symmetry, $G(x)$ indicates the probability of winning in the auction when the bidder bids $\beta(x)$, and $(x - \beta(x))$ is the resulting payoff.

Theorem **22** *A symmetric equilibrium in a first-price auction is given by*

$$\beta^I(x) = \frac{1}{G(x)} \int_0^x y g(y) dy = E[Y_1 | Y_1 < x],$$

where Y_1 is the highest of $n - 1$ independently drawn values.

Proof: Suppose every bidder except bidder 1 follows the suggested strategy. Let bidder 1 bid b. Notice that $b \leq \beta(w)$ in equilibrium since bidding any amount strictly greater than $\beta(w)$ cannot be an equilibrium - the bidder can always increase payoff by reducing the bid slightly but still larger than $\beta(w)$. Hence, bid amount of a bidder will lie between 0 and $\beta(w)$, and hence, there exists a $z = \beta^{-1}(b)$. Then the expected payoff from bidding $\beta(z) = b$ when his true value is x is

$$\begin{aligned}
\pi(b, x) &= G(z)\big[x - \beta(z)\big] \\
&= G(z)x - \int_0^z yg(y)dy \\
&= G(z)x - zG(z) + \int_0^z G(y)dy \\
&= G(z)\big[x - z\big] + \int_0^z G(y)dy,
\end{aligned}$$

where, we have integrated by parts in the fourth equality [5]. Hence, we can write

$$\pi(\beta(x), x) - \pi(\beta(z), x) = G(z)(z - x) - \int_x^z G(y)dy \geq 0.$$

Notice that the previous inequality holds whether $z \leq x$ or $z \geq x$. Hence, bidding according to $\beta(\cdot)$ is a symmetric equilibrium. ∎

We now prove that this is the unique symmetric equilibrium in the first-price auction. Now, consider any bidder, say 1. Assume that he realizes a true value x, and wants to determine his optimal bid value b using a symmetric (increasing and differentiable) bidding function β.

Notice that when a bidder realizes a value zero, by bidding a positive amount, he makes a loss. So, $\beta(0) = 0$. Bidder 1 wins whenever his bid $b > \max_{i \neq 1} \beta(X_i)$, equivalently $b > \beta(\max_{i \neq 1} X_i) = \beta(Y_1)$ (since $\beta(\cdot)$ is increasing). This is again equivalent to saying $Y_1 < \beta^{-1}(b)$ (since $\beta(\cdot)$ is increasing, an inverse exists). Hence, his expected payoff is

$$G(\beta^{-1}(b))(x - b).$$

A necessary condition for maximum is the first order condition, which is obtained by differentiating with respect to b.

$$\frac{g(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))}(x - b) - G(\beta^{-1}(b)),$$

where we used $g = G'$ is the density function of Y_1 and $\beta(\beta^{-1}(b)) = b$. At the equilibrium, $b = \beta(x)$, this should equal to zero, which reduces the above equation to

$$\begin{aligned}
G(x)\beta'(x) + g(x)\beta(x) &= xg(x) \\
\Leftrightarrow \frac{d}{dx}(G(x)\beta(x)) &= xg(x).
\end{aligned}$$

[5] To remind, integration by parts $\int h_1(y)h_2'(y)dy = h_1(y)h_2(y) - \int h_1'(y)h_2(y)dy$.

Integrating both sides, and using $\beta(0) = 0$, we get

$$\beta(x) = \frac{1}{G(x)} \int_0^x y g(y) dy = E[Y_1 | Y_1 < x].$$

Hence, this is the unique symmetric equilibrium in the first-price auction.

The equilibrium bid in the first-price auction can be rewritten as

$$\beta^I(x) = x - \int_0^x \frac{G(y)}{G(x)} dy.$$

This amount is less than x. From the proof of the Theorem 22, it can be seen that if a bidder with value x bids $\beta(z')$ with $z' > z$, then his loss in payoff is the shaded area above the $G(\cdot)$ curve in Figure 4. On the other hand, if he bids $\beta(z'')$ with $z'' < z$, then his loss in payoff is the shaded area below the $G(\cdot)$ curve in Figure 4.

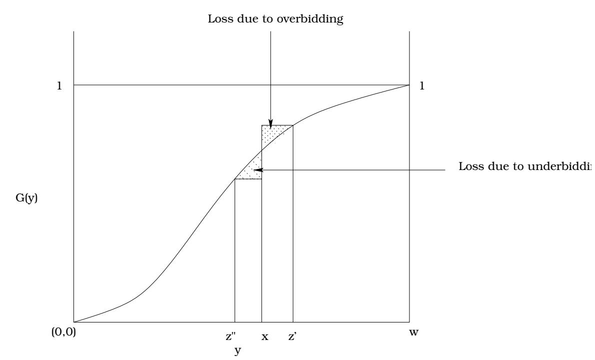

Figure 4: Loss in first-price auction by deviating from equilibrium

Another observation is that $\frac{G(y)}{G(x)} = (\frac{F(y)}{F(x)})^{n-1}$. Hence, the amount of lowering of bid vanishes to zero as the number of bidders increase, and the equilibrium bid amount approaches the true valuation.

Hence, the expected payment in the first price auction for a bidder with value x can be written as

$$\pi^I(x) = G(x)\beta(x) = G(x)E(Y_1 | Y_1 < x) = \int_0^x y g(y) dy = \pi^{II}(x).$$

It is instructive to look at some examples. Suppose values are distributed uniformly in $[0, 1]$. So, $F(x) = x$ and $G(x) = x^{n-1}$. So, $\beta(x) = x - \frac{1}{x^{n-1}} \int_0^x y^{n-1} dy = x - \frac{x}{n} = \frac{n-1}{n} x$. So, in equilibrium, every bidder bids a constant fraction of his value.

Let us consider the case of two bidders, and values distributed exponentially on $[0, \infty)$ with mean $\frac{1}{\lambda}$. So, $F(x) = 1 - \exp(-\lambda x)$ and for $n = 2$, $G(x) = F(x)$. So, $\beta(x) = E[Y_1 :$

$Y_1 < x] \leq E[Y_1] = E[X]$. If $\lambda = 2$, this means that $\beta(x) \leq 0.5$. This means that even if the bidder has a very high value of 100000000, he will not bid more than 0.5 in equilibrium. The intuition behind this is that even if the bidder has very high value, he has low probability of losing if he bids less than 0.5. So, it makes sense for him to bid low and get a larger expected profit.

13 Revenue Equivalence

The revenue from an auction is the sum of total *ex-ante payment* of all the bidders. Since the equilibrium interim payment of each bidder is the same in both the first price and the second price auction. It is immediate that the revenue from both the auctions are the same. This is sometimes termed as the *revenue equivalence* theorem.

Theorem **23** *Suppose bidders have private values with independent and identical distributions. Then any symmetric and increasing equilibrium of first-price and second-price auction yields the same expected revenue to the seller.*

This is a striking result because even though the actual payments in both the auctions can be quite different the expected payments turn out to be the same.

We can in fact compute an exact expression for the revenue in these auctions. The exact value of *ex ante* expected payment of the seller in the first-price auction can also be computed. This is equal to

$$\begin{aligned}
E(\pi^I(x)) = n\int_0^w \pi^I(x)f(x)dx &= n\int_0^w \Big(\int_0^x yg(y)dy\Big)f(x)dx \\
&= n\int_0^w \Big(\int_y^w f(x)dx\Big)yg(y)dy \\
&= n\int_0^w \big(1-F(y)\big)yg(y)dy \\
&= \int_0^w n(n-1)(1-F(y))F(y)^{n-2}yf(y)dy \\
&= E(\text{second highest value}).
\end{aligned}$$

The last equality can be explained as follows. Let us consider the random variable of the second highest number of n randomly drawn numbers using F, and denote its cumulative density function as $F^{(2)}$. Let us find the value $F^{(2)}(y)$. The probability that the second highest value is less than or equal to y can be broken into two disjoint events: (a) probability that all the values are less than y - which is $F(y)^n$, and (b) probability that exactly $n-1$ values are less than y - which $nF(y)^{n-1}(1-F(y))$. So, we can write

$$F^{(2)}(y) = F(y)^n + nF(y)^{n-1}(1-F(y)) = nF(y)^{n-1} - (n-1)F(y)^n.$$

This gives,

$$f^{(2)}(y) = n(n-1)F(y)^{n-2}f(y) - n(n-1)F(y)^{n-1}f(y) = n(n-1)F(y)^{n-2}f(y)(1-F(y)).$$

Since the expected second highest value is

$$\int_0^w yf^{(2)}(y)dy = \int_0^w n(n-1)(1-F(y))F(y)^{n-2}yf(y)dy,$$

which is exactly the expression we have. Hence, the total expected payment in the first-price auction is the expected second highest value of a bidder, which is also the total expected payment in the second-price auction.

The expected payment of a buyer with value x in the first-price auction or second-price auction can be written as

$$\begin{aligned}\pi^I(x) = \pi^{II}(x) = \int_0^x yg(y)dy &= xG(x) - \int_0^x G(y)dy \\ &= \text{expected value} - \text{expected profit}.\end{aligned}$$

Since $xG(x)$ is the expected value to a buyer with value x, the expected profit for him is $\int_0^x G(y)dy$. This is shown graphically in Figure 5.

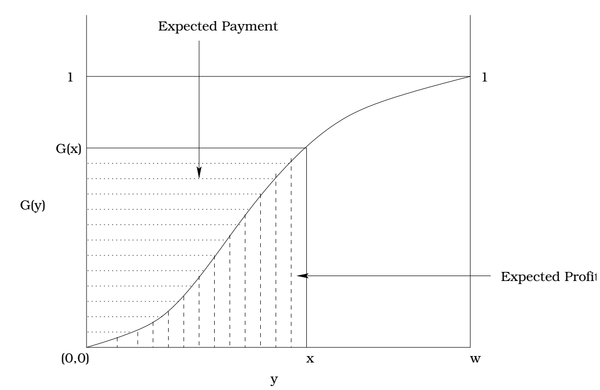

Figure 5: Expected profit and payment in the first-price or second-price auction

13.1 Uniform Distribution

For uniform distribution in interval $[0, w]$, $F(x) = \frac{x}{w}$ and $f(x) = \frac{1}{w}$. This gives $G(x) = \frac{x^{n-1}}{w^{n-1}}$. Hence, bid of a bidder with value x in first-price auction is

$$\beta^I(x) = x - \int_0^x \frac{y^{n-1}}{x^{n-1}} dy = x - \frac{x}{n} = \frac{n-1}{n} x.$$

The revenue in the first price and second-price auction is

$$n(n-1) \int_0^w (1 - \frac{x}{w})(\frac{x}{w})^{n-2} x \frac{1}{w} dx = \frac{n(n-1)}{w^n} \int_0^w (w - x) x^{n-1} dx = \frac{n-1}{n+1} w.$$

Note that with two bidders, the symmetric equilibrium of a bidder is to bid $\frac{1}{2}$ his value in the first-price auction. Hence, at a profile of valuations (v_1, v_2) with $v_1 > v_2$, agent 1 pays v_2 in the second-price auction and $\frac{v_1}{2}$ in the first-price auction. Clearly, there are some regions where one auction does better than the other in terms of revenue, but the expected revenue is the same in both the auctions.

13.2 Analysis of Bilateral Trading

The bilateral trading is one of the simplest model to study Bayesian games. It involves two players: a buyer (b) and a seller (s). The seller can produce a good with cost c and the buyer has a value v for the good. Suppose both the value and the cost are distributed *uniformly* in $[0, 1]$.

Now, consider the following Bayesian game. The buyer announces a price p_b that he is willing to pay and the seller announces a price p_s that she is willing to accept. Trade occurs if $p_b > p_s$ at a price equal to $\frac{p_b + p_s}{2}$. If $p_b \le p_s$, then no trade occurs.

The type of the buyer is his value $v \in [0, 1]$ and the type of the seller is his cost $c \in [0, 1]$. A strategy for each agent is to announce a price given their types. In other words, the strategy of the buyer is a map $p_b : [0, 1] \to \mathbb{R}$ and $p_s : [0, 1] \to \mathbb{R}$.

If no trade occurs, then both the agents get zero payoff. If trade occurs at price p, then the buyer gets a payoff of $v - p$ and the seller gets a payoff of $p - c$.

Theorem **24** *There is a Bayesian equilibrium (p_b^*, p_s^*) in the bilateral trading problem with uniformly distributed types in $[0, 1]$, where for every $v, c \in [0, 1]$,*

$$p_b^*(v) = \frac{2}{3} v + \frac{1}{12}, \quad p_s^*(c) = \frac{2}{3} c + \frac{1}{4}.$$

Proof: Suppose the seller follows strategy p_s^*. Then he never quotes a price above $\frac{2}{3}+\frac{1}{4}=\frac{11}{12}$. So, the buyer should never quote a price above $\frac{11}{12}$ in a best response. Similarly, the seller quotes a minimum price of $\frac{1}{4}$. Hence, the buyer should never quote a price below $\frac{1}{4}$ as best response. Suppose he quotes a price π_b when his value is v. Then, trade occurs if the $p_s^*(c) < \pi_b$ or $c < \frac{3}{2}\pi_b - \frac{3}{8}$. Note that since $\frac{1}{4} \le \pi_b \le \frac{11}{12}$, we have $0 \le \frac{3}{2}\pi_b - \frac{3}{8} \le 1$.

Let $x_b \equiv \frac{3}{2}\pi_b - \frac{3}{8}$. Then the expected payoff of buyer from bidding π_b at type v is

$$\begin{aligned}
\int_0^{x_b} \Big(v - \frac{\pi_b + p_s^*(c)}{2}\Big)dc &= \int_0^{x_b} \Big(v - \frac{\pi_b + \frac{2}{3}c + \frac{1}{4}}{2}\Big)dc \\
&= \big(v - \frac{\pi_b}{2} - \frac{1}{8}\big)x_b - \frac{1}{6}x_b^2 \\
&= \big(v - \frac{1}{3}x_b - \frac{1}{4}\big)x_b - \frac{1}{6}x_b^2 \\
&= (v - \frac{1}{4})x_b - \frac{1}{2}x_b^2.
\end{aligned}$$

This is a strictly concave function in π_b, hence, the first order condition gives the unique maximum. The first order condition gives $(v-\frac{1}{4}) - x_b = 0$. This implies that $x_b = \frac{3}{2}\pi_b - \frac{3}{8} = v - \frac{1}{4}$. Hence, $\pi_b = \frac{2}{3}v + \frac{1}{12}$ is a best response.

A similar optimization exercise solves the seller's problem. Suppose the buyer follows strategy p_b^*. Then, the buyer quotes a minimum of $\frac{1}{12}$ and a maximum of $\frac{3}{4}$. Then the seller should never quote less than $\frac{1}{12}$ because such a strategy will not maximize his expected payoff. Suppose he quotes π_c, then trade occurs if $\pi_c < \frac{2}{3}v + \frac{1}{12}$, which reduces to $v > \frac{3}{2}\pi_c - \frac{1}{8} \ge 0$ since $\pi_c \ge \frac{1}{12}$. Further, $\frac{3}{2}\pi_c - \frac{1}{8} \le 1$ since $\pi_c \le \frac{3}{4}$. Denote $x_c = \frac{3}{2}\pi_c - \frac{1}{8}$. Hence, the expected payoff of the seller at type c is given by

$$\begin{aligned}
\int_{x_c}^{1} \Big(\frac{\pi_c + \frac{2}{3}v + \frac{1}{12}}{2} - c\Big)dv &= \int_{x_c}^{1} \Big(\frac{1}{2}\pi_c + \frac{1}{24} - c + \frac{1}{3}v\Big)dv \\
&= \int_{x_c}^{1} \Big(\frac{1}{3}\pi_c + \frac{1}{12} - c + \frac{1}{3}v\Big)dv \\
&= \Big(\frac{1}{3}\pi_c + \frac{1}{12} - c\Big)(1 - x_c) + \frac{1}{6}(1 - x_c^2).
\end{aligned}$$

Again this is a strictly concave function and its maximum can be found by solving the first order condition. The first order condition gives us

$$\frac{1}{3}(1 - x_c) - \Big(\frac{1}{3}\pi_c + \frac{1}{12} - c\Big) - \frac{1}{3}x_c = 0.$$

This gives us $x_c = \frac{3}{2}\pi_c - \frac{1}{8} = c + \frac{1}{4}$, which gives the unique best response as $\pi_c = \frac{2}{3}c + \frac{1}{4}$. ■

There are other Bayesian equilibria of this game. However, this equilibrium can be shown to be unique in the class of strategies where players use strategies linear in their type. One

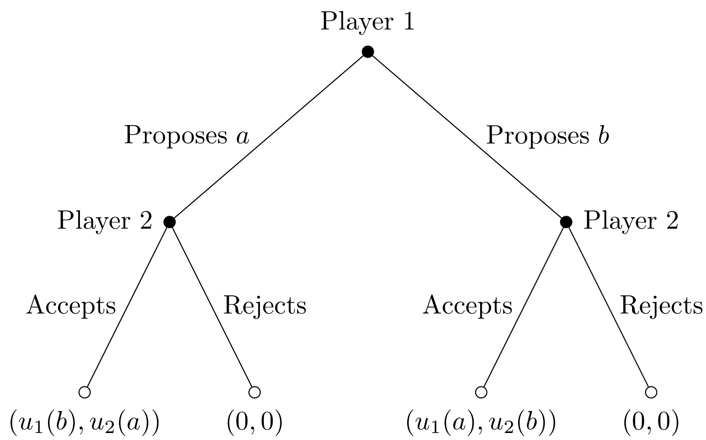

Figure 6: Extensive form game with perfect information

notable feature of this equilibrium is that trade occurs when $p_b^*(v) > p_s^*(c)$, which is equivalent to requiring $\frac{2}{3}v + \frac{1}{12} > \frac{2}{3}c + \frac{1}{4}$. This gives $v - c > \frac{1}{4}$. Note that efficiency will require trade to happen when $v > c$. Hence, there is some loss in efficiency. This is in general an impossibility - you cannot construct any Bayesian game whose equilibrium will have efficiency in Bayesian equilibrium in this model (more on this in some advanced course).

14 Extensive Form Games

In many situations strategic interactions between agents happen sequentially. Unlike in strategic form games, agents move sequentially in such games. We consider some examples first.

Suppose two players are deciding how to share two indivisible objects $\{a, b\}$. First, Player 1 proposes an allocation. Player 2 observes the proposal of Player 1 and then decides whether to accept or reject the proposal. If Player 2 rejects, then no player gets any object. If Player 2 accepts the proposal, then each receives the proposed allocation of Player 1. Each player $i \in \{1, 2\}$ only cares about his own object and has a utility function $u_i \equiv (u_i(a), u_i(b))$, indicating his utility for the objects.

This situation can be modeled as an extensive game of perfect information. This is usually depicted by a game tree.

An important feature of this game is that Player 2 has completely observed what Player 1 has proposed. His action is contingent on what he has observed so far in the game. Such games are called extensive form games with perfect information, i.e., where every player has perfectly observed what has happened so far in the game at every point. The outcomes of the game are realized after the game ends. Players assign payoffs to this terminal stages of the game - this will involve assigning payoffs to every possible sequence of moves in the game.

Figure 6 depicts the extensive form game using a tree. The payoffs of the agents are written in the *leaf* nodes.

A strategy in such a game is a complex object. It must state the action to be taken for every contingent path that can be taken in this game.

We now look at another example where perfect information is absent. Suppose two friends are trying to meet. Friend 1 observes the weather in his city, which is either rain or sunny. Then, he decides to either go to Friend 2's place or stay at home. If Friend 1 stays at home, Friend 2 does not do anything and the game ends. If Friend 1 comes to Friend 2's place, she either takes him for dinner or cooks at home. Crucial here is the fact that Friend 2 does not observe the weather in Friend 1's city, which Friend 1 has observed. However, Friend 2 observes whether he Friend 1 has come to her place or not. But Friend 2 does not know if Friend 1 has come from a sunny city or rainy city. In that sense, though the game has sequential nature, the information is not perfect in this game.

There is a way to represent this game as an extensive form game with imperfect information. This is done by introducing the dummy player (Nature) who creates the imperfect information. Nature makes the first move by taking either the action "Rainy" or "Sunny". The action of Nature is observed by Friend 1 but not by Friend 2. After observing the action of Nature, Friend 1 takes either of the actions "Stay home" or "Go to Friend 2". Friend 1 can now come to Friend 2 from a Sunny city or a Rainy city. This idea is captured by an *information set*, where a bunch of nodes in the game are combined together to capture Friend 2's uncertainty about where she is in the game. Irrespective of where she is in the game, she observes that Friend 1 has come to her place, and then she chooses one of the actions "go out" or "stay in".

Figure 7 shows the extensive form game with information set. The information set of Player 2 is shown in dashed rectangle - it consists of two nodes in the game tree. At this information set, Player 2 does not know if Player 1 has come from a sunny city or rainy city.

Each of the possible paths in the game are assigned a payoff for each player. Further, games of imperfect information also specify probabilities/priors of uncertain moves of Nature. These are used to compute expected payoffs on information sets.

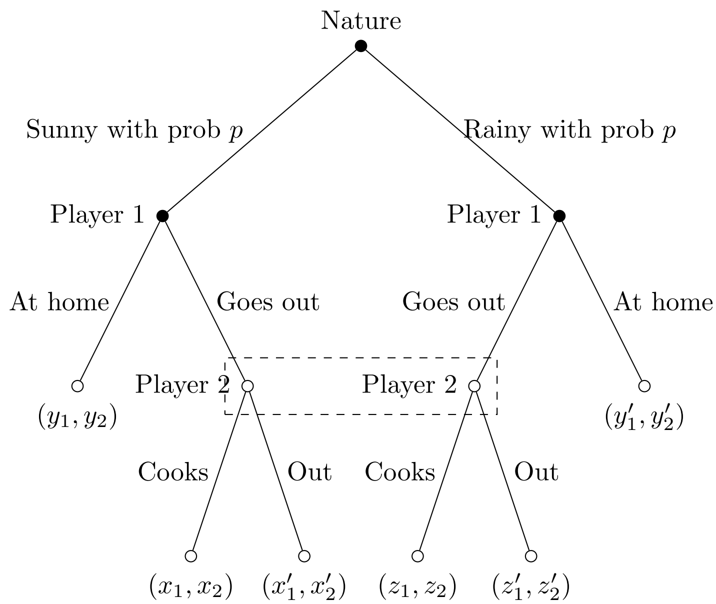

Figure 7: Extensive form game with information sets

15 Games with Perfect Information

We now formally define the notion of an extensive form game. We start from the most basic extensive game - a perfect information game, where every player at every node in the game knows what path/history has brought him to that node.

To formally define an extensive form game, we need to define a *cycle-free* graph. A graph $G = (V, E)$ is a set of a vertices V and subset of unordered pairs $E \subseteq V \times V$ such that for all $\{i, j\} \in E$, $i \neq j$. A cycle in a graph G is a sequence of distinct vertices $v_1, \ldots, v_k$ with $k > 2$ such that $\{v_1, v_2\}, \ldots, \{v_{k-1}, v_k\}, \{v_k, v_1\}$ are all edges of the graph. A graph G is cycle-free if there are no cycles in G.

A path in a graph G is a sequence of distinct vertices $v_1, \ldots, v_k$ such that $\{v_1, v_2\}, \ldots, \{v_{k-1}, v_k\}$ are all edges of the graph. A graph is connected if there is a path from every vertex to every other vertex. A connected and cycle-free graph is called a *tree*.

An important property of a tree graph is that there is a *unique* path from every vertex to every other vertex. From every tree $G = (V, E)$, we can construct a *rooted tree* by choosing a root vertex $r \in V$. A rooted tree is represented by $G \equiv (V, E, r)$. In a rooted tree, G, a vertex v is called the *child* of v' if there is an edge $\{v, v'\}$ and v' is in the unique path from root r to v. The set of all children of a vertex v is denoted by $C(v)$. Any vertex v with no children, i.e., $C(v) = \emptyset$ is called a *leaf* vertex.

An example of a rooted tree is shown in Figure 8. The root of this tree is shown. The leaves of the tree are $\{v_3, v_6, v_7, v_8, v_9, v_{10}\}$. For child: v_5 is the only child of v_2, where as v_1 has two children: $\{v_3, v_4\}$.

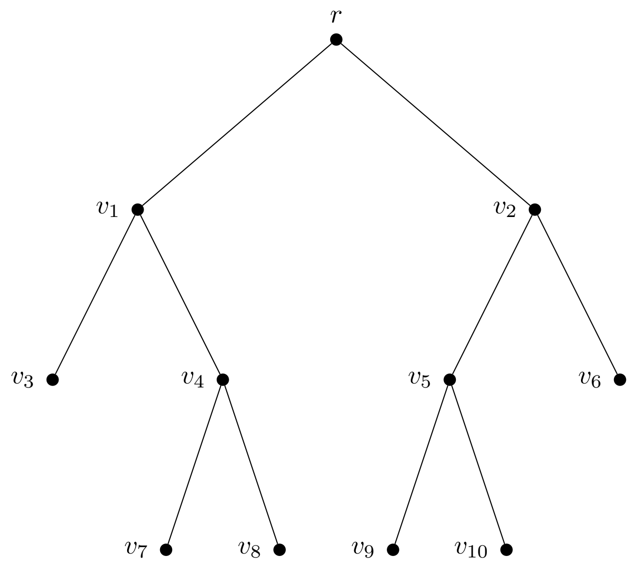

Figure 8: An example of a rooted tree

The backbone of an extensive form game is a rooted tree.

Definition **26** *An* **extensive form game of perfect information** *is*

$$\Gamma \equiv (N, V, E, r, \{V_i\}_{i\in N\cup\{0\}}, \{A(x)\}_{x\in V}, \{p_x\}_{x\in V_0}, \{u_i\}_{i\in N}),$$

where

- N *is the set of players*
- (V, E, r) *is a game tree, where*
  - *Each non-leaf vertex* $x \in V$ *specifies a player, called the decision maker at* x*, in* N *or Nature - Player* 0 *- who will take an action at this vertex.*
  - *Each edge* $\{x, y\} \in E$ *represents an action, in particular decision maker at* x *takes an action specified by this edge to reach vertex* y*. We will denote by* $A(x)$ *the set of actions available at vertex* x*.*
  - *Root vertex* r *specififes the first player in* $N \cup \{0\}$ *to take an action.*
- $A(x)$ *is the set of actions available at vertex* x*, and they map to the set of edges. Note that if* x *is a leaf vertex, then* $A(x)$ *is an empty set.*

- *$\{V_i\}_{i \in N \cup \{0\}}$ is a partitioning of the set of vertices that are not leaves. Hence, V_i represents the set of vertices where agent i takes action and V_0 represents the set of vertices where the Nature takes an action. It is possible that V_0 is empty, in which case, we say that this is an extensive form perfect information game without any chance moves.*

- *For every vertex $x \in V_0$, a probability distribution p_x over the set of actions $A(x)$ is known to all the players (perfect information).*

- *For every agent $i \in N$, $u_i(x)$ assigns a payoff for every leaf vertex x to Player i.*

We note here that the set of vertices/edges in a game tree may be infinite. This can happen because of two reasons: (1) the set of actions available at a vertex may be infinite and/or (2) the set of stages (i.e., lengths of paths) of the game may be infinite. At every vertex x in an extensive form game, the unique path from root r to vertex x conveys a lot of information: it contains information about who are the players who have taken what action to reach from r to x. It is standard to denote this information on the path as **history** h_x at vertex x. In fact, an alternate representation of an extensive form game is to just specify the history at every vertex.

Consider the following example of Figure 6. There is only one vertex, the root vertex, where Player 1 is the decision maker. For all other non-leaf nodes, Player 2 is the decision maker. Player 1 has two actions available to him - the two proposals he can make to Player 2. In each of his vertices, Player 2 has the same two actions (Accept, Reject) available to him. The payoffs of both the players are shown on the leaf vertices.

15.1 Strategy and Subgames

A strategy for a player in an extensive game must specify what he will do at each of his decision vertices. Hence, you can imagine a Player telling a computer to play on his behalf. In that case, he does not know ex-ante which decision vertices will be reached. So, he gives the computer a complete contingent plan of what actions must be taken at every decision vertex.

Formally, a strategy of player $i \in N$ is a map $s_i : V_i \to \cup_{x \in V_i} A(x)$ such that $s_i(x) \in A(x)$ for all $x \in V_i$.

Notice that there are certain games, where every player moves only once - these games are said to satisfy the *single move property.* However, there are games in which the single move property is not satisfied. In those games, if a strategy specifies a certain action at a decision vertex, that may ensure that certain decision vertex is never reached. But that does not exclude us from describing what action to take in those unreached vertices.

To see this, consider the game in Figure 9, where Player 2 moves twice. If Player 2 plays a strategy where he says he “Calls Player 1” at the first vertex, then exactly one more of his decision vertex will be reached. But a strategy for Player 2 must specify his action at *all* the decision vertices. This is crucial to evaluating his and his opponent’s options.

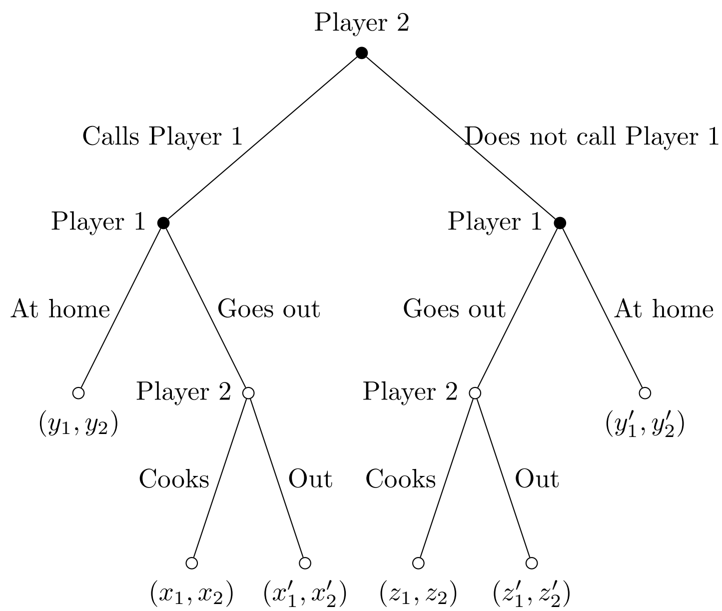

Figure 9: Extensive form game without single move property

The idea of a subgame is crucial to the analysis of an extensive form game. The **subgame** of an extensive form game of perfect information

$$\Gamma \equiv (N, V, E, r, \{V_i\}_{i \in N \cup \{0\}}, \{A(x)\}_{x \in V}, \{p_x\}_{x \in V_0}, \{u_i\}_{i \in N}),$$

starting at $x \in V$, where x is not a leaf vertex, is an extensive form game

$$\Gamma(x) \equiv (N, V(x), E(x), x, \{V_i(x)\}_{i \in N \cup \{0\}}, \{A(x')\}_{x' \in V(x)}, \{p_{x'}\}_{x' \in V_0(x)}, \{u_i\}_{i \in N}),$$

where the the (x) in the above notation means that the restriction of the original game starting from vertex x and its children, and children of its children etc.

Note that a game is a subgame of itself. So, every game has a subgame.

16 Equilibrium for Extensive Form Games

We now develop the theory of equilibrium for extensive form games. One naive way of doing that is to represent it as a strategic form game, and then apply the solution concepts of strategic form games. Representing an extensive form game as a strategic form game is quite easy: for every player i and every (pure) strategy of i in the extensive form game corresponds to a pure strategy in the strategic form game. The payoff from a strategy profile can then be computed from the game tree. This is because each strategy profile in the extensive form game maps to a unique terminal vertex of the game tree. This is called the **reduced normal/strategic form** of the extensive game. For a strategy profile s in an extensive form game Γ, we let x_s as the terminal vertex reached because of the strategy profile s. Then, the payoff of agent i from a strategy profile s is $u_i(x_s)$ and the payoff from a mixed strategy profile σ is given by

$$U_i(\sigma) := \sum_s u_i(x_s)\sigma_1(s_1) \times \ldots \times \sigma_n(s_n).$$

Definition **27** *A strategy profile σ is a Nash equilibrium of Γ if for all $i \in N$ and for all σ'_i*

$$U_i(\sigma) \geq U_i(\sigma'_i, \sigma_{-i}).$$

Nash equilibrium is not the correct solution concept for extensive form games because it misses the sequential move aspect of the game by treating it in strategic form. We illustrate this with an example.

Consider the game in Figure 10. The reduced strategic form representation of this game

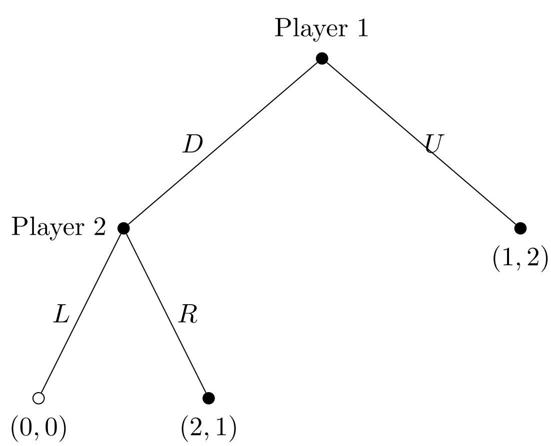

Figure 10: Nash equilibrium

is shown in Table 24. From this, one concludes that the game has two pure strategy Nash equilibria: (U, L) and (D, R).

	L	R
U	$(1,2)$	$(1,2)$
D	$(0,0)$	$(2,1)$

Table 23: Reduced strategic form of the game in Figure 10

But note that once the game has reached the information set of Player 2, he will play R. So, playing L is not *credible* for Player 2. Then, Player 1 can take this information into account while choosing his action. Player 1 clearly prefers playing D over U since Player 2 cannot threaten him credibly to play L. Hence, the equilibrium (U, L) is not a good prediction of the game.

The main idea here is that the equilibrium (U, L) specifies a strategy L for Player 2 which is not a credible strategy - once the decision vertex of Player 2 is reached, he will never play this.

As we discussed above, a strategy profile leads to a unique terminal vertex with a unique path from root to the terminal vertex. Hence, an equilibrium strategy profile will not touch on many decision vertices - these are called *off-equilibrium path* decision vertices. One primary requirement in extensive form game equilibrium is that action of every player must be optimal starting at every information set, not just information set reached on equilibrium path.

16.1 Subgame Perfect Equilibrium

We now discuss a *refinement* to Nash equilibrium for extensive form game. The idea of enforcing credibility is employed by using the notion of subgames.

DEFINITION **28** *A strategy profile σ is a* **subgame perfect equilibrium (SPE)** *of the extensive form game Γ if for every subgame of Γ the strategy profile σ restricted to that subgame is a Nash equilibrium of the subgame.*

Since Γ itself is a subgame of the game Γ, it follows that every SPE is a Nash equilibrium - hence, SPE is a refinement of Nash equilibrium. The game in Figure 10 has a unique SPE. To see this, the subgame starting from decision vertex of Player 2 has only one player. In that, Player 2 playing R is a dominant strategy. So, out of the two Nash equilibria of the entire game (subgame), only the one with R being played by Player 2 survives. Hence, (D, R) is the unique SPE.

16.2 Games with Perfect Information

In games with perfect information (and without any moves by Nature), the idea of a subgame perfect equilibrium is very compelling. We restrict attention to games with finite number of stages. We will show that it coincides with two other easy notions of equilibrium. In this section, we only focus attention on games with perfect information without any moves by Nature.

Figuring out Nash equilibrium of subgames can be quite a complicated task. In games with perfect information, this can be avoided because of a well known equivalence of subgame perfect equilibrium with two other notions. The first is the idea of *sequential rationality*.

Definition **29** *A strategy σ_i of Player i is* **sequentially rational** *given σ_{-i} if each decision vertex x of Player i, σ_i restricted to subgame at x is a best response to σ_{-i} restricted to the subgame at x.*

A strategy profile σ is **sequentially rational** *if for each Player i, σ_i is sequentially rational given σ_{-i}.*

The main difference between subgame perfect equilibrium and sequential rationality is that sequential rationality requires that at each subgame starting at decision vertex x, only the owner of decision vertex x must be choosing a best response. On the other hand, the subgame perfect equilibrium requires at every subgame, strategy of every player must be a best response given strategies of other players. Clearly, subgame perfection is more demanding, but we will show that both the ideas are the same.

Finally, an easy method to compute optimal behavior of agents in finite extensive form game is the following. Start with a decision vertex just before a terminal vertex. Specify an action that leads to the highest payoff for the owner of that vertex among all possible actions - in case of ties, all possible actions leading to highest payoff are specified. If such an optimal action leads to terminal vertex z, then replace this decision vertex and the subsequent subgame by terminal vertex z. Repeat this procedure. If indifferences occur, this will lead to multiple strategy profiles surviving. This procedure is called the **backward induction** procedure.

Definition **30** *A strategy profile that survives the above procedure is said to be a strategy profile* **surviving the backward induction procedure**.

We will prove the following theorem.

Theorem **25** *Let Γ be an extensive form game of perfect information with no Nature move and finite number of stages. Then the following are equivalent.*

1. σ is a subgame perfect equilibrium.

2. σ is sequentially rational.

3. σ survives the backward induction procedure.

Note here that (1) implies (2) because (1) requires best response from *all* players in the subgame but (2) only requires best response by the owner of the root vertex of the subgame. Also, (2) implies (3) because (2) allows owner of the root vertex to change his strategy at every decision vertex in the subgame but (3) only changes at one decision vertex at a time.

We will often refer to all these notions to be the definition of a subgame perfect equilibrium in such games. An immediate corollary of Theorem 25 is that a subgame perfect equilibrium in pure strategies always exist - this follows from the fact that the backward induction procedure always generates at least one pure strategy profile. If there are no indifferences in payoffs, the backward induction procedure generates a unique strategy profile, which is referred to as the backward induction *solution*.

16.2.1 Illustration of Backward Induction Procedure

In the game in Figure 10, Player 2 plays R. Then we replace the subgame starting at the decision vertex of Player 2 by payoff $(2, 1)$. Now, Player 1 chooses D in this new game. Hence, the unique outcome of the backward induction procedure is (D, R).

Consider the game in Figure 11. There are three players: two entrant firms and one incumbent firm. The entrants decide sequentially whether to stay out (O or o) or enter the market (E or e). If they stay out they get zero. If they enter, then the incumbent can fight ($f/f'/f''$) or accommodate ($a/a'/a''$). If both entrants stay out, the incumbent gets 5. If the entrant accommodates, the per firm profit is 2 for duopoly and -1 for triopoly. On top of this, if the incumbent fights, then it costs 1 for the incumbent and 3 for entrants. The game is described in Figure 11.

If we solve this game by backward induction procedure, then the incumbent always accommodates. Given this, entrant firm 2 enters in his left-most information set but stays out in the right-most information set. Given this, entrant firm 1 enters. This illustrates the idea of a first-mover advantage in extensive form games.

How do we describe the subgame perfect equilibrium of this game? We need to specify the actions at every information set: $(E, (e, o'), (a, a', a''))$. You can verify that there are many Nash equilibria of this game. Hence, Nash equilibrium has very less predictive power in this game but the subgame perfect equilibrium leads to a unique outcome.

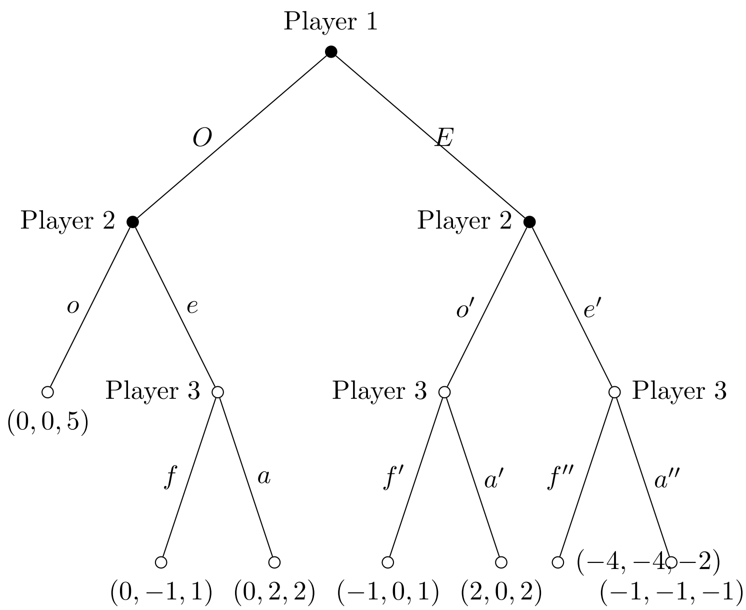

Figure 11: Backward induction

Backward induction can be a very demanding solution in games where players need to move many times. This is because it requires players to anticipate actions down the game tree. A sharp example of this fact is given a well known game called the **centipede game**. Two players start with 1 unit of money each. Each player can either decide to continue C or stop S. If anyone stops, then the game ends and each take their piles. If a player continues, then the opponent gets to take action but his pile is reduced by 1 while the opponent's pile is increased by 2. The play ends when any player reaches 100. Suppose Player 1 moves first. Unique prediction due to backward induction is Player 1 stops in the first chance resulting in $(1, 1)$. The subgame perfect equilibrium specifies action S at every decision vertex. This is also the unique Nash equilibrium of this game.

In lab experiments, agents have usually continued for some time. This is a general critique of equilibrium in extensive form game that no satisfactory refinement can predict such an outcome.

16.3 Indifferences and Backward Induction

If there are indifferences, then many pure and mixed strategies will survive backward induction and all of them will be subgame perfect equilibrium. To illustrate this, consider the following example in Figure 12.

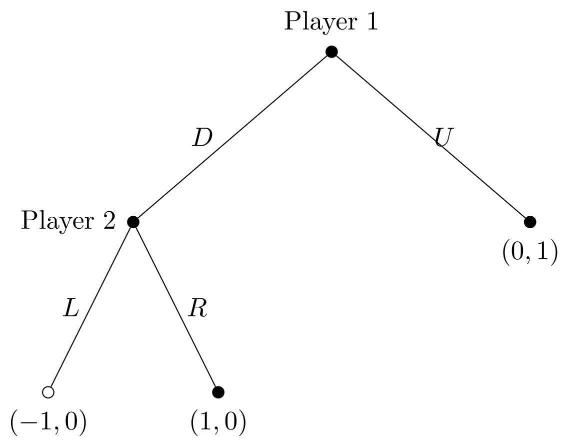

Figure 12: Backward induction with indifference

In the game in Figure 12, Player 2 is indifferent between his strategies L and R. Suppose he plays L, then optimal strategy for Player 1 is to play U. On the other hand if Player 2 plays R, then Player 1 chooses D. So, (U, R) and (D, L) are two subgame perfect equilibria. If Player 2 randomizes $\alpha L+(1-\alpha)R$. Player 1 gets 0 by playing U and $1-2\alpha$ by playing D. If $\alpha > \frac{1}{2}$, then Player 1 playing U is optimal. If $\alpha < \frac{1}{2}$, then Player 1 playing D is optimal. If $\alpha = \frac{1}{2}$, then Player 1 randomizing $\beta L + (1-\beta)D$ for any $\beta \in [0, 1]$ is optimal. All these correspond to subgame perfect equilibria of this game.

16.4 Proof of Theorem 25

For simplicity, we focus attention to pure strategy profiles. First thing, we note that is a redefinition of a strategy profile surviving backward induction procedure. Notice that backward induction requires that every decision vertex, the decision maker of that vertex chooses an optimal strategy given that the rest of the subgame has been chosen optimally. In particular, consider a subgame at decision vertex x with decision maker i. If s is a strategy profile that survives backward induction procedure, then denote its restriction to subgame at x as s^x. Backward induction requires that s_i^x is a better than any other strategy $\bar{s}_i^x$ for Player i given s_{-i}^x where $\bar{s}_i^x$ and s_i^x differ from each other by one decision vertex. Formally, s survives backward induction procedure if and only if this fact is true for every decision vertex x.

With the help of this redefinition of backward induction procedure, we can do the proof. The implications (1) $\Rightarrow$ (2) and (2) $\Rightarrow$ (3) are immediate from definitions. To see this note that sequential rationality requires that at every subgame starting with vertex x only the decision maker at x must best respond - but subgame perfect equilibrium needs that everyone must best respond. Hence, (1) $\Rightarrow$ (2). For (2) $\Rightarrow$ (3), consider the optimization problem done in the backward induction procedure and sequential rational strategies. Suppose Player i owns a decision vertex x. Denote the strategy profile s restricted to subgame from x as s^x. To verify sequential rationality of s_i^x given s_{-i}^x, we need to check for deviations at all decision vertices in the subgame. For backward induction, as we had argued earlier, we only need to check deviations of one decision vertex at a time. The meat of the proof lies in establishing the other directions.

(2) $\Rightarrow$ (1). Suppose s is a sequentially rational strategy profile. Let x be a decision vertex of agent i. We need to show that s^x is a Nash equilibrium of the subgame $\Gamma(x)$. Count the length of the paths from x to every possible decision vertex reachable from x, and denote the length of the maximal path by $\ell(x)$. We do the proof by induction on $\ell(x)$. If $\ell(x) = 1$, then the proof follows from sequential rationality itself. Assume $\ell(x) > 1$ and suppose that the claim is true for all y with $\ell(y) < \ell(x)$.

First, note that by definition, s_i^x is a best response to s_{-i}^x. Consider any player $j \neq i$. If j does not have a decision vertex in the subgame $\Gamma(x)$, then his strategy is vacuously a best response in this subgame. If j has a decision vertex in this subgame, let y be the first such decision vertex when we go from x to a terminal vertex. By induction and sequential rationality, s_j^y is a best response to s_{-j}^y.

Now, j's strategy in the subgame at x is the union of his strategies in each such y. Since each of them is a best response by induction, his strategy in the subgame $\Gamma(x)$ is also a best response. This shows that s is a Nash equilibrium.

(3) $\Rightarrow$ (2). For this direction, we will prove a general principle that is generally true in many variants of extensive form game. This is the one-shot deviation principle. Before doing so, note that we need to fix an agent i and strategies s_{-i} of other agents, and discuss about deviations of this agent. So, effectively, we are discussing a one-agent decision problem. We call a strategy s_i' a **one-shot deviation** of strategy s_i if it differs from s_i at exactly one vertex, say x. Further, the one-shot deviation strategy is **profitable** if it generates higher utility. Note that this is equivalent to requiring that it generates higher utility in the subgame $\Gamma(x)$.

Fixing the strategies of other players at s_{-i}, we will show that if for strategy s_i there is no strategy s_i' which is a one-shot profitable deviation, then s_i is sequentially rational given s_{-i}. Assume for contradiction that this is not true. Then, there is a decision vertex x of agent i such that s_i^x has a profitable deviation in subgame $\Gamma(x)$. Consider all paths from x to a decision vertex in the subgame $\Gamma(x)$ and let $L(x)$ denote the maximum number decision vertices along any such path that belongs to i. If $L(x) = 1$, then this will imply that s_i^x has a profitable one-shot deviation, contradicting survival from backward induction. So, $L(x) > 1$. We can now use induction on $L(x)$. Suppose the claim is true for all decision vertices y of agent i with $L(y) < L(x)$. Since a more than one-shot deviation in vertex x means a deviation in decision vertex y with $L(y) < L(x)$. But, this will mean a profitable deviation exists for decision vertex y also. This is a contradiction due to $L(y) < L(x)$ and our induction hypothesis.

16.5 Infinite Horizon and Action Sets

There are extensive games where the number of stages is infinite. For such games, the process of backward induction is not defined. However, the notion of subgame perfect equilibrium is still well defined. We need to consider subgames, and the strategies should consist of equilibrium behavior in each subgame.

Another important remark is that with finite number of stages, backward induction is well defined even if agents have infinite set of actions in a decision vertex. However, the optimal response may be empty with infinite set of actions. So, wherever the optimal response map is non-empty, we can easily define the backward induction process. The following example illustrates this point clearly.

16.6 Alternative Offers Bargaining

We now visit a classical application of subgame perfect equilibrium. In this problem, two players are bargaining over 1 unit of money. They will bargain for $T+1$ periods starting from period 0. In even periods (starting at 0), Player 1 offers a split $(o_t, 1 - o_t)$, where $o_t \in [0, 1]$ is Player 1's share. If Player 2 accepts, the game ends. Else, we move to the next period. In odd periods, Player 2 offers a split. If no split is accepted at the end of period T, then the game ends with each player getting 0. Money received in period t is discounted by δ^t, where $\delta \in (0, 1)$.

This game has perfect information, finite number of stages, but infinite set of actions at each decision vertex. There are many tied utilities too. But surprisingly, it has a unique subgame perfect equilibrium.

To understand the game better, consider just a one-period $T = 1$ case. Player 1 offers a split $(o_1, 1 - o_1)$ and Player 2 can either accept or reject. In all the decision vertices, where Player 2 gets a positive offer, he accepts. In the decision vertex where Player 2 gets zero offer, he is indifferent. Knowing this, we now apply backward induction on Player 1. Player

1's optimal is not clearly to give a positive split to Player 2 because that is dominated. If Player 2 rejects a zero offer with positive probability y, then Player 1 gets a payoff of $1-y$, which is dominated by Player 1 offering $(1-\frac{y}{2}, \frac{y}{2})$. Hence, again Player 2 rejecting a zero offer with positive probability and accepting a positive offer implies Player 1 has *no* optimal action at his decision vertex. Hence, the backward induction procedure does not provide any strategy of Player 1 for such a strategy of Player 2. On the other hand, if Player 2 accepts Player 1's zero offer with probability 1, then Player 1's optimal action is to offer $(1, 0)$. This will be a subgame perfect equilibrium. This forms the basis of the theorem below.

THEOREM **26** *In the alternative offers bargaining game, there is a unique subgame perfect equilibrium, where the initial offer is accepted. As $T \to \infty$, the equilibrium payoffs converge to $(\frac{1}{1+\delta}, \frac{\delta}{1+\delta})$.*

Proof: Suppose T is even. Then, in the last period, Player 1 offers. Consider the subgame from this period. It consists of a decision vertex for Player 1 where he offers a split $(o_T, 1-o_T)$ and a decision vertex for Player 2 for each offer of Player 1. In the decision vertex, Player 2 must accept any positive offer. But it can accept, reject, or randomize on zero offer. Then, consider the offer of Player 1. Player 1 cannot offer positive amount to Player 2 since he can improve it by giving half of that - hence, there is a one-shot deviation. So, Player 1 must offer 0 amount to Player 2. Now, if Player 2 rejects such an offer, then both get zero. Hence, if Player 2 randomizes with α probability reject and $(1-\alpha)$ probability accept, then Player 1 offering 0 gets a payoff of $(1-\alpha)\delta^T$. But Player 1 can do better by offering Player 2 an amount $\frac{1}{2}\alpha$ (which Player 2 will accept). Hence, if Player 2 rejects with positive probability, then offering 0 is not a best response of Player 1. So, offering 0 and getting rejected with some probability is not a subgame perfect equilibrium. Thus, offering 0 and accepting 0 is the unique subgame perfect equilibrium outcome from period T.

We now repeat this idea. Essentially, at each subgame an offer must be made such that the opponent is indifferent between accepting and rejecting and the opponent must accept. By backward induction, we proceed as follows.

1. In period T, Player 1 offers $(1, 0)$, which Player 2 accepts. Resulting payoffs are $(\delta^T, 0)$.

2. In period $(T-1)$, Player 1 can assure himself of δ^T. So, he accepts any offer giving him at least δ^T. So, Player 2 offers $(\delta, 1-\delta)$ which gives payoff $(\delta^T, \delta^{T-1} - \delta^T)$.

3. In period $(T-2)$, Player 2 can assure himself of $\delta^{T-1} - \delta^T$. So, Player 1 offers $(1 - \delta + \delta^2, \delta - \delta^2)$, which gives payoff $(\delta^{T-2} - \delta^{T-1} + \delta^T, \delta^{T-1} - \delta^T)$.

   Continuing in this manner, we get

4. In period 0, Player 1 offers $(1 - \delta + \delta^2 - \ldots + \delta^T, \delta - \delta^2 + \ldots - \delta^T) \equiv (\frac{1+\delta^{T+1}}{(1+\delta)}, \frac{\delta-\delta^{T+1}}{(1+\delta)})$, which is accepted by Player 2. Note that the limit of $T \to \infty$ is $(\frac{1}{1+\delta}, \frac{\delta}{(1+\delta)})$.

If T is odd, a similar analysis yields an offer by Player 1 equal to $(\frac{1-\delta^{T+1}}{(1+\delta)}, \frac{\delta+\delta^{T+1}}{(1+\delta)})$, whose limit $T \to \infty$ is also $(\frac{1}{1+\delta}, \frac{\delta}{(1+\delta)})$. ■

17 Games with Imperfect Information

In games with imperfect information a player may not observe the entire history at every decision vertex. Hence, when he reaches his decision vertex, there is uncertainty about which decision vertex he is really in. To make complete sense of this uncertainty, the set of actions available at each of these uncertain decision vertices must be same. This idea is captured by the notion of an information set.

Definition **31** *In an extensive form game the information set of Player i is a non-empty subset $U_i \subseteq V_i$ and a subset of actions $A(U_i)$, such that at each $x \in U_i$ we have $A(x) = A(U_i)$.*

The only additional information in an extensive form game with imperfect information is a specification of information sets. In particular, for every player i, we specify a partition $\{U_i^j\}_j$ of the decision vertices V_i of Player i, where each U_i^j is an information set. Now, set of actions are specified for each information set.

Definition **32** *An* **extensive form game of imperfect information** *is*

$$\Gamma \equiv (N, V, E, r, \{V_i\}_{i \in N \cup \{0\}}, \{U_i^j\}_{i \in N}^j, \{A(U_i^j)\}_{i \in N}^j, \{p_x\}_{x \in V_0}, \{u_i\}_{i \in N}),$$

where $\{U_i^j\}_{i \in N}^j$ is a partition of V_i and $A(U_i^j)$ specifies the actions available at each information set U_i^j for Player i.

Note that if every information set contains a single vertex, then the game is of perfect information.

The strategy and the idea of subgame is suitably changed in a game of imperfect information. Since the player is unsure about the vertex he has reached in an information set, his strategy must specify an action at every information set. We will denote by $\mathcal{U}_i \equiv \{U_i^1, \ldots, U_i^k\}$ the collection of information sets of Player i.

Formally, a strategy of player $i \in N$ is a map $s_i : \mathcal{U}_i \to \cup_{U_i^j \in \mathcal{U}_i} A(U_i^j)$ such that $s_i(U_i^j) \in A(U_i^j)$ for all $U_i^j \in \mathcal{U}_i$.

In the game in Figure 7, each player's information set is a singleton, except for Player 2, who has a single information set with two vertices. His strategy must specify what he will do at this information set.

The definition of a subgame is just the subtree starting from a decision vertex. If the game is of imperfect information, we need to worry about information sets. In particular, when we consider a subtree, for every Player and every information set of this player, all the vertices of this information set either belongs to the subtree or does not intersect with the subtree. So, $\Gamma(x)$ will be a subgame if for every $i \in N$ and for every $U_i^j \in \mathcal{U}_i$ either U_i^j lies in the subtree in $\Gamma(x)$ or it has an empty intersection with the subtree in $\Gamma(x)$.

The game in Figure 7 has only one subgame, i.e., the game itself. This is because every other subgame will only have part of the information set of Player 2.

17.1 Perfect Recall

Consider the following game in Figure 13. Player 2 is forgetful here. He forgets whether he had called Player 1 or not earlier. As a result, when Player 1 reaches his home, he does not know whether Player 2 has come because of his call or without his call. Thus, Player 2 has an information set consisting of two decision vertices.

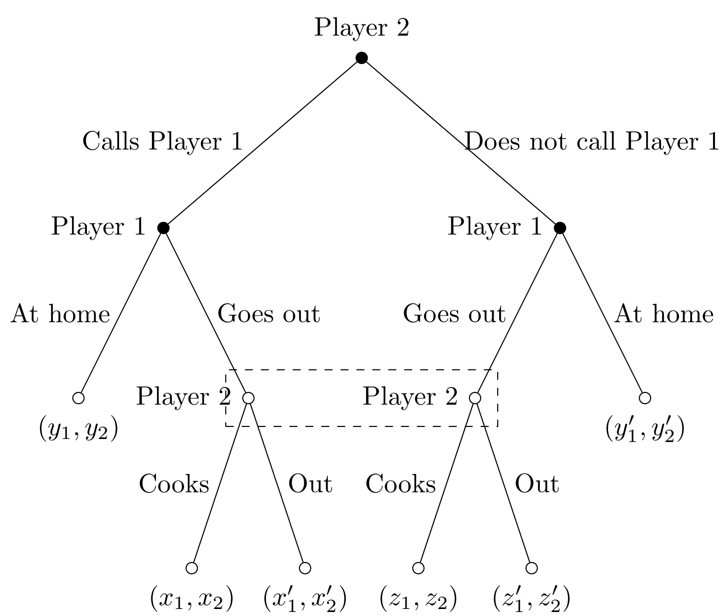

Figure 13: Extensive form game without perfect recall

Games in which players remember the entire sequence of information (history) from root to their every information set are players with **perfect recall**. Formally, Player i has perfect recall if at every information set U_i^j and every pair of vertices $x, x' \in U_i^j$, the information observed by Player i to reach x and x' from root are identical. An extensive form game in which all the players have perfect recall is called a game with perfect recall. We will exclusively focus attention on games in which all the players have perfect recall.

18 Mixed and Behavior Strategies

We have defined pure strategies in an extensive form game as a map that defines what action a player will take in each of his information sets. There are two natural ways to define *randomized* strategies in this environment. The first one says that we define a probability measure (distribution) over the set of all pure strategies. This is the notion of a mixed strategy. Formally, a **mixed strategy** of Player i is $\sigma_i \in \Delta \prod_{U_i^j \in \mathcal{U}_i} A(U_i^j)$.

Consider the game in Figure 14. Player 1 has two pure strategies - we roughly write it as $\{x, y\}$ to denote that in his only information set, he can either choose action x or action y. Similarly, the pure strategies of Player 2 can be written as $\{Aa, Ar, Ra, Rr\}$, where Aa indicates that in his left-most information set (decision vertex) he plays A and in the other information set he plays a - similar interpretations can be made for other pure strategies. A mixed strategy of Player 1 will be $\sigma_1(x), \sigma_1(y)$ such that $\sigma_1(x) + \sigma_1(y) = 1$. A mixed strategy of Player 2 will be $\sigma_2(Aa), \sigma_2(Ar), \sigma_2(Ra), \sigma_2(Rr)$ such that

$$\sigma_2(Aa) + \sigma_2(Ar) + \sigma_2(Ra) + \sigma_2(Rr) = 1.$$

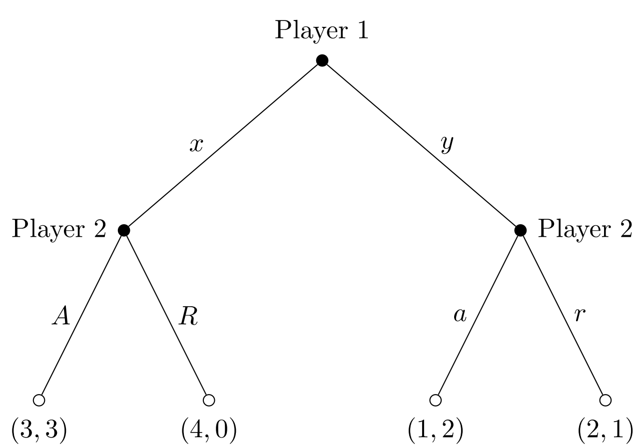

Figure 14: Extensive form game with perfect information

Another way to specify random behavior in this game is to specify a probability distribution at each information set. A **behavior strategy** of Player i specifies a probability distribution b_i^j over $A_i(U_i^j)$ for each of his information set U_i^j. Hence, $b_i \in \prod_{U_i^j \in \mathcal{U}_i} \Delta A_i(U_i^j)$. Notice that every behavior strategy naturally induces a probability distribution over pure strategies, and hence, is a mixed strategy.

In the game in Figure 14, Player 2 will have to specify two maps: $b_2^1(A), b_2^1(R)$ with $b_2^1(A)+b_2^1(R) = 1$ and $b_2^2(a), b_2^2(r)$ with $b_2^2(a) + b_2^2(r) = 1$. Note that the induced mixed strategy of Player 2 can be computed by multiplying the respective probabilities: for instance, $\sigma_2(Aa) = b_2^1(A)b_2^2(a)$. Thus, specifying randomization using a behavior strategy assumes independence across information sets - when a player reaches his information set he randomizes over the actions at that information set only.

Since mixed strategies allow for correlation, not every mixed strategy can be induced from behavior strategies. To see this, consider the game in Figure 14. Suppose $b_2^1(A) = \frac{1}{2} = b_2^1(R)$ and $b_2^2(a) = \frac{1}{3}$, $b_2^2(r) = \frac{2}{3}$. The mixed strategy generated is

$$\sigma_2(Aa) = \frac{1}{6}, \sigma_2(Ar) = \frac{1}{3}, \sigma_2(Ra) = \frac{1}{6}, \sigma_2(Rr) = \frac{1}{3}.$$

Now, consider the following mixed strategy of Player 2,

$$\sigma_2(Aa) = \frac{1}{3}, \sigma_2(Ar) = \frac{1}{6}, \sigma_2(Ra) = 0, \sigma_2(Rr) = \frac{1}{2}.$$

If there is a behavior strategy of Player 2 that generates this mixed strategy, then we must have $b_2^1(R) = 0$ or $b_2^2(a) = 0$, which will then imply that either $\sigma_2(Rr)$ or $\sigma_2(Aa)$ is zero, a contradiction. The main idea here is that behavior strategy does not allow for correlation present in this mixed strategy.

But such correlation is strategically unnecessary. This is because information sets are reached sequentially. To make ideas precise, fix a player i and a mixed strategy σ_{-i} of other players. By specifying a behavior strategy b_i, we induce a probability distribution over the terminal (leaf) vertices of the game tree by the play (b_i, σ_{-i}). Similarly, each σ_i also induces a probability distribution over terminal vertices by the play (σ_i, σ_{-i}).

Formally, let $\rho(x; \sigma)$ denote the probability that a terminal vertex x is reached by playing a strategy profile σ. How is ρ computed? It is computed by using the conditional probability formula. Formally, it is cumbersome to state. We illustrate with the above example. In the above example, suppose Player 1 plays the behavior/mixed strategy where he plays x and y with equal probability. Suppose Player 2 plays strategy σ_2. Then what is the probability of reaching the terminal vertex with payoff $(3, 3)$? It can be reached if Player 1 plays x and

Player 2 either plays Aa or Ar. Hence, the required probability is

$$\sigma_1(x) \times \big[\sigma_2(Aa) + \sigma_2(Ar)\big] = \frac{1}{4}.$$

A similar calculation reveals the following distribution over terminal vertices

$$(\frac{1}{4}, \frac{1}{4}, \frac{1}{6}, \frac{1}{3}),$$

where we have written the probabilities of terminal vertices from left to right.

A similar calculation for behavioral strategies can also be done. It can be verified that both the mixed strategy and the behavior strategies give rise to the same distribution over terminal vertices. When computing the probability of a terminal node, we somehow constructed a behavior strategy by adding up all the pure strategies in the support of the pure strategy that lead to this terminal vertex. It so turned out that it was indeed a behavior strategy that we had earlier stated.

DEFINITION **33** *A behavior strategy b_i and a mixed strategy σ_i of Player i are* **outcome equivalent** *if for every mixed strategy σ_{-i} of other players, the probability distributions induced over the terminal vertices by (b_i, σ_{-i}) and (σ_i, σ_{-i}) are the same.*

It is safe to conjecture that *for every mixed strategy, there is a behavior strategy that is outcome equivalent to it.* The conjecture is not exactly true. Consider the extensive form game in Figure 15. Player 1 is a player without perfect recall. He has two information sets: U_1 and U_2, where U_1 is the root vertex. So, a behavior strategy of Player 1 must specify $b_1^1(x), b_1^1(y)$ such that $b_1^1(x) + b_1^1(y) = 1$ and $b_1^2(a), b_1^2(r)$ such that $b_1^2(a) + b_1^2(r) = 1$.

Now, there are four pure strategies $\{xa, xr, ya, yr\}$. Consider the mixed strategy that puts equal ($\frac{1}{2}$) probability xa and yr but zero probability on the rest. This mixed strategy induces the following distribution on terminal vertices: $\frac{1}{2}$ probability on $(3,3)$ and $(2,1)$ and zero on the rest. But to get non-zero probability on the terminal vertex $(3,3)$, Player 1 has to choose a behavior strategy at his information set U_2 which puts positive probability on a. Similarly, to reach terminal vertex $(2,1)$, he has to put positive probability on y at U_1. As a result, vertex $(1,2)$ will be reached with positive probability. So, for this mixed strategy of Player 1, there are no behavior strategy.

This is a problem due to imperfect recall. If there is perfect recall in the previous example (xa) and (ya) will become (xa) and (ya'), where a is some outcome different from a' because they belong to different information sets. Formally, Harold Kuhn established the following theorem.

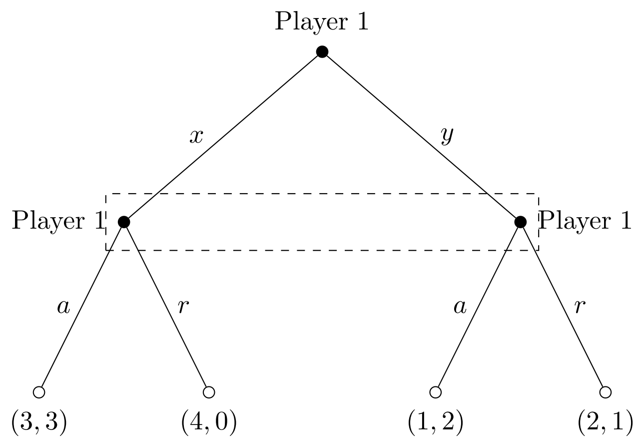

Figure 15: Extensive form game with imperfect recall

THEOREM **27** *In every extensive game, if a player i has perfect recall, then every mixed strategy of player i has an outcome equivalent behavior strategy.*

The proof involves constructing particular behavior strategies for every mixed strategy. Though the proof is notationally quite involved, the idea is relatively straightforward. We illustrate this with an example. Consider Player 2 in the game in Figure 16. Consider a mixed strategy of Player 2 as $\sigma_2(L\ell) = \sigma_2(Lr) = \frac{1}{3}$, $\sigma_2(R\ell) = \frac{1}{12}, \sigma_2(Rr) = \frac{1}{4}$. Suppose Player 1 plays p_u (for U) and p_d (for D) as his mixed strategy. Then the probability distribution induced on terminal vertices $(3,1),(3,0),(4,1),(2,2)$ respectively are $\frac{1}{3}, p_u\frac{1}{3}, p_u\frac{1}{3}, p_d\frac{2}{3}$.

Clearly, to achieve these probabilities Player 2 must play $\frac{1}{3}$ on R at his first information set. So, he plays L with probability $\frac{2}{3}$. Then, to ensure equivalent outcome, he should play ℓ and r with probability $\frac{1}{2}$ each. Hence, we computed behavior strategy of playing ℓ of Player 2 at his second information set by $\frac{\sigma_2(L\ell)}{\sigma_2(L\ell)+\sigma_2(Lr)} = \frac{1}{2}$. The proof of Kuhn's theorem formalizes this and shows that such computations are always possible.

19 Equilibria for Games of Imperfect Information

In games where there is imperfect information, subgame perfect equilibrium can still be applied but backward induction is not well-defined in such games. Moreover, subgame perfect equilibrium may be a useless solution concept in which there is imperfect information. To see this, consider the game in Figure 17. This game has only one subgame. Hence, the set of Nash equilibria are equivalent to the set of subgame perfect equilibria. The problem with subgame perfect equilibrium in this game is that it does not use any *beliefs* of Player 2. As a result, it puts no restriction on his optimal choice when his information set is reached. To

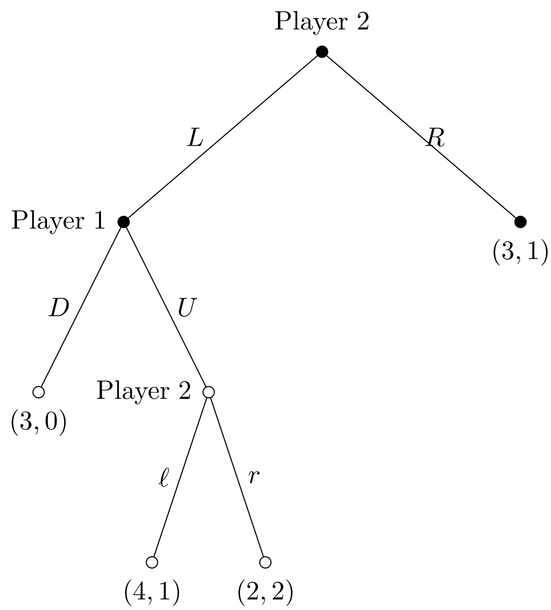

Figure 16: Extensive form game with perfect recall

appropriately define behavior in information sets, any equilibrium must also define beliefs and equilibrium choices must be consistent with these beliefs. This is the basic idea behind defining equilibrium refinements in games of imperfect information.

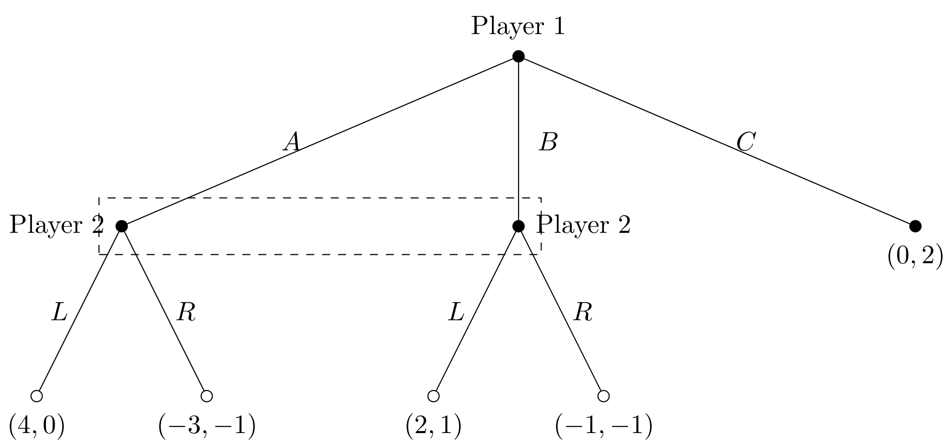

Figure 17: Imperfect Information

19.1 Perfect Bayesian Equilibrium

To understand the problem with subgame perfect equilibrium further in such games, consider the reduced-form strategic-form game of the game in Figure 17. It is shown in Table 24.

	L	R
A	$(4,0)$	$(-3,-1)$
B	$(2,1)$	$(-1,-1)$
C	$(0,2)$	$(0,2)$

Table 24: Reduced strategic form of the game in Figure 10

The Nash equilbria of this strategic-form game consists of (A,L), $(C, \alpha L + (1-\alpha)R)$, where $\alpha \leq \frac{1}{3}$. The idea of *sequential rationality* requires that each player must behave rationally once his information set is reached. To be able to do this, players must form beliefs about where they are inside their information set, and act optimally according to this belief. The nature of beliefs that is permissible results in different solution concepts.

For instance, if we specify a strategy profile, where Player 1 plays A with probability $\frac{1}{3}$ and B with probability $\frac{1}{2}$, then this equilibrium knowledge is enough to pin down the beliefs of Player 2. Remember, that Player 2 has correct belief about equilibrium behavior of Player 1. Hence, his belief of the information set can be deduced from this: total probability of reaching this information set is $\frac{5}{6}$, and individual conditional probabilities are $(\frac{2}{5}, \frac{3}{5})$. Of course, here we cannot apply this principle if a strategy profile does not reach a particular information set since conditional probabilities are not defined at those information sets. So, sequential rational behavior can be with respect to *any* belief at such information sets.

Formally, in an extensive form game with imperfect information, the belief of Player i is a map $\mu_i : U_i^j \rightarrow [0,1]$ for each j such that $\sum_{x \in U_i^j} \mu(x) = 1$ for all j.

Given a strategy profile σ, we can compute the probability with which each decision vertex is reached in an extensive form game. We denote this as $P_\sigma(x)$. The probability with with an information set U_i^j is reached given σ is $P_\sigma(U_i^j) = \sum_{x \in U_i^j} P_\sigma(x)$.

Definition **34** *Belief μ_i of Player i is* **Bayesian** *given a strategy profile σ if for every information set U_i^j reached with positive probability in the strategy profile σ, we have for all $x \in U_i^j$,*

$$\mu_i(x) = \frac{P_\sigma(x)}{P_\sigma(U_i^j)}.$$

Sequential rationality now extends to this setting as follows.

Definition **35** *A strategy σ_i of Player i at information set U_i^j is* **sequentially rational** *given strategies σ_{-i} and beliefs μ_i if for all σ_i', we have*

$$\sum_{x \in U_i^j} \mu_i(x) u_i(\sigma_i, \sigma_{-i}|x) \geq \sum_{x \in U_i^j} \mu_i(x) u_i(\sigma_i', \sigma_{-i}|x).$$

A strategy σ_i of Player i is **sequentially rational** *given σ_{-i} and μ_i if it sequentially rational at all information sets.*

An equilibrium here now involves specifying strategies and beliefs. Beliefs have to be consistent in the form of Bayesian and strategies have to be sequentially rational. The pair of strategy profile and belief profile is called an *assessment*.

Definition **36** *An assessment (σ, μ) is a* **perfect Bayesian equilibrium (PBE)** *if for every Player i*

- *μ_i is Bayesian given σ*
- *σ_i is sequentially rational given σ_{-i} and μ_i.*

In the game in Figure 17, for every belief of Player 2, L is a weakly dominant action. Given this, Player 1 must play A irrespective of his beliefs. Hence, the unique PBE of this game is $(A, L, \mu_2(B) = 1)$. In general, a PBE does not allow players to play a strictly dominated action, while a Nash equilibrium does not preclude this.

Remark. An easy fact to check is that Nash equilibrium requires optimal action (sequential rationality) on the path of play - so, it is silent on behavior on information sets that are not reached. The extra thing that PBE brings is optimal behavior on information sets that are not reached for *some* belief. Hence, σ is Nash equilibrium if and only if for every Player i (i) there are beliefs μ_i that are Bayesian given σ_{-i} and (ii) for each information set U_i^j with $P_\sigma(U_i^j) > 0$, we have σ_i is sequentially rational given σ_{-i} and μ_i. Hence, PBE is a *refinement* of Nash equilibrium.

Theorem **28** *Every PBE is a Nash equilibrium.*

However, PBE allows for any arbitrary beliefs off equilibrium path. This can lead to unsatisfactory predictions in certain games. The following example illustrates this. Consider the game in Figure 18. In this game, what beliefs of Player 2 induce him to play ℓ? Suppose he puts μ probability on his left decision vertex and $(1 - \mu)$ on the other. Then, his payoff by playing ℓ is $2 - \mu$ and his payoff from playing r is $3 - 4\mu$. So he plays ℓ if $\mu > \frac{1}{3}$, r if $\mu < \frac{1}{3}$, and mixes ℓ and r otherwise. But Player 1 plays his dominant strategy D in his second information set. So, what should Player 1 play in PBE in the first information set? Suppose he mixes $\alpha L + (1-\alpha)R$, where $\alpha > 0$. Then, $\mu = 1$ is the only Bayesian belief - note this information set is reached in equilibrium now. Then Player 2 must play ℓ. This means that $\alpha = 1$. If Player 1 plays R, then any belief is allowed for Player 2. But for Player 1 to choose R in equilibrium, Player 2 must play r - if he plays ℓ, then he is better of choosing L and then D to get payoff 2. For Player 2 to play r, the belief should be $\mu \leq \frac{1}{3}$. There are other PBE where Player 2 mixes also.

Now, let us consider the PBE $((R, D), r; \mu \leq \frac{1}{3})$. It is not reasonable to assume that Player 2 plays r in his information set since he knows that U is never played by Player 1. Another amazing feature of this game is its subgame perfect equilibrium. The subgame starting with the second information set of Player 1 has one Nash equilibrium - Player 1 chooses his dominant strategy D and Player 2 best responds with ℓ. Given this, Player 1 chooses L in the first information set. Hence, $((L, D), \ell)$ is a unique subgame perfect equilibrium of this game. Thus, the PBE is *not* a refinement of subgame perfect equilibrium.

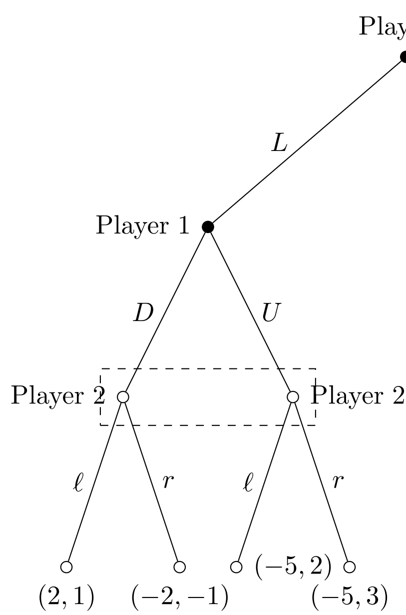

Figure 18: Problems with PBE

19.2 Sequential Equilibrium

To get rid of this unpleasant feature of PBE, a refinement is proposed. The refinement aims to put some consistent beliefs on information sets that are not reached in equilibrium.

Definition **37** *An assessment* (σ, μ) *is a* **sequential equilibrium** *if*

1. μ *is* **consistent** *given* σ*: There exists a sequence of completely mixed strategy profile* $\{\sigma^k\}_k$ *such that (i)* $\lim_k \sigma^k = \sigma$ *and if* μ^k *are the unique Bayesian beliefs for* σ^k*, then* $\lim_k \mu^k = \mu$.

2. σ *is sequentially rational given* μ.

The extra condition here from PBE is consistency, which requires that if Players make some small mistakes from equilibrium, the beliefs should be close to the Bayesian beliefs corresponding to those small mistakes. Note that the sequence we construct need not be unique, and different sequences may lead to different beliefs.

In extensive form games with imperfect information, the one-shot deviation principle continues to hold. Hence, in such games, it is enough to check for deviations at one information set at a time.

The following theorem, whose proof we skip, establishes that a sequential equilibrium is refinement of subgame perfect equilibrium.

Theorem **29** *Every sequential equilibrium is a subgame perfect equilibrium. Every completely mixed strategy Nash equilibrium is a sequential equilibrium.*

The second part of Theorem 29 follows trivially by taking the sequence of strategies same as the equilibrium strategy.

Let us now revisit the game in Figure 18. First, look at the subgame perfect equilibrium $((L, D), \ell)$. If we consider mixed strategies, where $\sigma_1^k(R) = \epsilon_R^k, \sigma_1^k(L) = 1 - \epsilon_R^k$ and $\sigma_1^k(D) = 1 - \epsilon_D^k, \sigma_1^k(U) = \epsilon_D^k$. Then,

$$\mu = \frac{(1 - \epsilon_k^D)(1 - \epsilon_R^k)}{1 - \epsilon_R^k} \to 1.$$

Note that perturbation of Player 2's strategy is not necessary here. Hence, $\mu = 1$ is a consistent belief given this strategy profile. We already know that this strategy profile is sequentially rational given μ. Hence, it is a sequential equilibrium.

Now, can there be a sequential equilibrium where Player 1 chooses (R, D) and Player 2 chooses r. If we perturb the strategies of Player 1, then we reach the information set of Player 2 with positive probability where the belief on the (L, D) decision vertex must be very high. As a result, Player 2 must choose ℓ here to be sequentially rational. Hence, no sequential equilibrium will choose Player 2 playing r with positive probability if Player 1 plays (R, D).

A comment about existence of PBE and sequential equilibrium is that if games have perfect recall, then these equilibria always exist.

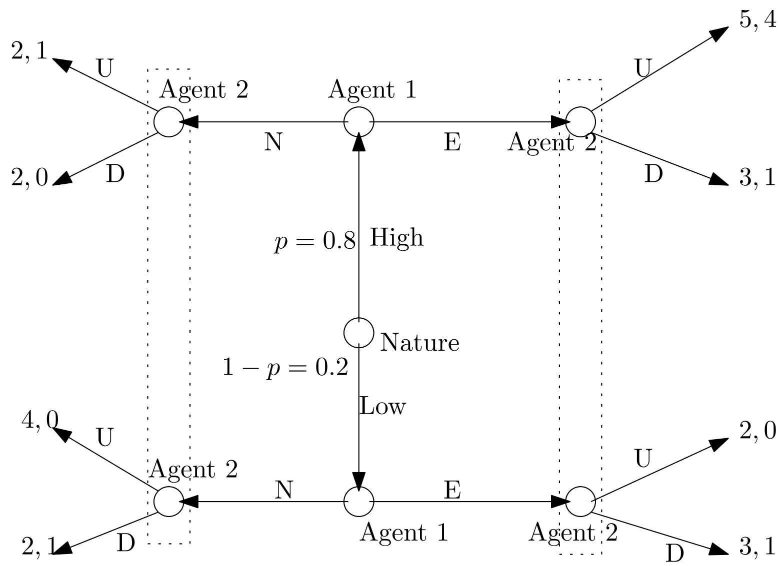

Figure 19: Signaling game

19.3 Example: A signaling game

We give an example to illustrate the notions of PBE and sequential equilibrium. This example is usually called a simpler version of the *signaling game*. There are two agents in this example - see Figure 19. Agent 1 has two types - High or Low, their probabilities are as shown in Figure 19. Agent 1's type is not observed by Agent 2 but his action, which is either N or E, is observable by Agent 2. After observing Agent 1's action, Agent 2 takes an action, which is either U or D. The payoffs are as shown in Figure 19.

We now compute some of the PBE of this game. Before doing so, we observe that Agent 1 of type High strictly prefers E to N. Hence, in any PBE, Agent 1 must choose E at his decision vertex corresponding to High type. We now look at various PBE of this game. Denote the belief of Agent 2 on his left information set as μ_L for the top decision vertex and $1 - \mu_L$ for the bottom decision vertex. Similarly, denote the belief of Agent 2 on his right information set as μ_R for the top decision vertex and $1 - \mu_R$ for the bottom decision vertex.

- **Separating PBE.** High type Agent 1 chooses E but Low type Agent 1 chooses N. If such a PBE exists, then all the information sets of Agent 2 is reached in equilibrium. By Bayesian rationality, Agent 2's belief must satisfy: $\mu_L = 0, \mu_R = 1$. Then, sequential rationality of Agent 2 implies that he must choose D in the left information set and U in the right information set. Finally, we verify that Agent 1 is sequentially rational. As argued, the High type choosing E is sequentially rational. For the Low type, choosing N gives a payoff of 2 and choosing E gives a payoff of 2 also. Hence, Agent 1's strategy is sequentially rational. So, we can describe the separating PBE as:

$$(High : E, Low : N, Left : D, Right : U, \mu_L = 0, \mu_R = 1).$$

This PBE is trivially a sequential equilibrium since every information set is reached with positive probability in this equilibrium.

- **Pooling PBE.** Both High and Low type Agent 1 choose E. If such a PBE exists, then left information set of Agent 2 is not reached in equilibrium and right information set is reached with probability 1. By Bayesian rationality, Agent 2's belief in right information set must be: $\mu_R = p = 0.8$. Then, sequential rationality of Agent 2 in the right information set implies he must choose U: choosing U gives a payoff equal to 0.8(4) compared to a payoff of 1 by choosing D. For Agent 1 to choose N when he is of Low type, Agent 2 must choose D - this is because if Agent 2 chooses U, then Agent 1 is better off choosing N when he is of Low type. So, sequential rationality of Low type Agent 1 forces Agent 2 to choose D in his left information set. But such a choice is possible with sequential rationality if $1 - \mu_L \geq \mu_L$ or $\mu_L \leq 0.5$.

  Hence, there is a class of pooling PBE:

  $$(High : E, Low : E, Left : D, Right : U, \mu_L \leq 0.5, \mu_R = p = 0.8).$$

  Any such PBE is also a sequential equilibrium. Fix a particular PBE with a particular value of μ_L. For this, we think of a perturbation of Agent 1's actions to reach the left information set of Agent 2. But this perturbation must generate beliefs μ_L in the limit. A possible way to generate this belief is to choose perturbations as follows:

  $$High : \epsilon' N + (1 - \epsilon')E; Low : \epsilon N + (1 - \epsilon)E,$$

  where $\epsilon' = \epsilon \frac{\mu_L}{4(1-\mu_L)}$. Notice that this choice of ϵ and ϵ' exactly generates μ_L belief by Bayesian rationality. Hence, as $\epsilon \to 0$ (and, hence, $\epsilon' \to 0$), we get the beliefs approaching μ_L.

- **Mixing at Low type.** High type Agent 1 chooses E but Low type agent mixes N and E. If such a PBE exists, then let Low type Agent 1 mixes as $\sigma_E E + (1 - \sigma_E)N$, where $\sigma_E \in (0, 1)$. As a result, all information sets of Agent 2 is reached in equilibrium. Bayesian rationality implies that

  $$\mu_L = 0, \mu_R = \frac{0.8}{0.8 + 0.2\sigma_E}.$$

Then, sequential rationality of Agent 2 requires that he must choose D in the left information set. Sequential rationality of Agent 1 at Low type requires that he must be indifferent between N and E (because he mixes). This is only possible if Agent 2 chooses U at his right information set. But then, $4\mu_R \geq 1$ or $3.2 \geq 0.8 + 0.2\sigma_E$ or $\sigma_E \leq 1.2$, which is always true. Hence, independent of the mixing probability of Agent 1 of Low type, Agent 2 prefers U at his right information set. So, for any $\sigma_E \in (0, 1)$, we have the following PBE:

$$(High : E, Low : \sigma_E E + (1 - \sigma_E)N, Left : D, Right : U, \mu_L = 0, \mu_R = \frac{0.8}{0.8 + 0.2\sigma_E}).$$

Since every information set is reached with positive probability in such PBE, they are also sequential equilibria.

20 Repeated Games

20.1 Basic Ideas - The Repeated Prisoner's Dilemma

Consider the Prisoners' Dilemma (PD) game in Table 25. Recall that a dominant strategy equilibrium of this game is (L_1, L_2), and it is the unique Nash equilibrium of the game.

	L_2	R_2
L_1	2,2	6,1
R_1	1,6	5,5

Table 25: Prisoner's Dilemma

Now, suppose the game is played twice with the actions at the end of every stage is observed by all the players, and the payoff of a player at the end of the game is the sum of payoff at the end of each stage. The game can be represented in extensive form now. A subgame perfect equilibrium of this extensive form game requires that the players play a Nash equilibrium in the second stage, and they play a Nash equilibrium of the entire game. Since the unique Nash equilibrium of the game is (L_1, L_2), the players will play (L_1, L_2) in second stage in any subgame pefect equilibrium. Given this, the players now know that they will get a payoff of 1 in the second stage. So, the we can add $(1, 1)$ to the payoff matrix in the first stage, and then compute a Nash equilibrium. This still gives a unique Nash equilibrium of (L_1, L_2). Hence, the outcome of this game in a subgame perfect equilibrium is (L_1, L_2).

This argument can be generalized. Let $G = (N, \{A_i\}_{i \in N}, \{u_i\}_{i \in N})$ denote a strategic-form game of complete information. The game G is called the **stage game** of the repeated game.

Definition 38 *Given a stage game G, let $G(T)$ denote the* **finitely repeated game** *in which G is played T times with actions taken by of all players in the preceding stages observed before the play in the next stage, and payoffs of $G(T)$ are simply the sum of payoffs in all T stages.*

Our arguments earlier lead to the following proposition (without formally defining notions of equilibrium).

Proposition 1 *If the stage game G has a unique Nash equilibrium, then for any finite repetition of G, the repeated game $G(T)$ has a unique subgame perfect outcome: the Nash equilibrium of the stage game G is played in every stage.*

There are two important assumptions here: (a) the stage game has a unique Nash equilibrium and (b) the stage game is repeated finite number of times. We will see that if either of the two assumptions are not present then it is possible for players to get better payoffs.

We now modify the PD game by introducing a new strategy for every player. The new PD game is shown in Table 26. There are two Nash equilibria of this game: (L_1, L_2) and (R_1, R_2).

	L_1	M_1	R_1
L_2	1,1	5,0	0,0
M_2	0,5	4,4	0,0
R_2	0,0	0,0	3,3

Table 26: A Game with Multiple Nash Equilibrium

Now, suppose the stage game in Table 26 is repeated twice. Then, using the arguments earlier, we can say that in every stage playing either of the Nash equilibria is subgame perfect. But, we will show that there exists a subgame perfect equilibrium in which (M_1, M_2) is played in the first stage.

Consider the following strategy of the players: if (M_1, M_2) is played in the first stage, then play (R_1, R_2) in the second stage; if any other outcome happens in the first stage, then play (L_1, L_2) in the second stage. This means, in the first stage of the game, the players are looking at a payoff table as in Table 27, where second stage payoff $(3, 3)$ is added to (M_1, M_2) and second stage payoff $(1, 1)$ is added to all other strategy profiles. The addition of different payoffs to different strategy profiles changes the equilibria of this game. Now, we have three pure strategy Nash equilibria in Table 27: (L_1, L_2), (M_1, M_2), and (R_1, R_2). Hence, $((M_1, M_2), (R_1, R_2))$ constitute a subgame perefect equilibrium of this repeated game. Thus, existence of multiple Nash equilibrium in the stage game allowed us to achieve cooperation in the fist stage of the game. Notice that (M_1, M_2) is not a Nash equilibrium of the stage game.

	L_1	M_1	R_1
L_2	2,2	6,1	1,1
M_2	1,6	7,7	1,1
R_2	1,1	1,1	4,4

Table 27: Analyzing Payoffs of First Stage

This is part of a general argument: if G is a static game of complete information with multiple Nash equilibria, there may be subgame perfect outcomes of the finitely repeated game $G(T)$ in which for any stage $t < T$, the outcome in stage t is not a Nash equilibrium.

20.2 A Formal Model of Infinitely Repeated Games

Let $G \equiv (N, \{A_i\}_{i\in N}, \{u_i\}_{i\in N})$ be a strategic form game. When we repeat such a stage game G, we will assume that players observe all the actions taken in each period. At any period, let a^t denote the action profile chosen by players. The sequence of actions profile $(a^1, \ldots, a^{t-1})$ that leads to current period will be called the history of period t.

An **infinitely repeated game** of G is defined by $G^\infty \equiv (G, H, \{u_i^*\}_{i\in N})$, where

- $H = \cup_{t=1}^{\infty} A^t$ are the set of all possible histories, with $A^1 \equiv \emptyset$ denoting the null history, A^t denoting the possible histories till period t, and A^∞ denoting all infinite length histories.
- $u_i^* : A^\infty \to \mathbb{R}_+$ for every $i \in N$ is a utility function that assigns every infinite history a payoff for Player i.

A history is terminal if and only if it is infinite. Note that an infinitely repeated game is a special type of infinite extensive game.

Strategies in a Repeated Game.

What is a strategy of a player in an infinitely repeated game? Remember, a strategy needs to assign an action for every *possible* situation. This means that we need to assign an action at every period for every possible history. Thus, strategy of Player i is a collection of infinite maps $\{s_i^t\}_{t=1}^{\infty}$, where

$$s_i^t : A^t \to A_i.$$

Since a strategy seems to be a really complicated (infinite) object here, it is difficult to imagine it. One easy way to think of a strategy is a machine (or automaton). The machine for Player i has the following components.

- A set Q_i of **states**.
- An element $q_i^0 \in Q_i$, indicating the initial state.
- A function $f_i : Q_i \to A_i$ that assigns an action to every state.
- A transition function $\tau_i : Q_i \times A \to Q_i$ that assigns a state for every state and every action profile.

States represent situations that Player i cares about. We give an example showing how a strategy in Prisoner's Dilemma can be modeled as a machine. The strategy we consider is called a **trigger** strategy. It chooses the cooperate action C as long as the history consists of all players choosing C. Else, it chooses D. We only care about two states here: whether everyone chosen C in the past or not. We will denote this as $\mathcal{C}$ and $\mathcal{D}$ respectively. Since we want to choose C in the first period, we set $q_i^0 := \mathcal{C}$. Now, $f_i(\mathcal{C}) = C$ and $f_i(\mathcal{D}) = D$. The transition function looks as follows:

$$\tau_i(\mathcal{C}, (C, C)) = \mathcal{C}, \tau(\mathcal{X}, (X, Y)) = \mathcal{D} \text{ if } (\mathcal{X}, (X, Y)) \neq (\mathcal{C}, (C, C)).$$

This is an example of a strategy which is relatively simple. Note that the number of states here is finite. As one can see that we can construct strategies that care about more number of states (possibly infinite). For our purposes, the kinds of strategies that we will use will require machines with finite state space.

Payoffs in Repeated Games.

Fix a strategy profile of players $s \equiv (s_1, \ldots, s_n)$. This strategy profile leads to outcomes in each stage/period. Denote by v_i^t, the payoff due to this strategy profile in period t. So, agent i has an infinite stream of payoffs $\{v_i^t\}_{t=1}^{\infty}$ from this strategy profile. Similarly, if there is another strategy profile s', then it will generate an infinite stream of payoffs $\{w_i^t\}_{t=1}^{\infty}$. As a result, if Player i has to compare outcomes of two strategy profiles, it compares two infinite streams of payoffs: $\{v_i^t\}_{t=1}^{\infty}$ and $\{w_i^t\}_{t=1}^{\infty}$.

There are many ways to make this comparison. The most standard way is to use a *discounted criterion*. In this way, we have a discount factor $\delta \in (0,1)$ which is same for all the players. Player i attaches a payoff equal to

$$\sum_{t=1}^{\infty} \delta^{t-1} v_i^t,$$

to the payoff stream $\{v_i^t\}_{t=1}^{\infty}$. For instance, if there is a payoff stream that generates payoffs $v \equiv (1,1,1,\ldots)$, then the payoff from this stream is $1(1+\delta+\delta^2+\ldots) = \frac{1}{1-\delta}$. Note that even though the payoff is 1 in each period, we get a higher payoff overall. It is often convenient to assign a payoff of

$$(1-\delta)\sum_{t=1}^{\infty} \delta^{t-1} v_i^t,$$

to the payoff stream $\{v_i^t\}_{t=1}^{\infty}$. This normalizes the payoff and makes it easy to compare it with the stage game payoff. Note that comparisons across two infinite stream of payoffs still remain the same.

Obviously, discounting puts different weights on payoffs of different periods. Particularly, future is valued less than present. Note that changes in payoff in a single period may matter in the discounting criteria. To see this, compare $v \equiv (1,1,\ldots)$ and $w \equiv (1+\epsilon, 1-\epsilon, 1-\epsilon, \ldots)$, where $\epsilon \in (0,1)$. Payoff from v is 1 and payoff from w is $(1+\epsilon)(1-\delta)+(1-\epsilon)\delta = 1+\epsilon-2\epsilon\delta = 1+\epsilon(1-2\delta)$. This is greater than 1 if and only if $\delta > \frac{1}{2}$. [6]

Similarly, look at the payoff streams $v \equiv (1,-1,0,0,\ldots)$ and $w \equiv (0,0,0,\ldots)$. The payoff from w is zero but the payoff from v is $(1-\delta)^2$. Hence, for any $\delta \in (0,1)$, v is preferred to w. However, consider the stream $v' \equiv (-1,1,0,0,\ldots)$. This generates a payoff of $(1-\delta)(-1+\delta) = -(1-\delta)^2$. Hence, v' is worse than w. This shows that the discounting puts more emphasis on current payoffs than future payoffs.

This is contrasted in the following two streams of payoffs $v \equiv (0,0,0,\ldots,1,1,1,\ldots)$ and $w \equiv (1,0,0,\ldots)$. The payoff stream v has M zeros and then all 1s. The payoff from v is δ^M and from w is $(1-\delta)$. For every δ, there is a M such that w is preferred to v. But for a fixed M, we can find δ close to 1 such that v is preferred to w.

Given a strategy profile, $s \equiv (s_1,\ldots,s_n)$, we get a unique stream of action profiles $\{a^t\}_{t=1}^{\infty}$ associated with this strategy profile. Note how this action profile is obtained - first, each player i plays $a_i^1 := s_1^1(\emptyset)$. Having generated the action profiles $h^t \equiv (a^1,\ldots,a^{t-1})$, player i

[6]Sometimes, discounting is interpreted differently. A discount δ means that the stage game continues to next period with probability δ.

plays $a_i^t \equiv s_i^t(h^t)$. From this, we can compute the utility of Player i as

$$u_i^*(s) := (1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}u_i(a^t).$$

Having defined strategies and payoffs, we are now ready to define the equilibrium concepts for repeated games.

Definition **39** *A strategy profile $s \equiv (s_1, \dots, s_n)$ is a* **Nash equilibrium** *of the infinitely repeated game G^∞ if for every $i \in N$, for every s_i', we have*

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

A strategy profile s is a **subgame perfect equilibrium** *if its restriction from any period t is a Nash equilibrium of the subgame starting from that period.*

20.3 Folk Theorems: Illustrations

There are two interesting take-aways from the results of repeated games. First, repeated games allow for a large set of payoffs to be achieved in Nash and subgame perfect equilibrium. Such theorems are called Folk Theorems. The second take-away is the kind of strategies that support such equilibrium payoffs. Such strategies are very common in many social interactions. To be able to establish folk theorems using such common real-life strategies give a strong foundation for such results.

We will now illustrate the basic idea behind the folk theorems using the Prisoner's Dilemma example - see Table 28. We first show that there are subgame perfect equilibria where cooperation can be achieved.

	L_2	R_2
L_1	1,1	-1,2
R_1	2,-1	0,0

Table 28: Prisoner's Dilemma

Proposition **2** *Suppose $\delta \geq \frac{1}{2}$. Then, there is a subgame perfect equilibrium in the Prisoner's Dilemma game (Table 28), where both the players play (L_1, L_2) in every period.*

Proof: We describe the following strategy. Each player i follows L_i if the history consists of both players playing (L_1, L_2). If the history is different from (L_1, L_2) play in each period in the past, i plays R_i. The strategy stated here is called a *trigger strategy*. Fix Player 1 and assume that Player 2 is following the trigger strategy stated in the Proposition. We show that following the trigger strategy is optimal for Player 1. We need to consider two types of subgames.

CASE 1. We consider a subgame where the history so far has been (L_1, L_2). In that case, following L_1 gives Player 1 a payoff of 1. Playing R_1 in some periods has the following consequence. In the first period he plays R_1 he gets a payoff of 2 since Player 2 plays L_2. But in subsequent periods Player 2 plays R_2. So, he gets a maximum payoff of 0. As a result, his payoff is less than $(1-\delta)\big(1+\delta+\ldots+\delta^{t-1}+2\delta^t\big)$, where t is the first period from this subgame where he deviates. Remember the truthful payoff stream is $(1,1,1,\ldots)$. The deviated payoff stream payoff is less than the payoff stream $(1,1,\ldots,2,0,0,0,\ldots)$. Then, it is sufficient to compare the payoff streams $(1,1,1,\ldots)$ and $(2,0,0,\ldots)$. The later one gives a payoff of $2(1-\delta)$. But $\delta \geq \frac{1}{2}$ implies that $1 \geq (1-\delta)2$. Hence, no deviation is profitable in this subgame.

CASE 2. We consider a subgame where the history involves action profiles other than (L_1, L_2). In that case, Player 2 is repeatedly playing R_2 in this subgame. But if Player 2 is playing R_2, Player 2 gets a payoff stream of $(0,0,\ldots)$ by Playing R_1 in every period but gets a payoff stream where in every period he gets payoff less than or equal to 0 by playing some other strategy.

Hence, the specified strategy is a Nash equilibrium in this subgame. ■

20.4 Nash Folk Theorem

The trigger strategies used in Proposition 2 can be used to establish a general result about what payoffs can be achieved in a Nash equilibrium of G^{∞}.

The important payoff for folk theorems is the minmax value. Define the **minmax value** of player i in the stage game G as

$$\underline{v}_i = \min_{a_{-i}} \max_{a_i} u_i(a_i, a_{-i}),$$

where (a_i, a_{-i}) denotes an action profile of the stage game. [7] This is the minimum payoff player i can be held to by its opponents (using pure actions), given that he plays best response

[7]The minmax and maxmin payoff of a player can be quite different. Please construct examples to see that the minmax is different from maxmin. In the early parts of the lectures, I used $\underline{v}_i$ to denote the maxmin

to the action profile a_{-i}. Let $u_i(\underline{a}_i, \underline{a}_{-i}) = \underline{v}_i$ for player i. Then, we call $\underline{a}_{-i}$ the **minmax action profile** against player i.

The reason minmax values are important is the following lemma.

LEMMA **6** *Player i's payoff is at least $\underline{v}_i$ in any pure action Nash equilibrium of the stage game G and the infinitely repeated game G^∞, regardless of the value of δ.*

Proof: Let a be a Nash equilibrium of the stage game. Then for every $i \in N$,

$$u_i(a) = \max_{a_i} u_i(a_i, a_{-i}) \geq \max_{a_i} \min_{a'_{-i}} u_i(a_i, a'_{-i}) = \underline{v}_i.$$

Hence, Player i's payoff is at least $\underline{v}_i$ in any Nash equilibrium of the stage game.

Now, suppose player i plays a best response to the actions of other players in each period of G^∞. This guarantees him $\underline{v}_i$ in every period irrespective of the strategy played by other players. Hence, a player i is guaranteed of a payoff of $\underline{v}_i$ by this strategy in G^∞. So, any strategy that does not guarantee $\underline{v}_i$ will have a deviation where Player i just best responds to the actions of other players in every period. ■

Hence, Player i is guaranteed to get at least $\underline{v}_i$ payoff in any pure action Nash equilibrium of the repeated game.

DEFINITION **40** *A payoff profile $v = (v_1, \ldots, v_n)$ is* **strictly enforceable** *if for every $i \in N$, we have $v_i > \underline{v}_i$.*

We now give a weaker version of Folk Theorem.

THEOREM **30 (Pure Nash Folk Theorem)** *Suppose v is a strictly enforceable payoff profile and there exists an action profile a in the stage game G such that $u_i(a) = v_i$ for all $i \in N$. Then, there exists a $\underline{\delta}$, such that for all $\delta \geq \underline{\delta}$, there is a Nash equilibrium of G^∞ with discount δ where a is played in every period.*

Proof: Suppose v is a strictly enforceable feasible payoff profile and there exists an action profile a in the stage game G such that $u_i(a) = v_i$ for all $i \in N$. Consider the following strategy. It is described by three states: (a) normal state (b) i-punishment state, and (c) more-punishment state. The initial state is normal state. In normal state, the strategy recommends playing a_i to each Player i. Now, we inductively define the states at every

payoff of Player i in a strategic form game, but here I use it for minmax payoff of Player i. I apologize for this confusion.

history. For every history, there is a unique *predecessor* history where actions lead to the current history. If the predecessor history is normal and everyone plays the recommended action profile a, then the state remains normal.

If the predecessor history is normal and a single player i does not play a_i, then the current history becomes i-punishment state. If the predecessor history is i-punishment, then the current history remains i-punishment irrespective of the action taken by the players. In the i-punishment state, the strategy recommends playing the minmax action profile of Player i.

If the predecessor history is normal and more than one player do not play a, then the current history becomes more-punishment state. If the predecessor history is more-punishment, then the current history remains more-punishment irrespective of the action taken by the players. In the more-punishment state, the strategy recommends playing a fixed action profile - this need not minmax any particular player.

The strategy is shown in Table 29.

Predecessor state	Action profile in predecessor	Current state	Recommended action profile
Normal	a	Normal	a
Normal	(a'_i, a_{-i})	i-punishment	Minmax for Player i
Normal	$(a'_S, a_{N\setminus S})$ with $\lvert S\rvert > 1$	more-punishment	Any fixed action profile
i-punishment	a'	i-punishment	Minmax for Player i
more punishment	a'	more punishment	Any fixed action profile

Table 29: Trigger strategy for Nash folk theorem

To see this strategy profile can be sustained in Nash equilibrium, first observe that the payoff from equilibrium is v_i for Player i. Suppose all the other players except i follows the prescribed strategy. Let $\bar{v}_i = \max_{a'_i \in A_i} u_i(a'_i, a_{-i})$. If Player i deviates, then he gets a maximum payoff of $\bar{v}_i$. This maximum payoff he gets in the first period he deviates and thereafter he is punished, and hence, gets a payoff of $\underline{v}_i$. Hence, if he deviates in period t, his maximum possible payoff from deviation is

$$(1-\delta)\big(v_i + \delta v_i + \ldots + \delta^{t-1}\bar{v}_i + \delta^t \underline{v}_i + \delta^{t+1}\underline{v}_i + \ldots\big)$$

For deviation to be not profitable, we need to ensure that

$$v_i \geq (1-\delta)\big(v_i + \delta v_i + \ldots + \delta^{t-1}\bar{v}_i + \delta^t \underline{v}_i + \delta^{t+1}\underline{v}_i + \ldots\big).$$

Expanding the LHS, we get

$$(1-\delta)\big(v_i + \delta v_i + \delta^2 v_i + \ldots\big).$$

Canceling common terms in expanded LHS and RHS, we need to ensure that

$$\delta^{t-1}\bar{v}_i + \delta^t \underline{v}_i + \delta^{t+1}\underline{v}_i + \ldots \leq \delta^{t-1}v_i + \delta^t v_i + \delta^{t+1}v_i + \ldots.$$

This means, we need to ensure that $\bar{v}_i(1-\delta) + \delta\underline{v}_i \leq v_i$.

This is equivalent to ensuring

$$\delta \geq \frac{\bar{v}_i - v_i}{\bar{v}_i - \underline{v}_i}.$$

Define

$$\underline{\delta} := \frac{\bar{v}_i - v_i}{\bar{v}_i - \underline{v}_i}.$$

Note that by assumption $\bar{v}_i > v_i > \underline{v}_i$. Hence, $\underline{\delta} \in (0,1)$. This proves the claim. ■

The exact version of folk theorems will be discussed later - they involve use of mixed behavior strategies by players.

One of the issues with the Nash folk theorem is the strategies required to sustain the Nash equilibrium is very extreme - it requires you to punish the deviant for infinite number of periods. This may not be a reasonable threat. For instance, consider the game in Table 30. Theorem 30 says that (T, L) is achievable in Nash equilibrium of G^∞ for sufficiently patient players as long as the Column player can punish deviations by action R. This will hurt the Row player but the Column player is also badly hurt. This motivates the next set of results that require subgame perfect equilibrium - even punishments need to happen in equilibrium.

	L	R
T	6,6	0,-100
B	7,1	0,-100

Table 30: A Stage game

20.5 The One-Shot Deviation Principle

The one-shot deviation principle is a useful tool in the repeated games setting. Two strategies s_i and s'_i are one-shot deviations of each other if they differ from each other by actions chosen at one period for one history, i.e., $s_i^t(h^t) \neq \bar{s}_i^t(h^t)$ but $s_i^{t'}(h^{t'}) = \bar{s}_i^{t'}(h^{t'})$ for all $(t', h^{t'}) \neq (t, h^t)$. The one-shot deviation principle says that, fixing Player i and strategies s_{-i} of other players, if strategy s_i of Player i is optimal over all strategies $\bar{s}_i$ that are one-shot deviations from s_i, then it is optimal over all strategies.

To see why the one-shot deviation principle is true, consider Player i by fixing the strategies of other players at s_{-i}. Suppose strategy s_i is optimal over all one-shot deviations. Suppose another strategy s'_i differs from s_i at *finite* set of decision vertices (i.e., periods and histories). Then, we go the last period t where s_i and s'_i differ at some history h^t. In this subgame, s_i and s'_i differ from each other by one-shot deviation. Hence, s'_i cannot be profitable in this subgame. So, all the gains from s'_i must be occurring before this period. So, we restore s'_i to s_i in all histories in this period. We inductively repeat this procedure to reach a stage where s_i and s'_i are one-shot deviations. This is the same argument we have done for the backward induction procedure. Indeed, we did not use any specifics of repeated games in this argument.

The difference here is that s_i and s'_i can differ from each other at infinite number of decision vertices. Here, the discounted criteria of repeated games rescue us. Suppose strategy s_i is suboptimal. Then, there is some history h^t after which Player i can make a sequence of different moves than those prescribed by s_i. If the number of such different moves is finite, the previous argument applies. Else, let γ be the gain of Player i from this deviation, which starts in period t at history h^t. Let M be the best conceivable one-period *gain* in payoff to Player i by deviating from s_i. Choose a period $s > t$ such that $\delta^{s-t}M < \frac{\gamma}{2}$ - note that since M is finite and $\delta \in (0, 1)$, we can find such a s. Note that $\delta^{s-t}M$ is the maximum possible payoff gain from period s onwards - here, instead of multiplying the payoff in period t by δ^{t-1}, we multiply by 1 as if the game started from period t. Thus, gain from period s onwards cannot be more than $\frac{\gamma}{2}$. So, gain from period t to s must be at least $\frac{\gamma}{2}$. So, s_i can be modified at finite decision vertices such that we get a new strategy s''_i that is better than s_i. Moreover, s''_i differs from s_i at finite histories. But this contradicts our earlier argument that one-shot deviation principle guarantees deviations at finite decision histories.

20.6 Perfect Folk Theorem - Reversion to Nash

To make punishments credible, we must require Nash equilibrium at every subgame. This is the main motivation for using subgame perfect equilibrium. For every history, players must be playing Nash equilibrium actions. The following is quite immediate.

Proposition **3** *Suppose a is a Nash equilibrium of G. Then playing a at every period for every history is a subgame perfect equilibrium of G^∞.*

Proof: This follows from the one-shot deviation principle. If this strategy is not subgame perfect equilibrium, then there is some history h^t at which a Player i has a one-shot deviation, where he plays a_i'. But the payoff from such a deviation only differs from the the prescribed strategy by $u_i(a_i, a_{-i}) - u_i(a_i', a_{-i})$, which is positive because a is a Nash equilibrium. This completes the proof. ■

Now, denote by v_i^* the worst payoff of Player i over all Nash equilibria action profiles in G. We are now ready to state a mild version of the perfect folk theorem.

Theorem **31 (Pure Perfect Folk Theorem with Nash Reversion)** *Suppose a is any action profile such that $u_i(a) > v_i^*$ for all $i \in N$. Then, there exists a $\underline{\delta} \in (0,1)$ such that for all $\delta \in (\underline{\delta}, 1)$, there is a subgame perfect equilibrium of G^∞ where a is played in every period on equilibrium path.*

Proof: We describe a strategy that is a subgame perfect equilibrium. The strategy classifies each history into three possible states: (a) normal state, (b) i-punishment state (c) more-punishment state. If the state is normal then strategy recommends to play a.

The first period, null history is normal state. Now, we inductively define the state of any history. For every history in period t, there is a unique history in period $(t-1)$, where actions taken will lead to the history in period t. Call this the *predecessor* history. For every history, if the predecessor history is normal and a is played, then the current history becomes normal. If the predecessor state is normal and a_i is not played by a single player i but others play a_{-i}, then the state becomes i-punishment state. If the predecessor history is normal and more than one player do not play a, then the state becomes more-punishment state. If the predecessor state is i-punishment, it stays i-punishment and if the predecssor state is more-punishment, it stays more-punishment. In the i-punishment state, the strategy recommends playing the Nash equilibrium action profile of the stage game that gives Player i payoff v_i^*. In the more-punishment state, the strategy recommends playing some fixed Nash equilibrium action profile of the stage game. Denote this strategy as s. The strategy is shown in Table 31.

Predecessor state	Action profile in predecessor	Current state	Recommended action profile
Normal	a	Normal	a
Normal	(a_i', a_{-i})	i-punishment	Nash for v_i^*
Normal	$(a_S', a_{N \setminus S})$ with $\lvert S \rvert > 1$	more-punishment	Any fixed Nash
i-punishment	a'	i-punishment	Nash for v_i^*
more punishment	a'	more punishment	Any fixed Nash

Table 31: Trigger strategy for perfect Folk Theorem

In any history which is either a i-punishment state or a more-punishment state, the strategy recommends playing a Nash equilibrium. By Proposition 3, this is a Nash equilibrium of this subgame.

The only complicated history is the one which is in normal state. Fix a Player i and suppose others are following s_{-i}. If Player i follows s_i, then he gets a payoff of $u_i(a)$. By the one-shot deviation principle, we need to check deviations in one history of this subgame. Suppose Player i deviates and plays another action a'_i in some period. He gets a payoff of $u_i(a'_i, a_{-i})$ in this period, but we move to i-punishment state in the subsequent periods. As a result, he gets a payoff of v^*_i after that. Hence, his payoff from deviation is

$$(1-\delta)u_i(a'_i, a_{-i}) + \delta v^*_i.$$

Hence, to be a subgame perfect equilibrium, we will need that

$$u_i(a) \geq (1-\delta)u_i(a'_i, a_{-i}) + \delta v^*_i.$$

This can be assured if we make sure the following holds:

$$u_i(a) \geq (1-\delta) \max_{a''_i \in A_i} u_i(a''_i, a_{-i}) + \delta v^*_i.$$

Denote $\max_{a''_i \in A_i} u_i(a''_i, a_{-i}) = d_i(a_{-i})$. Then, we need to ensure that $u_i(a) \geq (1-\delta)d_i(a_{-i}) + \delta v^*_i$. This is true if

$$\delta \geq \frac{d_i(a_{-i}) - u_i(a)}{d_i(a_{-i}) - v^*_i} = \underline{\delta}.$$

Note that $d_i(a_{-i}) \geq u_i(a) > v^*_i$ ensures that $\underline{\delta} \in [0, 1)$. In other words, for $\delta \in [\underline{\delta}, 1)$, the recommended strategy is a subgame perfect equilibrium. This completes the proof. ■

20.7 Exact Versions of the Folk Theorems

Exact version of the Nash folk theorem and perfect folk theorem says that every strictly enforceable *feasible* payoff can be attained as a Nash equilibrium. The same statement is true for subgame perfect equilibrium under some additional conditions of the *feasible* payoff state.

Definition 41 *A payoff profile $v \equiv (v_1, \ldots, v_n)$ is* **feasible** *if for every action profile a in the stage game G, there exists $\lambda_a \in [0, 1]$ with $\sum_{a'} \lambda_{a'} = 1$ and for every $i \in N$*

$$v_i = \sum_{a'} \lambda_{a'} u_i(a').$$

The set of all feasible payoff profiles is denoted as $Conv(V)$. These are payoffs that can be obtained by taking convex combination of different pure action profiles. In particular, if $V = \{v : v = u(a) \; \forall \; a \in A\}$, then $Conv(V)$ is just the convex hull of V - all vectors obtained by taking convex combination of vectors in V.

One way to interpret the feasible payoffs is that these are all the payoffs that can be obtained by playing *correlated strategies*. Correlated strategies require a public randomization device. So, achieving payoffs in $Conv(V)$ requires public randomization. This requires mixed/correlated behavior strategies. A mixed behavior strategy of an agent chooses a mixed action profile at every period. Now, the minmax payoff is determined using mixed action profiles. The problem with mixed actions is that it is difficult to detect deviations. This has led to a wide literature on *monitoring* technologies in repeated games. We give some informal idea about how the folk theorems look.

	L	R
T	3,0	1,-2
B	5,4	-1,6

Table 32: A Stage game

Consider the game in Table 32. We draw its feasible payoff vector in Figure 20. The minmax values of both the players are also shown in Figure 20. It is possible that the number of extreme points of this polytope is less than the number of action profiles. Check for a game with two players and two pure actions with payoffs: $(1, 1), (2, 2), (3, 3), (4, 4)$. Here, the feasible payoff vector set is a straight line joining $(1, 1)$ and $(4, 4)$.

It is clear that any action profile of the stage game leads to a feasible payoff vector. But if the players choose their mixed actions *independently*, then it is possible that some feasible payoff vector may not be attained - this is something we have seen earlier.

For this reason to achieve any payoff in the feasible payoff vector, the players should use *public randomization device*, and everyone observes the outcome of this device, and play a strategy according to this. The public randomization device randomizes amongst the (pure strategy) payoff vectors of the stage game. Based on the payoff vector chosen by the randomizing device, everyone chooses the corresponding strategy. An analogous proof to Theorem 30 and its subgame perfect version using public randomization device can be done to establish the exact folk theorems. They will say that every strictly enforceable feasible payoff can be achieved in Nash and subgame perfect equilibrium. The subgame perfect version of these theorems use more detailed "punishment and reward" strategies and extra technical condition. We state the theorem without a proof - the theorem is due to Fudenberg

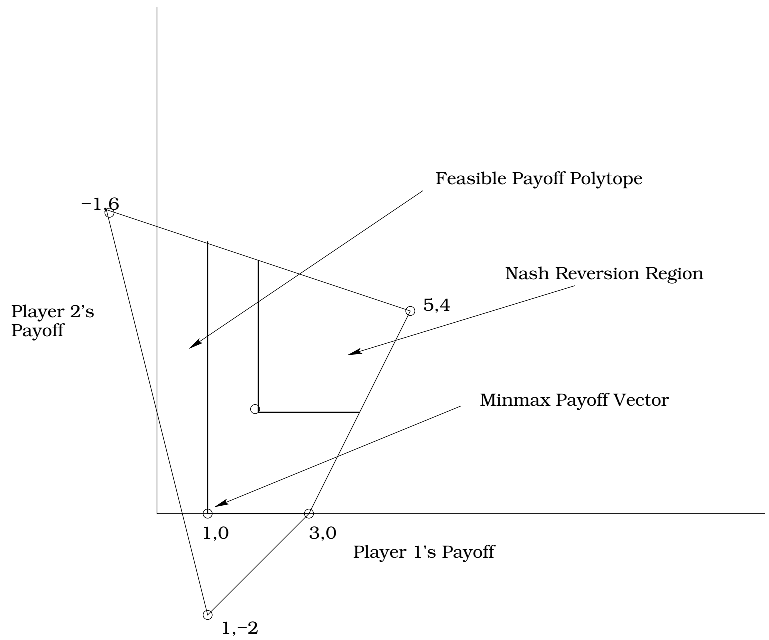

Figure 20: Feasible Payoff Vectors and Minmax Values

and Maskin.

THEOREM **32** *Suppose either $Conv(V)$ has dimension n or $n = 2$. Then, for every strictly enforceable feasible payoff vector, there is a discount factor (sufficiently close to 1) such that the infinitely repeated game generates the same payoff vector in a subgame perfect equilibrium.*

The proof of theorem uses a different type of strategy, which we illustrate below using an example. The stage game is shown in Table 33.

	L	C	R
T	2,2	2,1	0,0
M	1,2	1,1	-1,0
B	0,0	0,-1	-1,-1

Table 33: A Stage game

Notice that the minmax payoff vector is $(0, 0)$. The unique pure Nash equilibrium is (T, L). Using Theorem 31 is not so useful here. But the exact version of the folk theorem assures that $(T, L), (T, C), (M, L), (M, C)$ are possible to get in a subgame perfect equilibrium. We show below how (M, C) is possible.

Theorem **33** *Suppose $\delta \geq \frac{1}{2}$. Then, there is a subgame perfect equilibrium of the infinitely repeated game of the stage game in Table 33 such that (M, C) is played in every period in equilibrium.*

Proof: The strategy used classifies each history in each period as two states: (a) normal state (b) punishment state. A normal state recommends agents to play (M, C) and a punishment state recommends agents to play (B, R). The initial period (with null history) is a normal state.

Now, we can inductively define the state of every history. For every history in period t, there is a history in period $(t-1)$ that leads to this history, called the *predecessor*. If the predecessor is in normal state, and agents play (M, C), the current history (of period t) becomes a normal state. If the predecessor is in punishment state, and agents play (B, R), the current history becomes a normal state. Else, the current history becomes punishment state.

In other words, deviations (both in normal and punishment state) are punished for one period by staying in punishment state.

Hence, we can classify each history as a normal state or punishment state and look at deviations in each of them. Since the game is symmetric, we fix Player 1 without loss of generality and assume that Player 2 follows this strategy. If Player 1 follows the strategy, then he gets a payoff of 1. We consider two types of subgames.

Normal state. This is a subgame which starts from a normal state history. If the recommendation is followed, then player 1 gets 1. By the one-shot deviation principle, we only need to consider deviation in one period. If Player 2 plays C, then the maximum payoff of Player 1 by deviating is 2 in that period. Since this is a one period deviation, Player 1 follows the strategy from next period onwards. Since the next period will have a punishment history, he will undergo punishment and receive -1, and then normal state prevails, and he gets 1 from there onwards. The total payoff from deviation is thus computed as:

$$(1-\delta)\big(2+\delta(-1)+\delta^2+\delta^3+\ldots\big) = (1-\delta)(1-2\delta)+1.$$

Since $\delta \geq \frac{1}{2}$, this expression is less than or equal to 1. Hence, deviation is not profitable.

Punishment state. This is a subgame which starts from a punishment state history. If the recommendation is followed, then Player 1 gets punished in this period and gets (-1), which is followed by normal state that gives 1 in each period. So, the total payoff is

$$(1-\delta)\big(-1+\delta+\delta^2+\ldots\big) = 1-2(1-\delta).$$

The one-shot deviation will mean that Player 1 deviates in this period. Best deviation is to play T get 0. But this will result in a punishment in the next period and normal play from there on. Thus, the resulting payoff is

$$(1-\delta)\big(0+\delta(-1)+\delta^2+\delta^3+\ldots\big)=1-(1+2\delta)(1-\delta).$$

Note that since $\delta \geq \frac{1}{2}$, we have $1+2\delta \geq 2$. Hence, deviation is not profitable.

So, we conclude that deviation in any subgame is not profitable. This implies that the recommended strategy is a subgame perfect equilibrium. ■

The proof of the perfect Folk Theorem uses similar ideas but the punishment phase can last for more than one period (this is because the result is for general games). The number of periods the punishments last depend on the parameters of the problem.