Notes on Game Theory

Debasis Mishra

August 30, 2024

1 Introduction

Agents (individuals, firms, countries etc.) often engage in strategic interactions. Outcomes of such interactions depend on what actions are taken by all the interacting agents. For instance, outcome of an election depends on the votes of every voter. Trade between a buyer and a seller will depend on the what prices they quote. Level of emission in the environment will depend on how much emission each country commits to. How should voters vote in an election? What prices should buyers and sellers quote in trading? What level of emissions should each country commit to?

Game theory is a *toolkit* to analyze such strategic interactions. A critical element of analyzing such settings is *uncertainty*. A fundamental uncertainty is *what actions will other players take?* This is crucial in deciding how a voter will vote in an election or what price a buyer will quote in a trade. A primary objective of game theory is to formalize various ways to think about such uncertainties to take optimal actions.

Game theory can help in analyzing a variety of strategic situations. The course will be divided into various parts and each part will discuss a specific way of modeling and analyzing different strategic situations.

What we study in this course is usually referred to as *non-cooperative* game theory. We will not cover *cooperative* game theory, which models coalitional interactions and how coalitions divide gains from cooperation.

A question appropriate at this point is why not analyze situations using *general equilibrium*. There are many reasons, but two primary reasons are worth stating. General equilibrium assumes *price-taking* behavior of agents, i.e., agents choose an optimal alternative given prices. However, this theory does not consider agents being strategic. For instance, it does not consider how an equilibrium price is reached and how such prices interact with actions of agents. Another shortcoming of general equilibrium approach of modeling interaction is that it implicitly assumes that alternatives can be priced. This need not be possible in many settings: voting, matching of schools to students, organ transplant etc.

2 Games in Strategic Form

This section introduces one of the fundamental models of analyzing a strategic interactions. These games are referred to as strategic-form or normal-form games. It captures interactions in *reduced form*. For instance, think of a situation where each agent has access to a computer where it can write a detailed program. Also, assume that each agent is completely aware of all contigencies that can arise in the strategic interaction. Then, this agent can write a computer program where it can instruct the computer what it should do in each continegncy. This is the idea of a normal-form game. It is a meta-form of modeling strategic interactions.

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. When S_i is finite for each $i \in N$, we will refer to Γ as a **finite** game. The assumption is that players make binding committments about taking certain actions without observing others actions. 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. Hence, the game itself may involve players moving in sequence. But the strategic form game analysis says that players write down what they will do in *every* possible contingency of the game and they follow this as the game unfolds. For instance, consider the game of chess, suppose both the players write down what moves they will play for *every* possible position of the chess board, and as the game progresses they just follow this "plan" or "strategy". So, even though chess is a situation where players move one after the other, we are analyzing it in strategic form (or, normal form). So, the strategic form game is a "reduced form" (i.e., looking at situations from the very start) approach at analyzing strategic interactions.

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.

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

   Table 1: The Prisoner's Dilemma

   As in most models, a player's utility is just a *representation* of her preferences. Here, both the players have a ranking of outcomes. For instance, every strategy profile results in an outcome and Player 1 has the following ranking: (C, d) outcome is preferred to (D, d) outcome, which in turn is preferred to (C, c), which in turn is preferred to (D, c). The utility functions are just representation of this ranking, i.e., choosing some other numbers in Table 1 such that the same ranking is maintained will not change the analysis of these games.

2. Two shops are competing to locate themselves on a street - represented by the compact interval $[0, 1]$. Suppose consumers are uniformly located on the street. Once shops are located, the consumers go the nearest shop - with ties broken using equal probability. The utility of a shop is the *measure* of consumers it gets. Here the set of strategies are

the points in $[0,1]$ - an infinite set. If location of shop 1 is x_1 and shop 2 is x_2, then the payoff of shop 1 is

$$u_1(x_1, x_2) = x_1 + \frac{x_2 - x_1}{2} \qquad \text{if } x_1 \leq x_2$$
$$u_1(x_1, x_2) = (1 - x_1) + \frac{x_1 - x_2}{2} \qquad \text{if } x_1 > x_2$$
$$u_2(x_1, x_2) = 1 - u_1(x_1, x_2).$$

3. There are two shops on a street selling the same product. On Day 1, each shop privately observes its demand: shop $i \in \{1, 2\}$ observes a demand d_i. Given their demands, they set a price on Day 2. The payoff each shop i is $u_i(p_1, p_2)$, where p_1 and p_2 are prices of set by the two shops.

   To model this as a game, note that the two shops are the two players. What is a strategy of a shop? Naive answer will be a price. However, the price depends on the demand. The various possible demands are the contingencies of a player in the game. Hence, the strategy of a player i is a **map (function)**

$$s_i : \Re_+ \to \Re_+$$

   where $s_i(d_i)$ indicates the price of shop i when it observes demand d_i. An example of a strategy is $s_i(d_i) = \sqrt{d_i}$. The set of strategies S_i of each player consists of *all possible* maps like this.

   How should we evaluate payoff from a strategy profile in this case? Hence, if a demand vector (d_1, d_2) is observed, the payoff from a strategy profile (s_1, s_2) given by $u_i(s_1(d_1), s_2(d_2))$. But the payoff from strategy profile (s_1, s_2) is *expected payoff* across all demand vectors:

$$\sum_{(d_1, d_2)} u_i(s_1(d_1), s_2(d_2)) \text{ Prob } (\text{demand} = (d_1, d_2))$$

4. Consider the Prisoner's dilemma game being repeated twice (across two periods). In each period, the players can choose to confess or defect. The payoff of a player is the sum of payoffs across the two periods. Note that players observe the actions taken

by all the players in period 1 before taking an action in period 2. Potentially, the action they take in period 2 can *depend* on the actions observed in period 1. What is a strategy of a player? Strategy must specify two things in this repeated game:

- The action to take (confess or defect) in period 1: $s_i^1 \in \{\text{Confess}, \text{Defect}\}$
- The action to take (confess or defect) in period 2 as a function of observed actions in period 2:

$$s_i^2 : \{\text{Confess}, \text{Defect}\} \times \{\text{Confess}, \text{Defect}\} \to \{\text{Confess}, \text{Defect}\}$$

So, the strategy consists of these pairs (s_i^1, s_i^2) and the set of strategies S_i consists of all such pairs.

An example of such a strategy is to play confess in first period. If confess is played by both in first period, play defect in period 2; else play confess in period 2. So, $s_i^1 = \text{Confess}$ and $s_i^2(C, c) = \text{Defect}$; else $s_i^2(\cdot, \cdot) = \text{Confess}$. The payoff from this particular strategy for a player is $u_i(s_1, s_2) = 1 + 2 = 3$: this is because if players play this strategy, first period outcome is (C, c) and second period outcome is (D, d). Notice that the strategy specified actions to be taken in all contingencies (i.e., it specified what action to take if period 1 actions were $(C, d), (D, c), (D, d)$) but the only contingency required to compute payoff is the one that arose by play of the strategies. As we will see later, even then it is important to specify the actions to be taken in all contigencies.

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 we showed in these examples, a strategic form game *does not* mean that players move "simultaneously" in the game. Any strategic interaction can be viewed as a strategic form game. In a strategic form game, players make binding committments about what they will do under each contingency, and that forms their strategy. Hence, many strategic interactions can be reduced to such strategic form by specifying the strategies appropriately. As we go along in the course, we will see that strategies have different meaning and definitions in different types of interactions agents can have. But, in all such cases, the common thread that will run is: *a strategy will describe what an agent must do in all possible contingencies.*

3 Beliefs of Players

One of the objectives 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 that 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:[1]

- 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.

Note. From now on, unless stated otherwise, we will assume S_i for all i to be finite sets. Many results, with the help of extra notations and mathematics, extend to the case where strategy sets are not finite.

Notation. Throughout, we will refer to any n-tuple of strategies, one for each player, as a **strategy profile**. So, notationally, we will write this as $s \equiv (s_1, \ldots, s_n)$ or (s_i, s_{-i}). The s_{-i} notation is used to denote a strategy profile of $(n-1)$ players that does not include Player i. The notation S_{-i} is the cartesian product of sets of strategies of all the players excluding Player i, i.e., $S_{-i} = \prod_{j \neq i} S_j$. As an example, suppose $N = \{1, 2, 3\}$, $S_1 = \{a, b\}, S_2 = \{x, y, z\}, S_3 = \{p, q, r\}$. Then

$$S_{-1} = S_2 \times S_3 = \{(x,p),(x,q),(x,r),(y,p),(y,q),(y,r),(z,p),(z,q),(z,r)\}$$

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}. In the previous example, μ_1 will assign a probability distribution over the set $S_2 \times S_3$. For instance, $\mu_1(x,p) = \frac{1}{3}, \mu_1(y,r) = \mu_1(z,p) = \mu_1(z,q) = \frac{2}{9}$ is a valid belief.

[1] The use of expected utility to evaluate uncertainty has a strong foundation. You will learn this in microeconomics.

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. To see this consider Table 2. Suppose Player 1 believes Player 2 will play c. Then, her optimal strategy is to play C. But suppose Player 2 believes Player 1 will play D. Then, her optimal strategy is d. Here, the beliefs of both the players are *not consistent* with the optimal actions of other players.

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

Table 2: Belief inconsistency

Game theory aims to provide prediction of play of players in games. This is referred to as the solution concept. Broadly, there are two kinds of solution concepts in game theory: (a) solution concepts that work independent of beliefs and (b) solution concepts that assume correct and consistent beliefs. The former is sometimes referred to as a *non-equilibrium* solution concept, while the latter is referred to as an *equilibrium* solution concept.

4 Belief-free approach: 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. The basic question that drives this idea is: *how do we compare a pair of strategies?*

4.1 Strict domination

We start with the strongest possible way to compare strategies.

Definition 1 *A strategy* s_i *of Player* i **strictly dominates** *a strategy* t_i **with respect to** *belief* $\mu_i \in \Delta(S_{-i})$ *if*

$$\sum_{s_{-i}} u_i(s_i, s_{-i})\mu_i(s_{-i}) > \sum_{s_{-i}} u_i(t_i, s_{-i})\mu_i(s_{-i}).$$

In this case, we write $s_i \succ_{\mu_i} t_i$*.*

The relation $\succ_{\mu_i}$ crucially depends on belief μ_i. To see this, consider the game in Table 2. Take two beliefs: $\mu_1(c) = 1, \mu_1(d) = 0$ and $\mu_1'(c) = 0, \mu_1'(d) = 1$. We see that for Player 1: $C \succ_{\mu_1} D$ but $D \succ_{\mu_1'} C$. The following definition formalizes a way to compare strategies which is belief-free.

Definition 2 *A strategy* s_i *of Player* i **strictly dominates** *strategy* t_i *if*

$$s_i \succ_{\mu_i} t_i \qquad \forall\ \mu_i \in \Delta S_{-i}.$$

Thus, independent of what Player i belives other players will play, she prefers s_i to t_i. So, she can unambiguously choose s_i over t_i. This definition of dominance is equivalent to the following.

Lemma 1 *Suppose* $\Gamma = (N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N})$ *is a finite game. A strategy* s_i *for Player* i *strictly dominates strategy* t_i *if and only if*

$$u_i(s_i, s_{-i}) > u_i(t_i, s_{-i}) \qquad \forall\ s_{-i} \in S_{-i}. \tag{1}$$

Proof: Suppose Inequalities (1) hold for s_i and t_i. Fix a belief $\mu_i \in \Delta S_{-i}$. Now, using Inequality (1), we immediately get

$$\sum_{s_{-i}} u_i(s_i, s_{-i})\mu_i(s_{-i}) > \sum_{s_{-i}} u_i(t_i, s_{-i})\mu_i(s_{-i})$$

For the other direction, suppose s_i strictly dominates t_i. Now, choose some s_{-i} and consider the belief that $\mu_i(s_{-i}) = 1$ and $\mu_i(s'_{-i}) = 0$ for all $s'_{-i} \neq s_{-i}$. Then, it follows that

$$u_i(s_i, s_{-i}) = \sum_{t_{-i}} u_i(s_i, t_{-i})\mu_i(t_{-i}) > \sum_{t_{-i}} u_i(t_i, t_{-i})\mu_i(t_{-i}) = u_i(t_i, s_{-i})$$

■

Lemma 1 gives a straightforward way to check if strategy s_i strictly dominates t_i or not. Unfortunately, in many games there may be no strategy that strictly dominates another strategy. An example is the game in Table 2. However, in the Prisoner's Dilemma game in Table 1, C strictly dominates D for Player 1 and c strictly dominates d for Player 2. If there is a game in which there is a single strategy which strictly dominates every other strategy for a player, then that strategy is called a strictly dominant strategy.

DEFINITION **3** *A strategy $s_i \in S_i$ for Player i is* **strictly dominant** *if s_i strictly dominates t_i for every other strategy $t_i \in S_i \setminus \{s_i\}$.*

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 3, there is no strictly dominant strategy for either of the players. Notice that for Player 1, T strictly dominates M but T does not strictly dominate B.

	L	C	R
T	$(2,2)$	$(6,1)$	$(10,10)$
M	$(1,3)$	$(5,5)$	$(9,2)$
B	$(3,3)$	$(4,2)$	$(8,8)$

Table 3: Domination

4.2 Weak domination

Since it is difficult to find games where we can compare strategies based on strict domination, a natural weakening is worth considering.

Definition 4 *A strategy* s_i **weakly dominates** *strategy* t_i *if for every* $s_{-i} \in S_{-i}$*,*

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

with strict inequality holding for some s_{-i}*.*

A strategy s_i *is* **weakly dominant** *if it weakly dominates every other strategy* $t_i \in S_i \setminus \{s_i\}$*.*

Lemma 2 *A strategy* s_i *weakly dominates strategy* t_i *if and only if for every belief* $\mu_i \in \Delta S_{-i}$

$$\sum_{s_{-i}} u_i(s_i, s_{-i})\mu_i(s_{-i}) \geq \sum_{s_{-i}} u_i(t_i, s_{-i})\mu_i(s_{-i}) \tag{2}$$

with strict inequality holding for some μ_i*.*

Proof: Suppose s_i weakly dominates t_i (in the sense of Definition 4). Then, for every s_{-i} we have

$$u_i(s_i, s_{-i}) \geq u_i(t_i, s_{-i}), \tag{3}$$

with strict inequality for some s^*_{-i}. Fix any arbitrary belief mu_i. By multiplying by $\mu_i(s_{-i})$ each inequality (3) we get

$$\sum_{s_{-i}} u_i(s_i, s_{-i})\mu_i(s_{-i}) \geq \sum_{s_{-i}} u_i(t_i, s_{-i})\mu_i(s_{-i}) \tag{4}$$

If we consider the belief μ^*_i, where $\mu^*_i(s^*_{-i}) = 1$, inequality (4) becomes $u_i(s_i, s^*_{-i})$ on LHS and $u_i(t_i, s^*_{-i})$ on RHS. This is strict by our assumption. For the other direction, suppose inequality (2) holds for every belief μ_i and holds strictly for some belief μ^*_i. Fix any arbitrary s_{-i} and consider the belief $\mu_i(s_{-i}) = 1$. Then, (2) reduces to $u_i(s_i, s_{-i}) \geq u_i(t_i, s_{-i})$. Hence, the weak inequality in Definition (4) holds for all s_{-i}. Now, for some belief μ^*_i, we know that

$$\sum_{s_{-i}} u_i(s_i, s_{-i})\mu^*_i(s_{-i}) > \sum_{s_{-i}} u_i(t_i, s_{-i})\mu^*_i(s_{-i}) \tag{5}$$

Assume for contradiction $u_i(s_i, s_{-i}) \leq u_i(t_i, s_{-i})$ for all s_{-i}. Then, multiplying by $\mu_i^*(s_{-i})$ for each s_{-i}, we get

$$\sum_{s_{-i}} u_i(s_i, s_{-i})\mu_i^*(s_{-i}) \leq \sum_{s_{-i}} u_i(t_i, s_{-i})\mu_i^*(s_{-i}) \tag{6}$$

Inequalities (6) and (5) contradict each other. Hence, there must be some s_{-i}^* such that $u_i(s_i, s_{-i}^*) > u_i(t_i, s_{-i}^*)$. Hence, the inequality in Definition (4) is strict for some s_{-i}^*. This shows that s_i weakly dominates t_i in the sense of Definition (4). ■

While strictly dominant strategies are difficult to find, it is not difficult to find a weakly dominant strategy in some games of interest. We give two examples below.

The idea of strict or weak domination has another special appeal. When we introduced the definition of a strategic form game, we assumed that all the players know *everything* specified in the definition of the game. For instance, every Player i knows the payoffs obtained by every Player j. This will be crucial in doing any analysis we do later. However, for comparing strategies using strict or weak domination, we *only require players know their own payoffs*, but they need not know the payoff of other players or do not need to form any beliefs about what other players will play. For instance, consider the game in Table 3 where we argued that strategy T strictly dominates strategy M for Player 1. Now, consider the same game in Table 4, where we have removed all payoffs of Player 2. We can again conclude that strategy T strictly dominates strategy M: vector $(2, 6, 10)$ strictly dominates $(1, 5, 9)$ component-wise.

	L	C	R
T	$(2, \cdot)$	$(6, \cdot)$	$(10, \cdot)$
M	$(1, \cdot)$	$(5, \cdot)$	$(9, \cdot)$
B	$(3, \cdot)$	$(4, \cdot)$	$(8, \cdot)$

Table 4: Domination

This observation is true for weak domination too. Hence, to be able to compare any pair of strategies using (strict or weak) domination, we do not require players to have knowledge of payoffs of other players or have any beliefs about other players' strategies. This motivates our following two examples, where we are in a setting where players do not know each others' payoffs and we can still apply the notion of weak domination.

4.2.1 An Auction Example

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 with probability q_i and pays a price p_i, then his utility is $q_i(v_i - p_i)$. If a buyer does not win the object, she does not pay anything and her payoff is zero.

Note that buyers need not know each others values in this model, which is natural in this case. However, we can still apply the idea of domination in this case.

Lemma **3** *In the Vickrey auction, it is a weakly dominant strategy for every buyer to bid his value.*

Proof: Fix a buyer i and suppose each of the other bidder $j \neq i$ bids b_j – so, we have fixed an arbitrary strategy profile of other bidders $\{b_j\}_{j \neq i}$. We will argue whatever this strategy profile may be bidder i weakly prefers v_i to every other strategy.

Before proceeding with the proof, consider Figure 1. It plots the payoff of a buyer i along the Y-axis and bid of the buyer i along the X-axis. The payoff of the buyer i is zero if it bids below $\max_{j \neq i} b_j$. Otherwise (if he bids above $\max_{j \neq i} b_j$),

- if the value of the buyer i is above $\max_{j \neq i} b_j$, then its payoff of the buyer is given by the blue line (line above Y-axis),
- if the value of the buyer i is below $\max_{j \neq i} b_j$, then its payoff of the buyer is given by the red line (line below Y-axis).

Hence, each buyer i, independent of its value can partition its strategies into two sets: (i) below $\max_{j \neq i} b_j$ and (ii) above $\max_{j \neq i} b_j$. It gets the same payoff by bidding anything in each of these sets. A buyer whose value is above $\max_{j \neq i} b_j$ prefers the blue part to the orange part in Figure 1, but a buyer whose value is below $\max_{j \neq i} b_j$ prefers the orange part to the blue part in Figure 1.

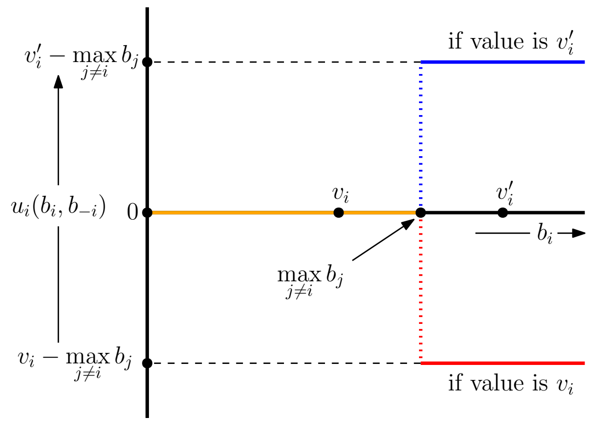

Figure 1: Weakly dominant strategy in Vickrey auction

Figure 1 gives an idea on why bidding value maximizes payoff of any buyer. Below, we formally show that it is indeed a weakly dominant strategy. Suppose buyer i has value v_i. We consider two cases.

CASE 1. $v_i > \max_{j \neq i} b_j$. In this case, the payoff of buyer i from bidding v_i is $v_i - \max_{j \neq i} b_j > 0$. As long as he bids more than $\max_{j \neq i} b_j$, buyer i's payoff remains the same: she still wins the object and pays the same. By bidding strictly less than $\max_{j \neq i} b_j$ she does not win the object and gets a payoff of zero. By bidding equal to $\max_{j \neq i} b_j$, she gets the object but with some probability $q \leq 1$ and pays $\max_{j \neq i} b_j$. Hence, her payoff is $q(v_i - \max_{j \neq i} b_j)$, which is not more than what she was getting by bidding v_i.

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. This is because either she is not getting the object (in which case his payoff is zero) or she is sharing the object in which case she is paying $\max_{j \neq i} b_j = v_i$. This is the case for all bids strictly less than $\max_{j \neq i} b_j$. If she bids greater than or equal to $\max_{j \neq i} b_j$, she wins (with some probability) but pays $\max_{j \neq i} b_j \geq v_i$. Hence, her payoff is non-positive. Hence, bidding v_i is at least as good as bidding anything else.

Finally, for weakly dominant strategy, we need to show that v_i is better than any other strategy for *some* bid vector of other players. For this, fix v_i and some strategy $b_i \neq v_i$. As we saw from the two cases, if $b_i > v_i$, then when $v_i < \max_{j \neq i} b_i < b_i$, it is strictly better for buyer i to bid v_i. Similarly, if $b_i < v_i$, then when $b_i < \max_{j \neq i} b_i < v_i$, it is strictly better for buyer i to bid v_i. ■

4.2.2 A voting example

We now consider an example from voting. Besides highlighting weakly dominant strategies, it also emphasizes that games can be *ordinal*, i.e., devoid of any utility representation.

In the voting problem, there is a finite set of candidates A. The candidates are ordered (ranked) by some parameter exogenously. For instance, candidates are ordered according to their idealogical position. We denote this ordering over A as $\succ$. Agents in N are voters. Voters have preference (strict ranking) over candidates in A. In particular, we will assume that each voter $i \in N$ has a strict ranking P_i over A *and* P_i satisfies **single-peakedness**.

Informally, a single peaked preference says that as we go away from the *peak* (top ranked candidate) using the exogenous order $\succ$, we prefer candidates less. Before formally defining the preference, consider Figure 2. The exogenous order $\succ$ is: $a_7 \succ a_6 \succ a_5 \succ a_4 \succ a_3 \succ a_2 \succ a_1$. The peak of this preference is a_3. According to single-peakedness, a_4 is preferred to a_6 because a_4 is closer to a_3 then a_6. A pictorial description of a single peaked preference is shown in Figure 2. In particular, a single peaked preference cannot have the following: a_3 at the top but a_5 is better than a_4. Note that we do not restrict preferences across the peak in any way. More formally, single-peakedness is defined as follows. Suppose $P_i(1)$ is the peak

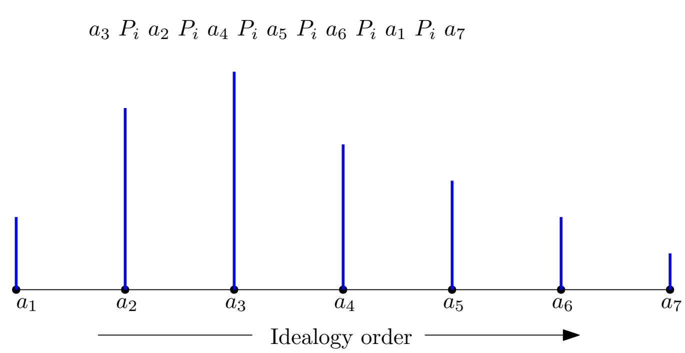

Figure 2: Single-peaked preference

of agent i when his preference is P_i. Then, P_i is single-peaked with respect to $\succ$ if for all $a, b \in A$ if $a \succ b \succ P_i(1)$ or $P_i(1) \succ b \succ a$, then $b\ P_i\ a$ (i.e., b is preferred to a).

Voting game. The game assumes that each agent has single-peaked preference with respect $\succ$. Each agent can submit a candidate - the strategy set $S_i = A$ for each $i \in N$. Denote the submitted candidate of agent/voter i as $s_i \in A$. Based on the submitted profile of candidates $(s_1, \ldots, s_n)$, the winner is determined by looking at the **median** of the submitted candidates with respect to $\succ$, i.e.,

$$W(s_1, \ldots, s_n) := \text{median}(s_1, \ldots, s_n),$$

where in case of n being even, we choose the $\frac{n}{2}$-th candidate in $(s_1, \ldots, s_n)$ according to $\succ$. For instance, in Figure 2, if there are three agents and they submit $(s_1 = a_3, s_2 = a_1, s_3 = a_5)$, then the median is agent 1's submitted candidate. Here, we do not need a utility representation for agents. Agents can compare outcomes in the game using their preferences. In particular, if agent's preferences are $(P_1, \ldots, P_n)$, then a strategy profile $(s_1, \ldots, s_n)$ is a **weakly dominant strategy** if for every $i \in N$ and every $\hat{s}_{-i}$,

$$W(s_i, \hat{s}_{-i})\ P_i\ W(s'_i, \hat{s}_{-i}) \text{ or } W(s_i, \hat{s}_{-i}) \ = \ W(s'_i, \hat{s}_{-i})\ \forall\ s'_i \in A \setminus \{s_i\},$$

with the strict relation holding for some $\hat{s}_{-i}$. We show that every agent has a weakly dominant strategy.

LEMMA 4 *In the voting game each agent $i \in N$ submitting $s_i = P_i(1)$ is a weakly dominant strategy.*

Proof: Before doing the proof, consider an example with five agents. Figure 3 plots the set of alternatives on X-axis and also on Y-axis. X-axis shows the vote of the fifth agent and Y-axis shows the outcome of the voting game (median). The votes of the other four agents are fixed and shown on the X-axis. As the fifth agent moves its vote from 0 to the right, the median does not change till it crosses the second agent, but then the median moves with the vote of the fifth agent. When the vote of the fifth agent crosses the vote of the third agent, the median becomes the vote of the third agent and stays that way further right. Hence, if agent 5 has her true peak between second and third agents' votes, she gets her own peak and cannot do better. If agent 5 has her peak to the left of agent 2's vote, she can only shift the median to the right, which she prefers less due to single-peakedness. Similarly, if agent 5 has her peak to the right of agent 3' vote, she can only shift the median to the left, which she prefers less due to single peakedness.

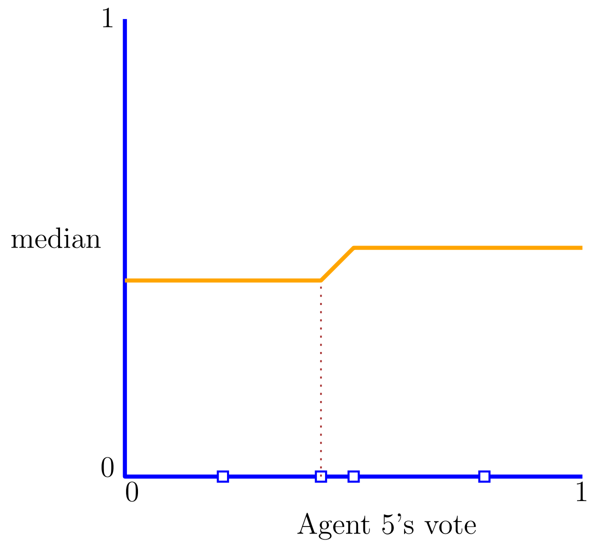

Figure 3: Median as a function of vote

Suppose other agents submitted $\hat{s}_{-i}$. If agent i submits $s_i = P_i(1)$, let the candidate $a = W(s_i, \hat{s}_{-i})$ be chosen. So, a is the median of $(s_i, \hat{s}_{-i})$. If $a = P_i(1)$, then agent i cannot do better by submitting some other candidate. Suppose a is to the "left" of $P_i(1)$: $P_i(1) \succ a$. As long as i submits s'_i such that $s'_i \succ a$ or $s'_i = a$, the median of $(s'_i, \hat{s}_{-i})$ remains a. If i submits s'_i such that $a \succ s'_i$, then the median of $(s'_i, \hat{s}_{-i})$ can only move to the "left" of a, i.e., $a \succ W(s'_i, \hat{s}_{-i})$. But $P_i(1) \succ a \succ W(s'_i, \hat{s}_{-i})$ implies that $a \; P_i \; W(s'_i, \hat{s}_{-i})$. Hence, agent i does not prefer this outcome. So, in both cases, agent i cannot do better.

A similar proof can be done if $a \succ P_i(1)$.

What remains to be seen is that for some report $s'_i \neq s_i$, there is some report s_{-i} such that agent i strictly prefers s_i to s'_i. We illustrate the idea for odd n – a similar proof works for even also. Of the $(n-1)$ agents ($(n-1)$ is even), consider a report such that $\frac{n-1}{2}$ agents have peak at s'_i and the remaining $\frac{n-1}{2}$ agents have peak at s_i. Then, by reporting s_i, the median is s_i. But reporting s'_i, the median shifts to s'_i, which the agent strictly prefers less. ■

5 Belief-free approach: iterated deletion of strategies

In this section, we discuss two extensively discussed belief-free approaches of narrowing down the strategies of players. Both the approaches are closely related but require more than rationality of players.

5.1 Dominated strategies

Consider the game in Table 5. Irrespective of the strategy played by Player 2, Player 1

	L	C	R
T	$(4, 2)$	$(6, 1)$	$(10, 10)$
M	$(1, 3)$	$(5, 5)$	$(9, 2)$
B	$(3, 3)$	$(7, 2)$	$(8, 8)$

Table 5: Dominated strategies

always gets less payoff in M than in T, i.e., T strictly dominates M. In such a case, we will say that Strategy M is strictly dominated.

DEFINITION **5** *A strategy $s_i \in S_i$ for Player i is* **strictly dominated** *if there exists $s'_i \in S_i$ such that s'_i strictly dominates s_i, i.e., for every $s_{-i} \in S_{-i}$, we have*

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

An implication of rationality is that 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. Note that belief of a player about others' strategies influences his choice of optimal strategy.

To see this, consider the example in Table 5. Strategy M is strictly dominated by Strategy T for Player 1. Hence, if Player 1 is rational, then he will not play M. If Player 2 does not know that Player 1 is rational, then he cannot eliminate M from the support of his belief of Player 1's strategies. Suppose **Player 2 knows that Player 1 is rational**. Then, he can conclude that Player 1 will not play M ever. As a result, his belief on what Player 1 can play must put probability zero on M. In that case, his Strategy C is strictly dominated by Strategy L. So, he will not play C. 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 B because it is now strictly dominated by T. Continuing in this manner, we will get that Player 2 does not play L. Hence, the only strategy profile surviving such elimination is (T, R).

The process we just described is called *iterated elimination of strictly dominated strategies*. It requires more than rationality. We do not provide a formal treatment of this topic. Loosely, a *fact* is **common knowledge** among players in a game if for any finite chain of players $(i_1, \ldots, i_k)$ the following holds: Player i_1 knows that Player i_2 knows that Player i_3 knows that ... Player i_k knows the fact. Iterated elimination of strictly dominated strategies, possibly for infinite rounds, require the following assumption. **Common Knowledge of Rationality (CKR)**: The fact that all players are rational is common knowledge. If we only require interated elimination of strictly dominated strategies for some finite number of rounds, then we will not require *full* common knowledge assumption – we will require the above chain of assumptions to hold till a finite "depth".

Let us consider another example in Table 6. 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 6: Domination

We now formally define the procedure of iterated elimination.

Definition 6 **(Iterated elimination of strictly dominated strategies (IESDS))** *Start with a strategic form game* $\Gamma = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$. *Let* $\Gamma^0 := \Gamma$, $S_i^0 := S_i$ *for each* $i \in N$. *At any iteration* $t > 0$,

- *for every player i, let $D_i^{t-1} \subseteq S_i^{t-1}$ be the set of strictly dominated strategies of Player i in Γ^{t-1};*

- *set $S_i^t := S_i^{t-1} \setminus D_i^{t-1}$ for all $i \in N$;*

- *set $\Gamma^t := (N, \{S_i^t\}_{i \in N}, \{u_i^t\}_{i \in N})$, where u_i^t is the restriction of u_i to strategy profiles in $S_1^t \times \ldots \times S_n^t$.*

- *If $\Gamma^t = \Gamma^{t-1}$, then* STOP*; else set $t := t + 1$ and repeat.*

If the IESDS process converges, let Γ^* be the game to which it converges, and the strategies in Γ^* for each player are called the strategies that **survive the IESDS process**.

If the original game is finite, the IESDS process terminates in finite iterations. Otherwise, this may lead to an infinite sequence of games. The convergence of these games is a technical matter and will not be discussed. But in most infinite games that we will study, we will see that the IESDS process converges.

In some games, the IESDS process leads to a unique outcome of the game. In those cases, we call it a **solution** of the game. However, absence of strictly dominated strategies will imply that no strategies can be eliminated. In such a case, the IESDS process only gives the original game: $\Gamma^* = \Gamma$. In some games, Γ^* will contain strictly smaller set of strategies than Γ but not necessarily unique. In that case, IESDS just says that players may play any strategy in Γ^*. In this sense, IESDS is not really a *unique prediction* of the game but a *set-valued prediction.*

We give below an example which illustrates the idea of IESDS process in an infinite game.

5.2 Example: Bertrand Game

There are two firms $\{1, 2\}$ producing the same good. Both the firms choose prices in $[0, 1]$. Depending on the prices chosen, p_1 and p_2, demand function for each firm $i \in \{1, 2\}$ is given by

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

where $j \neq i$ is the other firm. Suppose the marginal costs for both the firms are zero, then the utility function of firm i is

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

Given any p_j, the best strategy of firm i is the unique maximum point of the above strictly concave function, which can be obtained by taking the first order condition:

$$1 - 4p_i + p_j = 0. \tag{7}$$

or $p_i = \frac{1}{4}(1 + p_j) \in [\frac{1}{4}, \frac{1}{2}]$. In particular, given p_j, the utility function of Player i is strictly increasing in $[0, \frac{1}{4}(1 + p_j))$ and strictly decreasing in $(\frac{1}{4}(1 + p_j), 1]$. Also, notice that this critical point is increasing in p_j. So, if Player i does not know which p_j is chosen by Player j, then he needs to be sure about strategies which does worse over a strategy *irrespective* of the value of p_j chosen. For $p_j = 0$, the utility function is increasing in $[0, \frac{1}{4})$. Hence, for all $p_j \in [0, 1]$, strategy $p_i = \frac{1}{4}$ strictly dominates any strategy in $[0, \frac{1}{4})$. Similarly, the utility function is decreasing in the interval $(\frac{1}{2}, 1]$ irrespective of the value of $p_j \in [0, 1]$. So, the first iteration of elimination gives strategies, $[\frac{1}{4}, \frac{1}{2}]$ to both the players.

Now, when the lowest value of p_j is $\frac{1}{4}$, strategies $[\frac{1}{4}, \frac{5}{16})$ will be strictly eliminated as the utility function of i for all possible strategies of j is increasing in this region. Also, strategies $(\frac{3}{8}, \frac{1}{2}]$ are eliminated since utility function of i is decreasing in this region for all possible stragies of j. Hence, second iteration of elimination suggests, strategies $[\frac{1}{4}, \frac{5}{16})$ and $(\frac{3}{8}, \frac{1}{2}]$ are strictly dominated. Hence, we are left with the strategies in $[\frac{5}{16}, \frac{3}{8}]$.

Continuing in this manner gives us the following sequence of strategy sets:

$$[\frac{1}{4}, \frac{1}{2}], [\frac{5}{16}, \frac{3}{8}], [\frac{21}{64}, \frac{11}{32}], [\frac{85}{256}, \frac{43}{128}], \ldots.$$

So, the lower point of the intervals are

$$\frac{1}{4}, \frac{1}{4} + \frac{1}{4^2}, \frac{1}{4} + \frac{1}{4^2} + \frac{1}{4^3}, \ldots,$$

This sequence converges to $\frac{1}{3}$. The upper point of the intervals are

$$\frac{1}{4}+\frac{1}{4}, \frac{1}{4^2}+\frac{1}{4}+\frac{1}{4^2}, \frac{1}{4^3}+\frac{1}{4}+\frac{1}{4^2}+\frac{1}{4^3}, \ldots.$$

This sequence converges to $\frac{1}{3}$ too. Hence, iterated elimination of strictly dominated strategies lead to a unique outcome in this game: $(\frac{1}{3}, \frac{1}{3})$.

While this game produced a unique outcome due to IESDS, this need not hold in many games. For instance, consider the Vickrey auction game described above. Suppose there are three buyers with values $20, 10, 5$. What are the strictly dominated strategies in this game?

5.3 Order independence in IESDS

Instead of eliminating every strictly dominated strategy in a round, we may think of eliminating only a subset of strictly dominated strategies in a round. Does the set of strategies that survive change with this modification?

Formally, we consider a generalized version of IESDS.

Definition **7 (Generalized Iterated elimination of strictly dominated strategies (GIESDS))** *Start with a strategic form game $\Gamma = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$. Let $\Gamma^0 := \Gamma$, $S_i^0 := S_i$ for each $i \in N$. At any iteration $t > 0$,*

- *for every player i, let $D_i^{t-1} \subseteq S_i^{t-1}$ be the set of strictly dominated strategies of Player i in Γ^{t-1};*
- *set $S_i^t := S_i^{t-1} \setminus X_i^{t-1}$ for all $i \in N$, where $X_i^{t-1} \subseteq D_i^{t-1}$;*
- *set $\Gamma^t := (N, \{S_i^t\}_{i \in N}, \{u_i^t\}_{i \in N})$, where u_i^t is the restriction of u_i to strategy profiles in $S_1^t \times \ldots \times S_n^t$.*
- *If $\Gamma^t = \Gamma^{t-1}$, then* STOP*; else set $t := t+1$ and repeat.*

If the GIESDS process converges, let Γ^* be the game to which it converges, and the strategies in Γ^* for each player are called the strategies that **survive the GIESDS process***.*

In the above definition, by choosing X_i^t of each player in each round t differently, we can generate different sequence of games. For instance, we can order the players and ask a unique player to eliminate in a round. If there are two players, players can take turn to eliminate.

Theorem 1 *Suppose the game is finite. Then, GIESDS and IESDS terminates at the same game.*

Remark on the proof. The proof works for finite game through induction. I only give an informal sketch. The basic idea of the proof is the following: if a strategy s_i is strictly dominated by t_i in some round of IESDS, it is also strictly dominated (possibly by another strategy) in a new game, where we remove some strategies (perhaps t_i itself) of each player. Later, we will establish a more rigorous proof of this theorem.

Proof: Consider player i. Consider the strictly dominated strategies D_i^0 in the original game (i.e., in round 1). We argue that each strategy in D_i^0 is eliminated in GIESDS. Pick any $s_i \in D_i^0$. If s_i is strictly dominated by t_i, it continues to be dominated by t_i in any subgame. Further, if t_i is eliminated by a strategy, it must be strictly dominated by another strategy t_i', which also dominates s_i. Hence, s_i remains strictly dominated throughout the GIESDS, and it will be eliminated.

Now, we can do an induction argument. Suppose each strategy of each Player till round $k-1$ of IESDS is also eliminated in GIESDS. Pick a player i and a strategy s_i which is eliminated in round k of IESDS since it was dominated by t_i. By our induction argument, there exists a round in GIESDS, where all the strategies eliminated till round $(k-1)$ of IESDS do not exist. At such a round, s_i is either strictly dominated by t_i or t_i is already eliminated. If t_i was eliminated, it was dominated by another strategy which also dominates s_i in this round. In summary, s_i is still strictly dominated, and it will be eliminated in some round of GIESDS. This establishes the induction step. By the finiteness of the game, we are done. This establishes that GIESDS eliminates all the strategies that IESDS eliminates.

Can GIESDS eliminate more strategies? Consider the first round in GIESDS where a strategy s_i of Player i is eliminated which was not eliminated in IESDS. In that round, all the strategies of all the players present are also there in *some* round of IESDS. So if s_i is dominated by t_i in that round, it must have been dominated by t_i even in some round of IESDS. This is a contradiction. ■

5.4 Never best responses

We now introduce an alternate notion of deletion of strategies. This is based on the idea of *(never) best responses*. We remind the notation that $\mathcal{U}_i(s_i, \mu_i)$ denotes the utility of Player i from strategy s_i given her belief $\mu_i \in \Delta S_{-i}$ over other players' strategies, i.e., $\mathcal{U}_i(s_i, \mu_i) := \sum_{s_{-i}} u_i(s_i, s_{-i})\mu_i(s_{-i})$. Using this, we define the notion of a best response.

Definition 8 *A strategy $s_i \in S_i$ of Player i is a* **best response with respect to belief** $\mu_i \in \Delta S_{-i}$ *if*

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

Hence, if s_i is a best response with respect to belief μ_i, then it solves the optimization problem $\max_{s'_i \in S_i} \sum_{s_{-i}} u_i(s'_i, s_{-i})\mu_i(s_{-i})$. If S_i is finite, surely a maximum exists. If S_i is not finite, we can put sufficient conditions on S_i and u_i to ensure a best response strategy exists. If a strategy s_i is not a best response with respect to *any* belief, then we call it a never best response.

Definition 9 *A strategy $s_i \in S_i$ of Player i is a* **never best response** *if there exists no belief $\mu_i \in \Delta S_{-i}$ such that s_i is a best response with respect to μ_i. In other words, s_i is a never best response if for every belief $\mu_i \in \Delta S_{-i}$, there is a strategy t_i such that $\mathcal{U}_i(t_i, \mu_i) > \mathcal{U}_i(s_i, \mu_i)$.*

To verify if a strategy is a never best response, we need to consider all possible beliefs. The difference between never best response and strictly dominated strategy is subtle. It is merely the placement of *quantifiers* in the definition – in strictly dominated strategy, we want a strategy t_i which dominates s_i for all beliefs; in never best response, we want a strategy for each belief (so, this strategy may change with beliefs) which dominates s_i.

Consider the two player game in Table 7 - the table only shows payoff of Player 1. Strategy C is **not** strictly dominated. However, if we only consider *degenerate* beliefs of Player 1 (i.e., those for which she puts entire probability on either a or b), then C is not a best response. However, let Player 1 have belief where she puts $\frac{1}{2}$ probability on a and $\frac{1}{2}$ probability on b. Strategy C is a best response with respect to this belief. This shows that a strategy may not be a best response with respect to degenerate beliefs but it may be a best response with respect to beliefs which put positive probability on at least two strategies of other players.

	a	b
A	$(1, \cdot)$	$(0, \cdot)$
B	$(0, \cdot)$	$(1, \cdot)$
C	$(0.6, \cdot)$	$(0.6, \cdot)$

Table 7: Best response example

If s_i is a strictly dominated strategy, then it is a never best response.

Claim **1** *Every strictly dominated strategy is a never best response.*

Proof: Suppose s_i is a strictly dominated strategy for Player i, and suppose t_i strictly dominates s_i. Then, for every belief μ_i, we have $\mathcal{U}_i(t_i, \mu_i) > \mathcal{U}_i(s_i, \mu_i)$. Hence, s_i is a never best response. ■

However, a never best response strategy is not necessarily strictly dominated. For this, consider the game in Table 8. For Player 1, neither T nor M strictly dominate B. However, we argue that B is a never best respose. Consider an arbitrary belief μ_1 of Player 1 which puts p probability on L and $(1-p)$ on M. Then, her payoffs from this belief is

$$\mathcal{U}_1(T, \mu_1) = p$$
$$\mathcal{U}_1(M, \mu_1) = 1 - p$$
$$\mathcal{U}_1(B, \mu_1) = 0.4$$

If $p \geq 0.5$, then $\mathcal{U}_1(T, \mu_1) > 0.4$. If $p \leq 0.5$, then $\mathcal{U}_1(M, \mu_1) > 0.4$. Hence, for any arbitrary belief μ_1, B cannot be a best response. Hence, B is a never best response but not strictly dominated.

	L	M
T	$(1, \cdot)$	$(0, \cdot)$
M	$(0, \cdot)$	$(1, \cdot)$
B	$(0.4, \cdot)$	$(0.4, \cdot)$

Table 8: Never best response and strictly dominated

Contrary to this example, we will later see that the set of never best responses and strictly dominated strategies are "equivalent". For this equivalence, we need to define a *refined* notion of strict dominance, which we will discuss later. This will also capture the nuance about the extent of common knowledge of rationality required in iteratively deleting never best responses.

Similar to IESDS, we can also iteratively delete never best responses in every step. However, figuring out never best responses is not an easy task. Nevertheless, here is the straightforward adaptation of IESDS to never best response.

Definition **10 (Iterated elimination of never best responses)** *Start with a strategic form game $\Gamma = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$. Let $\Gamma^0 := \Gamma$, $S_i^0 := S_i$ for each $i \in N$. At any iteration $t > 0$,*

- *for every player i, let $K_i^{t-1} \subseteq S_i^{t-1}$ be the set of never best response strategies of Player i;*
- *set $S_i^t := S_i^{t-1} \setminus K_i^{t-1}$ for all $i \in N$;*
- *set $\Gamma^t := (N, \{S_i^t\}_{i \in N}, \{u_i^t\}_{i \in N})$, where u_i^t is the restriction of u_i to strategy profiles in $S_1^t \times \ldots \times S_n^t$.*
- *If $\Gamma^t = \Gamma^{t-1}$, then* STOP*; else set $t := t + 1$ and repeat.*

If the above process converges, let $\tilde{\Gamma}$ be the game to which it converges, and the strategies in $\tilde{\Gamma}$ for each player are called the strategies that **survive the iterative deletion of never best response process***.*

5.5 Beliefs: correlated and independent

Throughout, we assumed that belief is a probability distribution over strategy profiles of other players: $\mu_i \in \Delta S_{-i}$ or $\mu_i \in \Delta\Big(\prod_{j \neq i} S_j\Big)$. This allows a player to form correlated beliefs about other players's strategies, and hence, this is called a **correlated belief**. An **independent belief** of Player i is a probability distribution μ_i over other players strategies such that $\mu_i \in \prod_{j \neq i} \Big(\Delta S_j\Big)$. This means every Player i has a probability distribution over the strategies of each other Player j and her belief over the strategy profiles S_{-i} is just the product of these distributions.

If there are only two players in a game, independent and correlated beliefs are the same. But otherwise, independent beliefs are obviously a small subset of correlated beliefs. Consider a finite game with 3 players. Player 1 has three strategies: $\{A, B, C\}$, Player 2 has two strategies $\{a, b\}$, and Player 3 has two strategies $\{a', b'\}$. We only show the payoff of Player 1 in Table 9.

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

Table 9: Correlated beliefs example

Consider the following correlated belief of Player 1.

$$\mu_1(a, a') = \frac{1}{2},\ \mu_1(b, a') = 0,\ \mu_1(a, b') = 0,\ \mu_1(b, b') = \frac{1}{2}.$$

Note that such a correlated belief cannot be generated using an independent belief (why?). Now, an arbitrary belief of Player 1 is a function μ_1 which assigns the following non-negative numbers adding to 1:

$$\mu_1(a, a'),\ \mu_1(b, a'),\ \mu_1(a, b'),\ \mu_1(b, b').$$

Suppose we only consider independent beliefs: so, Player 1 believes that Player 2 plays a with probability p_2 and b with probability $(1 - p_2)$; Player 3 plays a' with probability p_3 and b' with probability $(1 - p_3)$. This results in the following beliefs:

$$\mu_1(a, a') = p_2 p_3,\ \mu_1(b, a') = (1 - p_2)p_3,\ \mu_1(a, b') = p_2(1 - p_3),\ \mu_1(b, b') = (1 - p_2)(1 - p_3).$$

The payoffs of Player 1 from this belief is given below:

$$\begin{aligned}
\mathcal{U}_1(A, \mu_1) &= 4p_2p_3 + 2\Big[(1 - p_2)p_3 + p_2(1 - p_3)\Big] + (1 - p_2)(1 - p_3) \\
\mathcal{U}_1(B, \mu_1) &= p_2p_3 + 2\Big[(1 - p_2)p_3 + p_2(1 - p_3)\Big] + 4(1 - p_2)(1 - p_3) \\
\mathcal{U}_1(C, \mu_1) &= 3\Big[p_2p_3 + (1 - p_2)(1 - p_3)\Big]
\end{aligned}$$

The difference in expected payoff are

$$\mathcal{U}_1(B,\mu_1) - \mathcal{U}_1(C,\mu_1) = 1 + p_2 + p_3 - 5p_2p_3$$
$$\mathcal{U}_1(A,\mu_1) - \mathcal{U}_1(C,\mu_1) = 4(p_2+p_3) - 2 - 5p_2p_3$$

We see that utility from B is better than A if and only if

$$\Big[\big[\mathcal{U}_1(B,\mu_1) - \mathcal{U}_1(C,\mu_1)\big] - \big[\mathcal{U}_1(A,\mu_1) - \mathcal{U}_1(C,\mu_1)\big] \geq 0\Big] \Leftrightarrow \Big[p_2 + p_3 \leq 1\Big].$$

But if $p_2+p_3 \leq 1$, then AM-GM inequality gives $1 \geq p_2+p_3 \geq 2\sqrt{p_2p_3}$. Hence, $p_2p_3 \leq \frac{1}{4}$. Using this, we get

$$\begin{aligned}\mathcal{U}_1(B,\mu_1) - \mathcal{U}_1(C,\mu_1) &= 1 + p_2 + p_3 - 5p_2p_3 \\ &= 1 + p_2(1-p_3) + p_3(1-p_2) - 3p_2p_3 \\ &\geq 1 - 3p_2p_3 \\ &\geq \frac{1}{4}.\end{aligned}$$

Similarly, if $p_2 + p_3 \geq 1$, then AM-GM inequality gives $(1-p_2)+(1-p_3) \leq 1$. Hence, $2\sqrt{(1-p_2)(1-p_3)} \leq (1-p_2)+(1-p_3) \leq 1$. This gives us $1 - 4(1-p_2)(1-p_3) \geq 0$. Now, use this to get:

$$\begin{aligned}\mathcal{U}_1(A,\mu_1) - \mathcal{U}_1(C,\mu_1) &= 4(p_2+p_3) - 2 - 5p_2p_3 \\ &= 4(p_2+p_3) - 3 - 4p_2p_3 + 1 - p_2p_3 \\ &\geq 1 + 4(p_2+p_3) - 4 - 4p_2p_3 \\ &= 1 - 4(1-p_2)(1-p_3) \\ &\geq 0.\end{aligned}$$

Hence, independent beliefs imply that C cannot be a best response to such beliefs.

However, consider the following correlated beliefs.

$$\mu_1(a,a') = \frac{1}{2},\ \mu_1(b,a') = 0,\ \mu_1(a,b') = 0,\ \mu_1(b,b') = \frac{1}{2}.$$

Payoffs of Player 1 are now: $\mathcal{U}_1(A, \mu_1) = \mathcal{U}_1(B, \mu_1) = 2.5, \mathcal{U}_1(C, \mu_1) = 3$. Hence, C is a best response. Notice that this correlated belief cannot be generated using independent beliefs.

The strict domination claims work out even if the beliefs are independent. For instance, consider the definition of strict dominance (Definition 2) and verify that Lemma 1 still works out with this definition.

5.6 Weakly dominated strategies

In many games, IESDS result in no strategies getting eliminated. This will happen if there are no strictly dominated strategies for any player. In such a case, the following weaker notion of weak domination may be considered.

DEFINITION **11** *Strategy s_i of Player i is* **weakly dominated** *if there exists another strategy t_i which weakly dominates s_i.*

We saw that common knowledge of rationality ensures that players can eliminate strictly dominated strategies iteratively. There is no such foundation for eliminating (iteratively or otherwise) weakly dominated strategies. Indeed, if we remove weakly dominated strategies iteratively, then strange things may happen. For instance, consider the following example in Table 10.

	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 10: Order of elimination of weakly dominated strategies

In IESDS, we allowed all players to eliminate all strictly dominated strategies in each iteration. The result will be same if we allowed each player to eliminate only a subset of strategies in each iteration in IESDS. This is not the case for weakly dominated strategies.

The game in Table 10, 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).

6 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 11. No strategy in this game is dominated.

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

Table 11: No dominated strategies

We 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 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 **12** *A strategy profile* $(s_1^*, \dots, s_n^*)$ *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 in the above definition may be a finite or an infinite game. The definition above requires that given strategies of other players s_{-i}^*, a unilateral deviation by Player i is not profitable. A belief based definition is also possible. In particular, the following is an equivalent definition. A strategy profile $(s_1^*, \ldots, s_n^*)$ is a Nash equilibrium if for all $i \in N$,

$$\mu_i(s_{-i}^*) = 1$$
$$\mathcal{U}_i(s_i^*; \mu_i) \geq \mathcal{U}_i(s_i; \mu_i) \qquad \forall\, s_i \in S_i.$$

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. If a player finds certain unilateral deviation profitable, then such a steady state cannot be sustained (and, hence, it cannot be a steady state).

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

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

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

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

Hence, Nash equilibrium requires non-emptyness of best response set at the equilibrium strategy profile.

The following observation is immediate.

CLAIM **2** *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 strictly dominant strategy profile, it is a* unique *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.

6.1 Examples

We give various examples of games where a Nash equilibrium exist. In Table 12, we consider the Prisoner's Dilemma game. By Claim 2, (A, a) is a Nash equilibrium of this game since it is the outcome in strictly dominant strategies.

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

Table 12: Nash equilibrium in Prisoner's Dilemma

Consider now the game (called the *coordination game*) in Table 13. 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 13: 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 14. 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 14: 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. Since the utility functions are strictly concave (verify this!), these will be the unique optimal solutions. We can also directly verify that it is a Nash equilibrium. For this, first note that

$$u_1(q_1^*, q_2^*) = (q_1^*)^2$$
$$u_2(q_1^*, q_2^*) = (q_2^*)^2$$

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 the unique Nash equilibrium.

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 15. 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 15: The Matching Pennies game

6.2 Elimination of Dominated Strategies

We first show how elimination of certain strategies does not lead to elimination of Nash equilibria. First, we show 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. In all the games below, we write Γ as a strategic form game - this may be a finite game or an infinite game or mixed extension of a finite game. The claims remain valid in all these cases.

Lemma **5** *Let Γ be a 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'. ■

Though eliminating arbitrary strategies will not eliminate original Nash equilibria, it may introduce new Nash equilibria. The game in Table 16 has a unique equilibria (B, L). Now, suppose we eliminate strategy T of Player 1. Then, Player 2 is indifferent between playing L and R. Hence, both (B, L) and (B, R) are Nash equilibria.

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

Table 16: Elimination may introduce new equilibria

An important feature of this example was that we eliminated an undominated strategy of a player. The following claim shows that this is not possible if weakly dominated strategies are eliminated.

LEMMA **6** *Let* Γ *be a 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 17. 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 Lemma 6).

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

Table 17: Elimination may lose equilibria

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

THEOREM **2** *Let* Γ *be a game in strategic form and* Γ^{IESDS} *be the game resulting from IESDS of* Γ*. If* $\text{NE}(\Gamma)$ *and* $\text{NE}(\Gamma^{\text{IESDS}})$ *are the set of Nash equilibria in* Γ *and* Γ^{IESDS} *respectively.*

Then,

$$\text{NE}(\Gamma) = \text{NE}(\Gamma^{\text{IESDS}})$$

Proof: Since every strictly dominated strategy is a weakly dominated strategy, by Lemma 6, if s^* is a Nash equilibrium of Γ^{IESDS}, it is also a Nash equilibrium of Γ. Hence,

$$\text{NE}(\Gamma) \supseteq \text{NE}(\Gamma^{\text{IESDS}}) \tag{8}$$

Next, we show that if s^* is a Nash equilibrium of Γ, then s^* survives IESDS. Assume for contradiction this is not the case. Then, there must exist a first round of the IESDS process where for some player i strategy s_i^* is eliminated. Importantly, since this is the first round where such a strategy is eliminated, the strategies $^*_{-i}$ of other players still exist. Since s_i^* is eliminated, there exists a strategy t_i which strictly dominates s_i^* in this stage of IESDS, i.e., $u_i(s_i^*, s_{-i}^*) < u_i(t_i, s_{-i}^*)$ (this inequality holds since s_{-i}^* is available in this round of IESDS). But this contradicts the fact that s^* is a Nash equilibrium (and hence, s_i^* is a best response to s_{-i}^*). Hence, s^* survives IESDS.

Then, by Lemma 5, s^* is a Nash equilibrium of Γ^{IESDS}. Hence, we get

$$\text{NE}(\Gamma) \subseteq \text{NE}(\Gamma^{\text{IESDS}}) \tag{9}$$

Putting (8) and (9) together, we get $\text{NE}(\Gamma) = \text{NE}(\Gamma^{\text{IESDS}})$. ■

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, we can iteratively eliminate all strictly dominated strategies of the players.

7 Existence of Nash Equilibrium

In many games Nash equilibria exist. The following natural question arises:

What are some sufficient conditions on the game that ensures existence of Nash *equilibrium?*

We will discuss some classes of games where we will show that a Nash equilibrium exists. The first one is somewhat technical in nature - but general enough to be applied to a large variety of games. The last one is somewhat simpler in nature, and the existence in those class of games were proved by Nash himself - in fact, the this class of games is a subclass of the first class of games, but even then we discuss it because of other reasons. All these classes of games have one thing in common: the strategy sets of each player has a lot of structure (geometrical) and the utility functions are well-behaved over these strategy sets.

An existence result is a technical result and may not appeal to everyone. However, it has its own beauty and importance. First, it shows that in some class of games, we can begin to think of computing and describing Nash equilibria. Second, it illustrates that the game is consistent with *some* steady state solution - though the precise steady state(s) are not found by proving an existence result.

All the existence results rely on some kind of **fixed point** result. We elaborate on this a little bit before proceeding further. Let X be some non-empty set and $f : X \to X$. We say $x \in X$ is a **fixed point** of f if

$$x = f(x).$$

A fixed point theorem identifies conditions on X and f such that a fixed point exists. These versions of fixed point theorems are indirectly useful.[2]

However, a set-theoretic version (or, correspondence version) of the fixed point theorem is immediately useful. As before, fix a set X and let $f : X \to 2^X$, where 2^X is the set of all subsets of X. So, for every $x \in X$, the function value $f(x)$ gives a subset of X. Such a function f has a **fixed point** $x \in X$ if

$$x \in f(x).$$

A fixed point theorem here would identify conditions on X and f such that a fixed point exists.

The usefulness of correspondence version of fixed point theorems is somewhat direct. Fix a strategy profile $s \in S$. Remember that the best response of agent i for s_{-i} is $B_i(s_{-i})$ and

[2]In one dimension, the *intermediate value theorem* is a stronger version of a fixed theorem.

it gives all the strategies that maximize agent i's payoff against s_{-i}. Define the function $B : S \to 2^S$ as follows: for every $s \in S$,

$$B(s) = B_1(s_{-1}) \times \ldots \times B_n(s_{-n}).$$

We refer to B as the **best response correspondence**.

Take the game in Table 18. Consider the strategy profile $s \equiv (s_1 = M, s_2 = L)$. Now, $B_1(s_2) = \{T\}$ and $B_2(s_1) = \{C, R\}$. Hence,

$$B(s_1, s_2) = \{T\} \times \{C, R\} = \{(T, C), (T, R)\}.$$

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

Table 18: Best response maps

The following claim establishes that such fixed point theorems will be useful for showing existence of Nash equilibrium.

Claim **3** *A Nash equilibrium exists if and only if the best response correspondence has a fixed point.*

Proof: If a Nash equilibrium s exists, then $s_i \in B_i(s_{-i})$ for all $i \in N$. Hence, $s \in B(s)$ - so, a fixed point of B exists. If a fixed point s of B exists, then $s \in B(s)$, which in turn implies that $s_i \in B_i(s_{-i})$. Hence, s is a Nash equilibrium. ■

Claim 3 forms the foundation for proving most of the existence results about Nash equilibrium. We will see this in next few sections.

7.1 Convex Strategy Sets with Concave and Continuous Utility Functions

The first such existence theorem is in a class of infinite games. The strategy space is assumed to have some geometric structure and the utility functions are assumed to be well-behaved.

THEOREM **3** *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. $u_i(s)$ *is continuous in* s.
3. $u_i(s_i, s_{-i})$ *is concave in* s_i.

Then, Γ *has a Nash equilibrium.*

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

THEOREM **4 (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 *is upper hemicontinuous, i.e., if* $\{x^k, y^k\} \to \{x, y\}$ *with* $y^k \in f(x^k)$ *for each* k*, then* $y \in f(x)$.[3]

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

[3] Compare this with the usual definition of continuity of a function. Suppose $g : X \to X$ is a function. Then, g is continuous if $\lim_k x^k = x$ implies $\lim_k g(x^k) = g(x)$. Put differently, if we denote $y^k = g(x^k)$. Then, it says if $\lim_k y_k = y$, then $f(x) = y$. The correspondence version is modified to reflect that $f(x^k)$ is subset of X. Hence, we take a sequence with $y^k \in f(x^k)$ and say f is upper hemicontinuous if whenever $\lim_k y_k = y$, we have $f(x) \ni y$. There is analogous definition for lower hemicontinuity, which says that if $y \in f(x)$, then there exists a sequence y^k converging to y such that $y^k \in f(x^k)$.

We use Theorem 4 in a straightforward manner to establish existence of Nash equilibrium. For every strategy profile s, we know by Claim 3 that s is a Nash equilibrium if and only if s is a fixed point of the best response correspondence B. We show that B satisfies all the conditions of Theorem 4, and we will be done.

1. Since each S_i is compact and convex, the set of strategy profiles $S_1 \times \ldots \times S_n$ 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.

3. Next, we show that $B(s)$ is convex. For this, we show that each $B_i(s_{-i})$ is convex. Since cartesian product of convex sets is a convex set, it will follow that $B(s)$ is convex. Pick $i \in N$ and s. Take $t_i, t'_i \in B_i(s_{-i})$ and let $t''_i = \lambda t_i + (1-\lambda)t'_i$ for some $\lambda \in (0,1)$. 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

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

   Now, pick any $s'_i \in S_i$ and note that $u_i(t_i, s_{-i}) \geq u_i(s'_i, s_{-i})$ and $u_i(t'_i, s_{-i}) \geq u_i(s'_i, s_{-i})$. Hence, $u_i(t''_i, s_{-i}) \geq \lambda u_i(t_i, s_{-i}) + (1-\lambda)u_i(t'_i, s_{-i}) \geq u_i(s'_i, s_{-i})$. This means $t''_i \in B_i(s_{-i})$. Hence, $B_i(s_{-i})$ is convex.

4. Finally, we show that B is upper hemicontinuous. To see this, assume for contradiction that B is not upper hemicontinuous. 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 $u_i(\bar{s}_i, t_{-i}) > u_i(\bar{t}_i, t_{-i})$ for some $\bar{s}_i \in S_i$. The argument then follows from continuity

of u_i in both his own strategy and the strategy of others. To show this, let us define $\Delta := u_i(\bar{s}_i, t_{-i}) - u_i(\bar{t}_i, t_{-i})$. Note that $\Delta > 0$ and choose $\delta > 0$ such that $\Delta > 2\delta$.

We find a tuple $(t^k, \bar{t}^k)$ which is arbitrarily close to $(t, \bar{t})$. The arbitrary closeness is such that the following relations are satisfied:

$$u_i(\bar{s}_i, t_{-i}) - u_i(\bar{s}_i, t^k_{-i}) < \delta \tag{10}$$

$$u_i(\bar{t}^k_i, t^k_{-i}) - u_i(\bar{t}_i, t_{-i}) < \delta \tag{11}$$

The inequalities follow from continuity of u_i. See Figure 4 for an illustration. To see precisely how this is possible, define a new function:

$$F(\bar{t}^k_i, t^k_{-i}) := \max\left(u_i(\bar{s}_i, t_{-i}) - u_i(\bar{s}_i, t^k_{-i}), u_i(\bar{t}^k_i, t^k_{-i}) - u_i(\bar{t}_i, t_{-i})\right)$$

Notice that $\lim_{k\to\infty} F(\bar{t}^k_i, t^k_{-i}) = 0$ as t^k_{-i} converges to t_{-i} and $\bar{t}^k_i$ converges to $\bar{t}_i$, and u_i is continuous. Then, there exists a k such that $F(\bar{t}^k_i, t^k_{-i}) < \delta$. Hence, (10) and (11) holds. This implies that

$$\left[u_i(\bar{s}_i, t_{-i}) - u_i(\bar{s}_i, t^k_{-i})\right] + \left[u_i(\bar{t}^k_i, t^k_{-i}) - u_i(\bar{t}_i, t_{-i})\right] < 2\delta.$$

Using the fact that $u_i(\bar{s}_i, t_{-i}) - u_i(\bar{t}_i, t_{-i}) = \Delta$, we get

$$u_i(\bar{s}_i, t^k_{-i}) - u_i(\bar{t}^k_i, t^k_{-i}) > \Delta - 2\delta > 0,$$

where the inequality follows from our assumption that $\Delta > 2\delta$. But this contradicts the fact that $\bar{t}^k_i \in B_i(t^k_{-i})$.

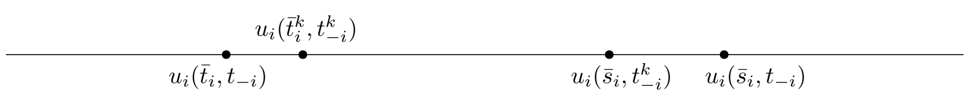

Figure 4: Illustration of proof of upper hemicontinuity

Now, we apply Kakutani's fixed point theorem (Theorem 4) to conclude that there exists s such that $s \in B(s)$. This implies that s is a Nash equilibrium. ■

To see how Theorem 3 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 for all values of $s_2 \neq \frac{1}{2}$. So, Theorem 3 cannot be applied here. But Nash equilibrium exists in such games - $s_1^* = s_2^* = \frac{1}{2}$ is a Nash equilibrium.

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 3. 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 3 guarantees that a Nash equilibrium exists.

8 Mixed Strategies

We now consider a game which is *derived* from a finite game. Formally, let

$$\Gamma := (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N}),$$

be a finite strategic form game (i.e., each S_i is finite). Consider the game derived from Γ by extending the strategy set of each player by allowing them to randomize over S_i.

Formally, the **mixed extension** of Γ is given by

$$\Delta\Gamma := (N, \{\Delta S_i\}_{i\in N}, \{U_i\}_{i\in N}),$$

where for all $i \in N$, the utility function U_i of Player i is a **linear extension** of his utility function u_i in Γ. In particular, if we consider a strategy profile $\sigma \in \prod_{i\in N} \Delta S_i$ in the mixed extension $\Delta\Gamma$, we have

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

where $\sigma_i(s_j)$ is the probability with which Player i plays strategy s_j of game Γ in the strategy σ_i of game $\Delta\Gamma$. Note that the mixed extension of a game is an infinite game - it includes all possible lotteries over pure strategies of a player.

For any finite strategy set S_i of Player i, every $\sigma_i \in \Delta S_i$ is called a **mixed strategy** of Player i. In this case S_i is called the set of **pure strategies** of Player i. In other words, mixed strategies are all the strategies of a player in the mixed extension. 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.

Consider the following game in Table 19. 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 19: 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 20. 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 20: Mixed strategies - probability of all pure strategy profiles

8.1 Extending the strategy space

Since $\Delta\Gamma$ is derived from Γ, the first question to ask is what happens to dominated and dominant strategies, and Nash equilibria of Γ when we consider $\Delta\Gamma$. This is a relevant question because the set of strategies in $\Delta\Gamma$ is larger than Γ. Hence, not only does a player have more strategies to play but also has more strategies to deviate. For instance, is it possible that a Player i may have a strictly dominant pure strategy s_i in Γ but in the mixed extension $\Delta\Gamma$, there may be a mixed strategy which s_i no longer strictly dominates? The following lemma is useful in understanding such relationships between a finite game and its mixed extension.

Lemma **7 (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 $U_i(s_i, \sigma_{-i}) = U_i(s_i', \sigma_{-i})$.

Suppose not and $U_i(s_i, \sigma_{-i}) > U_i(s_i', \sigma_{-i})$. Then,

$$\begin{aligned}
U_i(\sigma_i, \sigma_{-i}) &= \sum_{s_i'' \in S_i(\sigma_i)} U_i(s_i'', \sigma_{-i})\sigma_i(s_i'') \\
&= U_i(s_i, \sigma_{-i})\sigma_i(s_i) + U_i(s_i', \sigma_{-i})\sigma_i(s_i') + \sum_{s_i'' \in S_i(\sigma_i)\setminus\{s_i, s_i'\}} U_i(s_i'', \sigma_{-i})\sigma_i(s_i'') \\
&< U_i(s_i, \sigma_{-i})\big(\underbrace{\sigma_i(s_i) + \sigma_i(s_i')}_{\sigma_i'(s_i)}\big) + U_i(s_i', \sigma_{-i}) \cdot \underbrace{0}_{\sigma_i'(s_i')} + \sum_{s_i'' \in S_i(\sigma_i)\setminus\{s_i, s_i'\}} U_i(s_i'', \sigma_{-i}) \underbrace{\sigma_i(s_i'')}_{\sigma_i'(s_i'')} \\
&= U_i(\sigma_i', \sigma_{-i}),
\end{aligned}$$

where σ_i' is the new mixed strategy of Player i, where he plays s_i with probability $\sigma_i(s_i)+\sigma_i(s_i')$ and s_i' with probability zero, and every other strategy s_i'' in $S_i(\sigma_i)$ is played with probability $\sigma_i(s_i'')$. But this contradicts the fact that $\sigma_i \in B_i(\sigma_{-i})$.

This means that $U_i(s_i, \sigma_{-i}) = U_i(s_i', \sigma_{-i})$ for all $s_i, s_i' \in S_i(\sigma_i)$. We denote this utility as $\Pi_i(\sigma_{-i})$. Then,

$$U_i(\sigma_i, \sigma_{-i}) = \sum_{s_i'' \in S_i(\sigma_i)} U_i(s_i'', \sigma_{-i})\sigma_i(s_i'') = \Pi_i(\sigma_{-i}).$$

This proves the claim. ■

This allows us to state the following straightforward results.

THEOREM **5** *Suppose $s^* \equiv (s_1^*, \ldots, s_n^*)$ is a strategy profile in the finite game Γ. Then, the following are true.*

1. *If s^* is a Nash equilibrium of Γ, it is also a Nash equilibrium of the mixed extension $\Delta\Gamma$.*

2. *Every strictly dominant strategy of $\Delta\Gamma$ is a pure strategy, i.e., a strategy in Γ.*

Proof: PROOF OF (1). Suppose s^* is not a Nash equilibrium of $\Delta\Gamma$. Then, for some $i \in N$, $s_i^* \notin B_i(s_{-i}^*)$. Then, some $\sigma_i \in B_i(s_i^*)$. By Lemma 7, there is some strategy $s_i' \in S_i$ such that $\sigma_i(s_i') > 0$ and $s_i' \in B_i(s_{-i}^*)$. This means, $u_i(s_i', s_{-i}^*) > u_i(s_i^*, s_{-i}^*)$. This contradicts the fact that s^* is a Nash equilibrium of Γ.

Proof of (2). Suppose σ_i is a strategy in $\Delta\Gamma$ but not in Γ (i.e., σ_i is **not** a pure strategy) and σ_i is strictly dominant in $\Delta\Gamma$. Then, by Lemma 7, there are two strategies $s_i \neq s_i'$ belonging to Γ such that $\sigma_i(s_i) > 0$ and $\sigma_i(s_i') > 0$, and for all s_{-i},

$$U_i(s_i, s_{-i}) = U_i(s_i', s_{-i}) = U_i(\sigma_i, s_{-i}).$$

Hence, σ_i is not strictly dominant. ■

Theorem 5 has consequences in computing a Nash equilibrium in the mixed extension of Γ. It says that we can compute Nash equilibria and strictly dominant strategies of Γ, and they continue to maintain their properties in the mixed extension. The following remarks say that mixed extensions may create additional complications.

- 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 21. Strategy C is not dominated by any pure strategy for Player 1. However, the mixed strategy $\frac{1}{2}A$ and $\frac{1}{2}B$ 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 21.

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

Table 21: Mixed strategies may dominate pure strategies

- 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 22. 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.

8.2 Existence of Nash Equilibrium in Mixed Strategies

In this section, we prove Nash's seminal theorem. It is arguably the most famous theorem in game theory.

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

Table 22: Mixed strategies may be dominated

THEOREM **6 (Nash)** *The mixed extension of every finite game has a Nash equilibrium.*

Note that this theorem is a corollary of our earlier existence theorem - Theorem 3. This is because, it is not difficult to check that the strategy space in the mixed extension of a finite game is a convex set, the utility functions are linear in strategies, and hence, continuous and concave as desired. The proof below is based on a weaker fixed theorem due to Brower. It is also based on the original proof of Nash, and has a useful technique that can be applied in other settings. Historical fact: Theorem 3 was proved later (in 1952 by Debreu, Fan, and Gricksberg, all independently) and Nash's theorem came in 1950.

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 $s_i \in S_i$, we define

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

The interpretation of $g_i(s_i,\sigma)$ is that it is zero if Player i does not find deviating to s_i from σ profitable. Else, it captures the increase in payoff of Player i from (σ) to (s_i,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 7. This implies that σ is a Nash equilibrium if and only if $g_i(s_i,\sigma) = 0$ for all $i \in N$ and for all $s_i \in S_i$.

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

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

$$f_i(s_i, \sigma) := \frac{\sigma_i(s_i) + g_i(s_i, \sigma)}{1 + \sum_{s'_i} g_i(s'_i, \sigma)}.$$

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

It is easy to see that for each i and each s_i, $f_i(s_i, \sigma) \geq 0$. Further,

$$\begin{aligned}\sum_{s_i \in S_i} f_i(s_i, \sigma) &= \sum_{s_i \in S_i} \frac{\sigma_i(s_i) + g_i(s_i, \sigma)}{1 + \sum_{s'_i \in S_i} g_i(s'_i, \sigma)} \\ &= \frac{\sum_{s_i \in S_i} \sigma_i(s_i) + \sum_{s_i \in S_i} g_i(s_i, \sigma)}{1 + \sum_{s_i \in S_i} g_i(s_i, \sigma)} \\ &= 1.\end{aligned}$$

Denote by $f_i(\sigma)$ the vector of probabilities $\{f_i(s_i, \sigma)\}_{s_i \in S_i}$. Hence, $f_i(\sigma)$ is another mixed strategy of Player i. Further, f_i 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. 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 s_i,

$$g_i(s_i, \sigma) = \sigma_i(s_i) \sum_{s'_i \in S_i} g_i(s'_i, \sigma).$$

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

$$\begin{aligned}f_i(s_i, \sigma) &= \sigma_i(s_i) \\ &= \frac{\sigma_i(s_i) + g_i(s_i, \sigma)}{1 + \sum_{s'_i \in S_i} g_i(s'_i, \sigma)}.\end{aligned}$$

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 such that $g_i(s_i, \sigma) > 0$ - this uses Lemma 7 because we are claiming that a pure strategy gives more payoff. As a result $\sum_{s'_i \in S_i} g_i(s'_i, \sigma) > 0$. From the previous step, we know that for every s''_i, we have $\sigma_i(s''_i) > 0$ if and only if $g_i(s''_i, \sigma) > 0$. In other words, for every s''_i, $\sigma_i(s''_i) > 0$ if and only if $U_i(\sigma) < U_i(s''_i, \sigma_{-i})$, i.e., every pure strategy in the support of σ_i gives strictly more payoff than σ_i itself (when others play σ_{-i}). This is clearly a contradiction – here is the formal argument:

$$U_i(\sigma) = \sum_{s''_i \in S_i} \sigma_i(s''_i) U_i(s''_i, \sigma_{-i}) = \sum_{s''_i : \sigma_i(s''_i) > 0} \sigma_i(s''_i) U_i(s''_i, \sigma_{-i}) > U_i(\sigma) \sum_{s''_i : \sigma_i(s''_i) > 0} \sigma_i(s''_i) = U_i(\sigma)$$

This gives us a contradiction.

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. ■

The Brouwer's fixed point theorem is not simple to prove, but you are encouraged to look at its proof. In one-dimension, the Brouwer's fixed point theorem is the *intermediate value theorem*.

8.3 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 randomize deliberately. For instance, in zero-sum games with two players, players may randomize. In games like Poker, players have been shown to randomize.
- 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.
- 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 population.

9 Computing Nash Equilibrium

We discuss some issues related to computation of Nash equilibrium. The proof of Nash' theorem already provides a method to compute a Nash equilibrium by searching for a fixed point of an appropriate function. Usually, such fixed point calculations are not easy (and may take infinite iterations). Below, we highlight some simplifications one can do before computing Nash equilibrium. Then, we discuss a graphical method of computing Nash equilibrium in 2-person finite games with at most three pure strategies.

Before we do that we consider the procedure in Nash' proof for a particular game – battle of sexes (Table 23). This game has two pure strategy Nash equilibria: (T, L) and (B, R). So, if we start Nash' proof at any of these two profiles, we will immediately stop. So, lets start with $\sigma = (B, L)$. Then, we compute g_i as follows.

$$U_1(\sigma) = U_2(\sigma) = 0$$
$$g_1(T, \sigma) = 2, g_1(B, \sigma) = 0$$
$$g_2(L, \sigma) = 0, g_2(R, \sigma) = 2$$
$$f_1(T, \sigma) = \frac{0+2}{1+2} = \frac{2}{3}; f_1(B, \sigma) = \frac{1}{3}$$
$$f_2(L, \sigma) = \frac{1+0}{1+2} = \frac{1}{3}; f_2(R, \sigma) = \frac{2}{3}.$$

This is the first iteration and we can update our σ now:

$$\sigma = \Big(\frac{2}{3}T + \frac{1}{3}B, \frac{1}{3}L + \frac{2}{3}R\Big).$$

The payoffs of the players are as follows.

$$U_1(\sigma) = \frac{2}{3}, U_2(\sigma) = \frac{2}{3}$$
$$g_1(T, \sigma) = 0, g_1(B, \sigma) = 0$$
$$g_2(L, \sigma) = 0; g_2(R, \sigma) = 0.$$

Hence, this is a Nash equilibrium. Although for this game, finding a Nash equilibrium through Nash' technique was easy, it is not the case in general (you are encouraged to take an arbitrary game with two players and two strategies, and try to find a mixed strategy Nash equilibrium using Nash' technique).

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

Table 23: Battle of sexes game Nash equilibria

9.1 Mixed strategy equilibrium computation - examples

In general, computing mixed strategy equilibrium of a finite 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 7.

We start off by a simple example on how to compute all Nash equilibria of a game. Consider the game in Table 24.

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

Table 24: 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.

Consider the two player game in Table 25. Computing Nash equilibria of such a game can be quite tedious. However, we can be smart in avoiding certain computations.

In two player 3-strategy games, we can draw the best response correspondences in a 2-d simplex - Figure 5 represents the simplex of Player 1's strategy space for the game in Table 25.

	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 25: Nash equilibria computation

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.

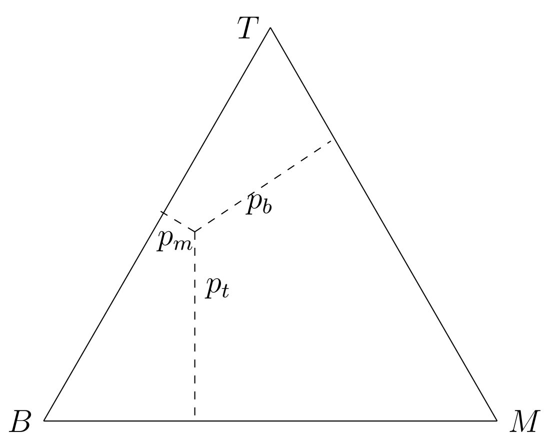

Figure 5: Representing probabilities on a 2d-simplex

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 6. 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 6.

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 6.

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 6 in the simplex of Player 2's strategy.

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

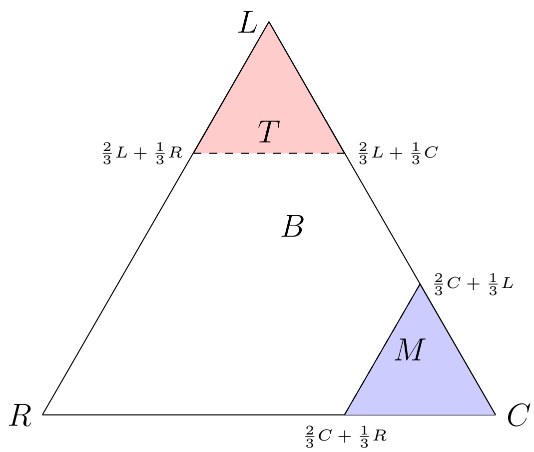

Figure 6: Best response map of Player 1

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 6. 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 6.

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

Computing Nash equilibria. To compute Nash equilibria, we see that there is no best

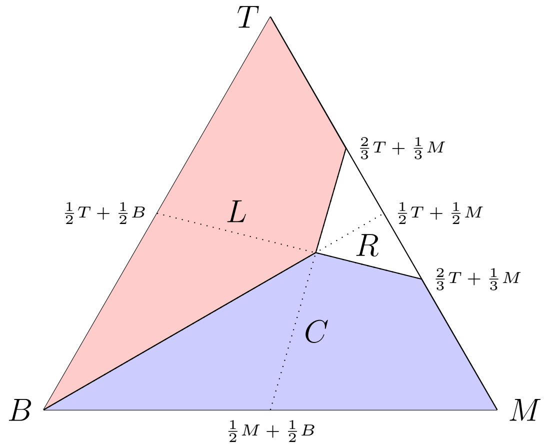

Figure 7: Best response map of Player 2

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 mixing of T and M is not 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 7. 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 6, 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 25 now:

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

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

10 Two Player Zero-Sum Games

The two player zero-sum games occupy an important 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 found in many real-life applications - examples include any real game where one player's loss is another player's gain. Before formally introducing the notion of a zero-sum game, we describe another concept that we use here.

10.1 The Maxmin Value

Consider a game shown in Table 26. 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 26: 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 13 *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, i.e., $\bar{s}_i$ is a maxmin strategy of Player i if*

$$\bar{s}_i \in \arg\max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i}).$$

Note that the above definition allows us to consider games which are mixed extensions of some finite game too. In that case, the max and min over strategy space is well defined because the set of strategies is a compact space and the utility function is linear in (mixed) strategies.

If s_i is a maxmin strategy for Player i, then it satisfies

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

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 26, 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 27. 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 27: 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 (such a strategy must be unique), then the profile of weakly dominant strategies constitute a unique profile of maxmin strategies. This was true, for instance, in the example involving the second-price sealed-bid auction.

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

THEOREM 8 *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

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

■

10.2 Zero-sum games

We now look into two-player zero-sum games. Formally, a zero-sum game is defined as follows.

DEFINITION 14 *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.

Consider the two player zero-sum game in Table 28. It is called the *matching pennies* game - the strategies are sides of a coin, if the sides match then Player 1 wins and pays

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

Table 28: Matching pennies

Player 2 Rs. 1, else Player 2 wins and pays Player 1 Rs. 1. There is no pure strategy Nash equilibrium of this game. However, once we start looking at its mixed extension, we observe some interesting facts. 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 mixed strategy 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 (in the mixed extension) can be computed as follows. If Player 2 plays h or t Player 1 gets a payoff of 0. Hence, whatever Player 2 plays (by mixing), Player 1 will get zero as payoff from his strategy $\frac{1}{2}H + \frac{1}{2}T$. So $\underline{v}_1 \geq 0$. By Theorem 8, we know that $0 \geq \underline{v}_1$. As a result, the maxmin value of Player 1 is at least zero. We know (by Theorem 8) that the Nash equilibrium payoff is at least the maxmin value. [4] 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

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

[4] Theorem 8 continues to hold even we allow consider the mixed extension of a finite game.

Any maxmin and minmax strategies of Player 1 and Player 2 respectively are called **optimal** strategies.

The main result for (mixed extension of) two person zero-sum game is the following.

THEOREM **9** *The following are true for mixed extension of any two player zero-sum game.*

1. *The payoff from any Nash equilibrium* (σ_1^*, σ_2^*) *corresponds to* $(\underline{v}_1, \underline{v}_2)$. *Hence, if* (σ_1^*, σ_2^*) *is a Nash equilibrium, they are also the optimal (max-min) strategies.*

2. $\underline{v}_1 + \underline{v}_2 = 0$, *i.e.,*

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

3. *If* (σ_1^*, σ_2^*) *are max-min strategies, they are also a Nash equilibrium strategy profile.*

Proof: PROOF OF (1). Let σ^* be a Nash equilibrium profile. Nash equilibrium condition for Player 1 implies,

$$u(\sigma_1^*, \sigma_2^*) = \max_{\sigma_1 \in \Delta S_1} u(\sigma_1, \sigma_2^*) \geq \min_{\sigma_2 \in \Delta S_2} \max_{\sigma_1 \in \Delta S_1} u(\sigma_1, \sigma_2) = -\underline{v}_2. \tag{12}$$

Note that by Theorem 8, $-u(\sigma_1^*, \sigma_2^*) \geq \underline{v}_2$. Hence, we have $-u(\sigma_1^*, \sigma_2^*) = \underline{v}_2$. Hence, the inequality in Inequality (12) must be an equality:

$$\underline{v}_2 = \min_{\sigma_1 \in \Delta S_1} [-u(\sigma_1, \sigma_2^*)].$$

Hence, σ_2^* is a max-min strategy giving Player 2 a payoff of $\underline{v}_2 = -u(\sigma_1^*, \sigma_2^*)$.

The proof for Player 1 is identical. Nash equilibrium condition for Player 2 implies that for all $\sigma_2 \in \Delta S_2$, we have $-u(\sigma_1^*, \sigma_2^*) \geq -u(\sigma_1^*, \sigma_2)$. Hence,

$$u(\sigma_1^*, \sigma_2^*) = \min_{\sigma_2 \in \Delta S_2} u(\sigma_1^*, \sigma_2) \leq \max_{\sigma_1 \in \Delta S_1} \min_{\sigma_2 \in \Delta S_2} u(\sigma_1, \sigma_2) = \underline{v}_1. \tag{13}$$

By Theorem 8, $u(\sigma_1^*, \sigma_2^*) \geq \underline{v}_1$. Hence, we get $u(\sigma_1^*, \sigma_2^*) = \underline{v}_1$. This implies that the inequality in Inequality (13) is an equality:

$$\underline{v}_1 = \min_{\sigma_2 \in \Delta S_2} u(\sigma_1^*, \sigma_2).$$

Hence, σ_1^* is a maxmin strategy.

Proof of (2). Every game has a Nash equilibrium in mixed strategies. [5] If σ^* is a Nash equilibrium of the zero-sum game, then the zero-sum game property ensures that

$$u_1(\sigma_1^*, \sigma_2^*) + u_2(\sigma_1^*, \sigma_2^*) = \sum_{s \in S} \big[u_1(s)\sigma^*(s) + u_2(s)\sigma^*(s)\big] = 0.$$

By (1) above, $u_1(\sigma_1^*, \sigma_2^*) + u_2(\sigma_1^*, \sigma_2^*) = \underline{v}_1 + \underline{v}_2 = 0$.

Proof of (3). Since (σ_1^*, σ_2^*) are max-min strategies of the players, we write

$$\underline{v}_1 = \min_{\sigma_2} u(\sigma_1^*, \sigma_2)$$
$$\underline{v}_2 = \min_{\sigma_1} -u(\sigma_1, \sigma_2^*) = -\max_{\sigma_1} u(\sigma_1, \sigma_2^*).$$

Using (2), we get $\underline{v}_1 + \underline{v}_2 = 0$, and hence,

$$\min_{\sigma_2} u(\sigma_1^*, \sigma_2) = \max_{\sigma_1} u(\sigma_1, \sigma_2^*).$$

But this implies that

$$u(\sigma_1^*, \sigma_2^*) \geq \min_{\sigma_2} u(\sigma_1^*, \sigma_2) = \max_{\sigma_1} u(\sigma_1, \sigma_2^*) \geq u(\sigma_1^*, \sigma_2^*).$$

Hence, the above inequalities are all equalities.

$$u(\sigma_1^*, \sigma_2^*) = \max_{\sigma_1} u(\sigma_1, \sigma_2^*) \quad \text{and} \quad u(\sigma_1^*, \sigma_2^*) = \min_{\sigma_2} u(\sigma_1^*, \sigma_2).$$

This is equivalent to writing

$$u(\sigma_1^*, \sigma_2^*) = \max_{\sigma_1} u(\sigma_1, \sigma_2^*) \quad \text{and} \quad -u(\sigma_1^*, \sigma_2^*) = \max_{\sigma_2} -u(\sigma_1^*, \sigma_2).$$

Hence, (σ_1^*, σ_2^*) is a Nash equilibrium. ■

[5] The original proof of this theorem, due to von-Neumann and Morgenstern predates Nash' existence theorem. Their proof does not use the fact that every game has a mixed strategy Nash equilibrium.

11 A Foundation for Iterated Elimination

In this section, we revisit the iterated elimination of strictly dominated strategies (IESDS) procedure. We look at IESDS in the mixed extension of a finite game but we only eliminate strictly dominated **pure** strategies. The objective is to develop a foundation for eliminating strictly dominated strategies in a finite game using IESDS. The foundation we provide relates it to a new concept in strategic form game called *rationalizability*.

To fix ideas, we are given a finite strategic form game: $\Gamma := (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$. We are going to consider the mixed extension of this game. But we will only be concerned with eliminating pure strategies from this mixed extension. To remind, Theorem 5 has told us that if a pure strategy s_i is strictly dominated for Player i, every mixed strategy σ_i with s_i in its support is also strictly dominated. Hence, eliminating a pure strategy also eliminates *some* strictly dominated mixed strategies. However, as we have seen earlier, it may not eliminate *all* strictly dominated mixed strategies. The foundation we provide here is about iterated elimination of strictly dominated *pure* strategies.

11.1 Never best response and strict domination

Before we formally analyze the IESDS procedure for pure strategies, we prove the equivalence of never-best response strategies and strictly dominated strategies. Though we provide this equivalence result for pure strategies, it also holds for mixed strategies.

We remind the definitions of strict domination and never best response in terms of beliefs of players. A belief of Player i is a probability distribution μ_i over the strategies S_{-i} of other players.

Definition 15 *A strategy s_i of Player i in Γ is a* **strictly dominated** *if there exists a mixed strategy σ_i such that for every belief $\mu_i \in \Delta S_{-i}$,*

$$\mathcal{U}_i(\sigma_i, \mu_i) > \mathcal{U}_i(s_i, \mu_i).$$

Note here that the belief of Player i is a correlated belief and we do not restrict ourselves to independent beliefs. Our original definition is belief based but we provide two more definitions in strategy space (something like this we had seen without considering the mixed extension).

Claim 4 *The following are equivalent for any $s_i \in S_i$:*

1. *there is a strategy $\sigma_i \in \Delta S_i$ such that $U_i(\sigma_i, \sigma_{-i}) > U_i(s_i, \sigma_{-i})$ for all $\sigma_{-i} \in \Delta S_{-i}$.*

2. *there is a strategy $\sigma_i \in \Delta S_i$ such that $U_i(\sigma_i, s_{-i}) > U_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}$.*

3. *there exists a mixed strategy σ_i such that for every belief $\mu_i \in \Delta S_{-i}$, $\mathcal{U}_i(\sigma_i, \mu_i) > \mathcal{U}_i(s_i, \mu_i)$.*

Proof: (1) $\Rightarrow$ (2). This is surely true since every s_{-i} also belongs to ΔS_{-i}.

(2) $\Rightarrow$ (3). Pick any belief $\mu_i \in \Delta S_{-i}$ and note that

$$\mathcal{U}_i(s_i, \mu_i) = \sum_{s_{-i}} U_i(s_i, s_{-i})\mu_i(s_{-i}) < \sum_{s_{-i}} U_i(\sigma_i, s_{-i})\mu_i(s_{-i}) = \mathcal{U}_i(\sigma_i, \mu_i).$$

Hence, s_i is strictly dominated.

(3) $\Rightarrow$ (1). This is obvious since every σ_{-i} corresponds to an independent belief of Player i: $\mu_i(s_{-i}) := \times_{j \neq i}\sigma_j(s_j)$. ■

So, Claim 4 says that even though we do not consider belief over mixed strategy profiles of all players, considering beliefs over pure strategies of other players is equivalent to considering domination in the sense of (1) in Claim 4.

Similarly, we remind the definition of a never best response.

Definition 16 *A strategy s_i is a* **never best response** *for Player i if for every belief $\mu_i \in \Delta S_{-i}$, there exists a strategy σ_i such that $\mathcal{U}_i(\sigma_i, \mu_i) > \mathcal{U}_i(s_i, \mu_i)$.*

Again, we consider mixed strategy σ_i but only beliefs over pure strategies of other players. Note here that the belief of Player i is a correlated belief and we do not restrict ourselves to independent beliefs.

Obviously, a strictly dominated strategy is a never best response – proof is similar to Claim 1. However, a never best response pure strategy may not be strictly dominated in the finite game Γ as the game in Table 29 shows. Here, strategy B of Player 1 is not strictly dominated but it is a never best response. This is not true if we consider the mixed extension

$\Delta\Gamma$. In the mixed extension $\Delta\Gamma$, strategy B is still a never best response but it is also strictly dominated by $\frac{1}{2}T + \frac{1}{2}M$.

	L	M
T	$(1, \cdot)$	$(0, \cdot)$
M	$(0, \cdot)$	$(1, \cdot)$
B	$(0.4, \cdot)$	$(0.4, \cdot)$

Table 29: Never best response and strictly dominated

This point is formally proved in the theorem below.

Theorem **10** *A pure strategy of a player in the mixed extension of a finite strategic form 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 $\bar{s}_j \in S_j$, which is a never best response in the mixed extension $\Delta\Gamma$. Consider a new (artificial) game in which there are just two players j and $-j$. The set of strategies available to Player j is $S'_j := S_j \setminus \{\bar{s}_j\}$ and to Player $-j$ is S_{-j} (i.e., every strategy profile of players in $N \setminus \{j\}$ is interpreted as a strategy of Player $(-j)$). The utility of Player j at strategy profile (s_j, s_{-j}) is:

$$v_j(s_j, s_{-j}) = u_j(s_j, s_{-j}) - u_j(\bar{s}_j, s_{-j}).$$

We will abuse notation and denote the payoff from a mixed strategy profile (σ_j, σ_{-j}) to Player j as $v_j(\sigma_j, \sigma_{-j})$. Note that in $\Delta\Gamma'$, a (mixed) strategy σ_{-j} of Player $(-j)$ is a belief of Player j in Γ and vice versa. For any mixed strategy profile (σ_j, σ_{-j}) we have

$$v_j(\sigma_j, \sigma_{-j}) = \mathcal{U}_j(\sigma_j, \mu_j = \sigma_{-j}) - \mathcal{U}_j(\bar{s}_j, \mu_j = \sigma_{-j})$$

The payoff to Player $-j$ is negative of payoff to Player j – hence, it is a zero-sum game. Denote this game as Γ' and consider its mixed extension $\Delta\Gamma'$.

Hence, strategy $\bar{s}_j$ is a never-best response in Γ implies for every belief μ_j of Player j there is a strategy σ_j such that $\mathcal{U}_j(\sigma_j, \mu_j) - \mathcal{U}_j(\bar{s}_j, \mu_j) > 0$. But each such belief μ_j must correspond to a strategy σ_{-j} of Player $-j$ in Γ'. Since never best response means the above inequality holds for **all** beliefs μ_j (including those which are not independent), we conclude that for **every** mixed strategy σ_{-j} of Player $-j$ in Γ', there exists a strategy σ_j such that $\mathcal{U}_j(\sigma_j, \sigma_{-j}) - \mathcal{U}_j(\bar{s}_j, \sigma_{-j}) > 0$, or, $v_j(\sigma_j, \sigma_{-j}) > 0$.

Let $(\sigma_j^*, \sigma_{-j}^*)$ be a Nash equilibrium of this game. Hence, $v_j(\sigma_j^*, \sigma_{-j}^*) > 0$. By Theorem 9, $v_{-j}(\sigma_j^*, \sigma_{-j}^*) < 0$. Since $(\sigma_j^*, \sigma_{-j}^*)$ is a Nash equilibrium of $\Delta\Gamma'$, we get that for every strategy σ_{-j} of Player $(-j)$, we have

$$-v_j(\sigma_j^*, \sigma_{-j}) \leq -v_j(\sigma_j^*, \sigma_{-j}^*) < 0,$$

which gives $v_j(\sigma_j^*, \sigma_{-j}) > 0$ for all σ_{-j}. Hence,

$$\mathcal{U}_j(\sigma_j^*, \sigma_{-j}) > \mathcal{U}_j(\bar{s}_j, \sigma_{-j}) \ \forall \ \sigma_{-j}.$$

But σ_{-j} corresponds to a belief of Player j in Γ. Hence, for *every* belief μ_j of Player j, we have

$$\mathcal{U}_j(\sigma_j^*, \mu_j) > \mathcal{U}_j(\bar{s}_j, \mu_j).$$

Hence, $\bar{s}_j$ is strictly dominated for Player j. ■

Remember that the equivalence in Theorem 10 is only valid if we allow for correlated beliefs – of course, for two-player games these correlated belief is same as independent belief. In Section 5.5, we had given an example of a three player game where there was a pure strategy which was a never best response strategy if we only considered independent beliefs but it was a best response if we considered correlated beliefs (that strategy was not strictly dominated).

Theorem 10 extends even if we consider mixed strategies. That is, for any Player i and any mixed strategy σ_i, we can prove that σ_i is strictly dominated if and only if it is a never best response. The proof of this fact is more intricate and skipped.

11.2 Correlated Rationalizability

The fact that pure strategy Nash equilibrium does not exist makes it problematic as a solution concept some times. Mixed strategies are not entirely convincing since players play pure strategies at the end of the game anyway. So, *what are the pure strategy profiles we can* *expect to see as outcome of a game?* The notion of correlated rationalizability is developed as a "set theoretic" pure strategy Nash equilibrium. Instead of predicting a unique strategy to be played by each player, we will say that a player *may* play any strategy from a set as long as it is best response with respect to *some* belief over the strategy sets chosen by other players.

DEFINITION **17** *A profile of set of strategies* $(Z_1, \ldots, Z_n)$*, where* $Z_i \subseteq S_i$ *for each* $i \in N$*, is* **rationalizable** *in the strategic form game* $(N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ *if for every* $i \in N$ *and every* $s_i \in Z_i$ *there is a belief* μ_i *whose support is a subset of* Z_{-i} *such that*

$$\mathcal{U}_i(s_i, \mu_i) \geq \mathcal{U}_i(s'_i, \mu_i) \qquad \forall \; s'_i \in S_i,$$

i.e., s_i *is a best response with respect to belief* μ_i*.*

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 (i.e., even with respect to strategies outside Z_i for each i). The story is that an outside observer collects data of the game being played many times. She observes that each Player i has played Z_i set of strategies. Can she **rationalize** the strategies played by the players? Here, rationalizability puts two restrictions: (1) players are allowed to form beliefs on the observed strategies only; (2) players have to play a strategy which is a best response with respect to **some** belief of theirs, where beliefs respect (1). The rationalizability question is:

what observation (data) of the observer can be rationalized?

The aim of this section is to answer this question.

Note that if a profile of set of strategies $(Z_1, \ldots, Z_n)$ is rationalizable and another profile of set of strategies $(Z'_1, \ldots, Z'_n)$ is rationalizable then the profile of set of strategies $(Z_1 \cup Z'_1, \ldots, Z_n \cup Z'_n)$ is also rationalizable (Why?). Hence, the set of rationalizable strategies is the largest collection of $\{Z_j\}_j$ that can be rationalized.

Consider the example in Table 30. $(\{A\}, \{a\})$ 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 30: Two Player Game

Strategies in the support of Nash equilibria are rationalizable.

LEMMA **8** *Let E_i denote the set of pure strategies of Player i which are in support of some (possibly mixed) Nash equilibrium. Then, $(E_1, \ldots, E_n)$ is rationalizable.*

Proof: Suppose s_i is a strategy of Player i in the support of Nash equilibrium σ^*. Now for every j, let Z_j, which is a subset of E_j, be the strategies in the support of the Nash equilibrium σ^*. Now, define the belief μ_i of Player i as the product

$$\mu_i(s_{-i}) := \times_{j \neq i} \sigma_j^*(s_j) \qquad \forall\ s_{-i} \in Z_{-i}.$$

By the definition of Nash equilibrium σ_i^* is a best response to σ_{-i}^*, which means that it is a best response with respect to belief μ_i. By the indifference lemma, each s_i in the support of σ_i^* is a best response of i with respect to the belief μ_i. ■

In general, finding the set of rationlizable strategies can be quite cumbersome. Below, we provide an easy method with the help of a cute result.

Couple of quick observations are worth making. First, if a strategy is strictly dominated, then it is not rationalizable. But we can say more. Our next theorem says that rationalizability and iterated elimination of strictly dominated pure strategies are equivalent.

THEOREM **11** *Suppose $(Z_1, \ldots, Z_n)$ is the profile of the largest set of rationalizable strategies. Let $(X_1, \ldots, X_n)$ be a profile of set of strategies available after iterated elimination of strictly dominated pure strategies. Let $(Y_1, \ldots, Y_n)$ be a profile of set of strategies available after* *iterated elimination of never best response strategies. Then,*

$$X_i = Y_i = Z_i \;\forall\; i \in N.$$

Proof: Let $(Z_1, \ldots, Z_n)$ be the *largest* set of rationalizable strategies for each player. We will argue that Z_i survives iterated elimination of strictly dominated strategies for each $i \in N$. Suppose not. Then, consider the first period t in the IESDS procedure where strategy $s_i \in Z_i$ of some Player i gets eliminated in iterated elimination procedure. Since this is the first period where such a strategy is getting eliminated, all the strategies Z_j of $j \neq i$ still exists in the game Γ^t in period t – as before, we denote the game in any period t of the IESDS procedure as Γ^t and the strategies available to Player i as X_i^t. Hence, s_i is strictly dominated in Γ^t implies (by Theorem 10) that s_i is a never-best-response strategy for this game. Since s_i is rationalizable, there is a belief μ_i over $\times_{j\neq i} Z_j$ such that s_i is a best response with respect to μ_i. This is a contradiction to s_i being a never-best-response in Γ^t since strategy profiles in Z_{-i} are available in Γ^t. Hence, for all $i \in N$, $Z_i \subseteq X_i$.

Now, we turn to the other direction, where we will show that for each Player i, $X_i \subseteq Z_i$. For this, we show that $(X_1, \ldots, X_n)$ is a set of rationalizable strategies, and hence, each X_i must belong to Z_i since Z_i is the largest collection of rationalizable set of strategies. Pick $s_i \in X_i$. By definition every strategy in X_i is not strictly dominated in the game Γ^T (where T is the final period in the IESDS procedure) with strategy sets X_i. So, by Theorem 10, every strategy in X_i is a best response among strategies in X_i to some belief μ_i over X_{-i}. So,

$$\mathcal{U}_i(s_i, \mu_i) \geq \mathcal{U}_i(\sigma_i', \mu_i) \qquad \forall\; \sigma_i' \in \Delta X_i^T.$$

Importantly, strategy profiles X_{-i} are available in each game of the IESDS procedure, and hence, μ_i is a valid belief of Player i in every period of the IESDS procedure. We will show that for all $t \in \{0, \ldots, T\}$,

$$\mathcal{U}_i(s_i, \mu_i) \geq \mathcal{U}_i(\sigma_i', \mu_i) \qquad \forall\; \sigma_i' \in \Delta X_i^t.$$

Suppose this is not true. Then, there is some period t where

$$\mathcal{U}_i(s_i, \mu_i) \geq \mathcal{U}_i(\sigma'_i, \mu_i) \qquad \forall\ \sigma'_i \in \Delta X_i^{t+1}. \tag{14}$$

but

$$\mathcal{U}_i(s_i, \mu_i) < \max_{\sigma''_i \in \Delta X_i^t} \mathcal{U}_i(\sigma''_i, \mu_i) = \mathcal{U}_i(\hat{\sigma}_i, \mu_i),$$

where $\hat{\sigma}_i$ is a best response in Γ^t for Player i with respect to belief μ_i. By the indifference lemma, every pure strategy in the support of $\hat{\sigma}_i$ is best response of Player i with respect to belief μ_i. By Theorem 10, each $\hat{s}_i$ in the support of $\hat{\sigma}_i$ is not strictly dominated in Γ^t. Hence, $\hat{\sigma}_i \in \Delta X_i^{t+1}$. But this contradicts Inequality (14).

This implies that $\mathcal{U}_i(s_i, \mu_i) \geq \mathcal{U}_i(\sigma'_i, \mu_i)$ for all $\sigma'_i \in \Delta X_i^0 = \Delta S_i$. Hence, s_i is a best response with respect to belief μ_i (with support $\times_{j \neq i} X_j$) in Γ. Hence, the collection of sets of strategies $(X_1, \ldots, X_n)$ is rationalizable.

This shows that $X_i = Z_i$ for all $i \in N$. Theorem 10 shows that $X_i = Y_i$ for all $i \in N$. ■

12 Correlated Equilibrium

Consider the mixed extension of 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 31. The story that accompanies this game is what is called the *game of chicken*. Suppose two countries have two policies: large defense investments or small defense investments. If both have small defense invesements, then they get a payoff of 6 each. If both have high defense investment, then they get a payoff of 0. If one of them has low investment and other has high investment, then the one who invests more 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 31: Game of chicken

There are three Nash equilibria of this game: $(T,R),(B,L),\left(\frac{2}{3}T+\frac{1}{3}B,\frac{2}{3}L+\frac{1}{3}R\right)$. Notice that the mixed strategy Nash equilibrium puts a probability of $\frac{1}{9}$ with which the worst possible payoff profile (B,R) will be played. 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 recommendations, players form beliefs by using Bayes' rule - they compute conditional probabilities.

In the game in Table 31, suppose the observer recommends pure strategy profiles in Nash equilibrium: (T,R) and (B,L) with probability p and $(1-p)$ respectively. 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 $p\in[0,1]$ is an equilibrium of this game. The set of payoffs that can be obtained are convex combination of $(7,2)$ and $(2,7)$. An important difference between a correlated equilibrium $\frac{1}{2}(T,R)+\frac{1}{2}(B,L)$ and the mixed strategy $(\frac{1}{2}T+\frac{1}{2}B,\frac{1}{2}R+\frac{1}{2}L)$ is that (i) the latter is not an equilibrium, and (ii) the latter produces outcome (T,L) and (B,R) with positive probability, which are not produced in the correlated equilibrium.

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.

12.1 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 an expert in the media recommending some investment strategy.

Consider the same example in Table 19. 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 32.

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

Table 32: 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 32. If Player 1 and Player 2 play mixed strategies that generates the same distribution over strategy profile as in Table 32, 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 32.

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:

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

12.2 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.

A correlated strategy p is a probability distribution over $S \equiv S_1 \times \ldots \times S_n$. We say strategy s_i of Player i is in the support of p if $\sum_{s_{-i}} p(s_i, s_{-i}) > 0$.

The definition of correlated equilibrium is the following.

Definition 18 *A correlated strategy p is a* **correlated equilibrium** *if for every $i \in N$*

and every s_i in the support of p and for every $s_i' \in S_i$, the following holds:

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

As such, it is difficult to interpret this definition. Below, we interpret via an extended game where an outside expert recommends strategies using a **public randomization device**. A public randomization device generates strategy profiles with pre-specified probabilities, and these probabilities are common knowledge amonth players. The equilibrium of this extended game will correspond to the definition of correlated equilibrium in Definition 18.

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 $\Gamma(p)$. Consider the battle of sexes game in Table 33. To remind, the game (mixed extension) has exactly three Nash equilibrium: $(A, a), (B, b)$, and $(\frac{2}{3}A + \frac{1}{3}B, \frac{2}{3}a + \frac{1}{3}b)$.

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

Table 33: Correlated equilibria of battle of sexes

Now, a correlated equilibrium is described by a probability distribution over pure strategy profiles: $p(A, a), p(A, b), p(B, a), p(B, b)$. The extended game is shown in Figure 8. You can think that the observer makes the first move of this game by announcing a recommendation - which you can think of as a "state". So, each recommendation, and hence, each pure strategy profile, defines a state. Once the state is defined, players can take pure strategies of the game Γ. Players are only given partial information about the state - this is because they only know about their own recommendation. This is shown in Figure 8 by grouping pairs of states. The RED groups of states are for Player 1 - here, Player 1 receives the same recommendation and he cannot distinguish between which of these states have occurred. Similarly, the BLUE groups of states are for Player 2.

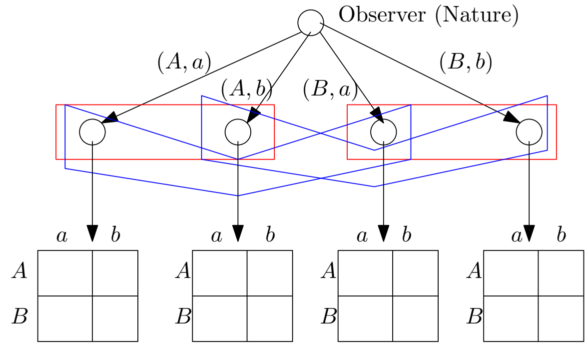

Figure 8: Representation of extended game

How do we analyze equilibria of such games? First step is Players are *Bayesian rational*. This means that given the probability distribution with which the Observer (Nature) draws the states, Players use Bayes' rule to compute their probability with which they are in each state. So, Player 1 after observing A, believes that the probability of state (A, a) is $\frac{p(A,a)}{p(A,a)+p(A,b)}$.

Formally a strategy in this extended game is a different object compared to the strategy in a strategic form game.

DEFINITION **19** *A strategy of Player i in the extended game $\Gamma(p)$ is a map $\psi_i : S_i \rightarrow S_i$, i.e., specifies an action for every possible recommended action.*

For instance, consider the strategy, which we call the **obedient** strategy - for every $s_i \in S_i$, $\psi_i^*(s_i) = s_i$ for each i. We are interested in studying obedient strategies to be an "equilibrium" of this extended game $\Gamma(p)$. An equilibrium in this game requires two things: (1) Bayesian rationality; and (2) Best response play, called *sequential rationality*. Once Player i receives a recommendation s_i, he forms belief about the state. Given this belief obedient strategy must maximize her payoff given that everyone else is obedient. In this case, we call ψ^* an **equilibrium** of the extended game $\Gamma(p)$.

We consider both the requirements in turn.

Bayesian rationality. If Player i receives recommendation s_i, then his conditional belief that state is (s_i, s_{-i})

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

where the denominator is positive from the fact that s_i is in the support of p. Bayesian rationality requires that players compute beliefs as above (using Bayes' rule).

Sequential rationality. Given the belief of Player i, his expected payoff from following $\psi_i^*(s_i) = s_i$ (given others are obedient) 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 obedient) 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 has to hold for all s_i which can come as a recommendation for player i, i.e., in the support of p and for all $s_i' \in S_i$.

This equilibrium condition of $\Gamma(p)$ leads to the exact same definition of a correlated equilibrium.

Definition **20** *A correlated strategy p over S is a* **correlated equilibrium** *if for every $i \in N$, for every s_i in the support of p and every $s_i' \in S_i$,*

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

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

Note that in the above definition if s_i is *not* in the support of p, then $p(s_i, s_{-i}) = 0$ for all s_{-i}. Hence, the inequality holds trivially. This means that p is a correlated equilibrium if

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

$$\sum_{s \in S_1 \times \ldots \times S_n} p(s) = 1$$

$$p(s) \geq 0 \qquad \forall\ s \in S_1 \times \ldots \times S_n.$$

The above sysstem is a system of linear inequalities. In fact, they describe a bounded system of linear inequalities. Hence, it is a *polytope*, a subset of $\mathbb{R}^{|S_1 \times \ldots \times S_n|}$ described by hyperplanes. The set of correlated equilibria is a polytope in this space. Such a polytope is always non-empty since Nash equilibria are in it as we discuss next.

12.3 Correlated and Nash equilibrium

Consider a correlated equilibrium for the game in Table 33 which has every strategy profile in its support:

$$p(A, a), p(A, b), p(B, a), p(B, b).$$

According to the definition, the following are the inequalities to be satisfied for the game in Table 33. Consider Player 1. If she receives A, then her expected payoff from playing A is

$$\frac{1}{p(A,a) + p(A,b)} \Big[1 \times p(A,a) + 0 \times p(A,b)\Big] = \frac{p(A,a)}{p(A,a) + p(A,b)}.$$

But her expected payoff from playing B is

$$\frac{1}{p(A,a) + p(A,b)} \Big[0 \times p(A,a) + 2 \times p(A,b)\Big] = \frac{2p(A,b)}{p(A,a) + p(A,b)}.$$

Hence, for this to be a correlated equilibrium, we will need

$$p(A, a) \geq 2p(A, b).$$

A similar calculation reveals that if a player receives recommendation B, then to play strategy B, we should have $2p(B,b) \geq p(B,a)$. Using similar calculations for Player 2, we conclude that if the following inequalities have a solution (in the interior), then p defines a correlated equilibrium.

$$\begin{aligned} p(A,a) &\geq 2p(A,b) \\ 2p(B,b) &\geq p(B,a) \\ 2p(A,a) &\geq p(B,a) \\ p(B,b) &\geq 2p(A,b). \end{aligned}$$

Hence, we get the following inequalities:

$$\begin{aligned} p(A,a) &\geq \max(2p(A,b), \frac{1}{2}p(B,a)) \\ p(B,b) &\geq \max(2p(A,b), \frac{1}{2}p(B,a)). \end{aligned}$$

In general, this is a linear system of inequalities. They can be described by their "extreme points". To get one of the extreme points, we can set $2p(A,b) = \frac{1}{2}p(B,a) = p(A,a) = p(B,b)$ and using the fact their sum is 1, we get

$$p(A,b) = \frac{1}{9}, p(B,b) = \frac{4}{9}, p(A,a) = p(B,b) = \frac{2}{9}.$$

Notice that this is the mixed strategy Nash equilibrium of the game. This generates a payoff of $\frac{2}{3}$ for each player. As it turns out, the set of all payoffs that can be achieved by correlated equilibrium is the convex hull of Nash equilibrium payoffs: $(1,2), (2,1), (\frac{2}{3}, \frac{2}{3})$.

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).

This is in contrast to a mixed strategy Nash equilibrium. Remember, a mixed strategy Nash equilibrium is a pair $(pA + (1-p)B, qa + (1-q)b)$ which satisfies for player 1 the following inequality:

$$u(A,a)q + u(A,b)(1-q) \geq u(A,a)p'q + u(A,b)p'(1-q) \\ + u(B,a)(1-p')q + u(B,b)(1-p')(1-q) \ \forall \ p' \in [0,1].$$

This inequality ensures that A is a best response to $qa + (1-q)b$. This is surely not a *linear* system – on the RHS, p' is multiplied with q terms and this has to hold for all p'. This is a much more complicated system – to be precise, as we have shown, finding Nash equilibria is equivalent to finding fixed points of the best response correspondence, which is a complicated problem computationally.

However, the notion of correlated equilibrium extends the idea of Nash equilibrium. The following discussion formalizes that. 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. This also shows that a correlated equilibrium always exists.

THEOREM **12** *For every Nash equilibrium σ^* of Γ, the induced correlated strategy p_{σ^*} is a correlated equilibrium.*

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, $s_i, s_i' \in S_i$ such that s_i is in the support of σ^*. We see that

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

Further,

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

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. ■

A simple corollary of Theorem 12 is that a correlated equilibrium always exists in a finte game.

COROLLARY **1** *Every finite game has correlated equilibrium.*

Proof: By Nash, every finite game has a mixed strategy Nash equilibrium. Such a mixed strategy Nash equilibrium is also a correlated equilibrium due to Theorem 12 ■

13 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 how much to invest in a project. Each agent privately observes a signal about the quality of the project. Based on the signal, an agent has a set of *actions* available to him, which can be thought of as the level investment. His utility will depend on the signals about the project and the actions chosen by both the agents. Here, signals are privately observed by the players. 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 quality of the project. Each agent uses a *common prior* to evaluate uncertainty using expected utility. In this example, there is a probability distribution about the quality of the project. Note that since an agent only observes his own signal, 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 **21** *A* **Bayesian game (game of incomplete information)** *is defined by*

- N*: a finite set of players,*
- T_i*: set of* **types** *(signals) for each player* i*, and* $T = \times_{i \in N} T_i$ *is the set of type vectors,*
- p*: a common probability distribution (***belief or prior***) over* T *with the restriction that* $\pi_i(t_i) := \sum_{t_{-i} \in T_{-i}} p(t_i, t_{-i}) > 0$ *for each* $t_i \in T_i$ *and for each* $i \in N$*,*
- $A_i(t_i)$*: the set of actions available to each Player* i *with type* t_i*,*
- $u_i(t, a)$*: the payoff assigned by each Player* i *at type profile* $t \equiv (t_1, \ldots, t_n) \in T$ *when action profile* $a \equiv (a_1, \ldots, a_n)$*, where* $a_j \in A_j(t_j)$ *for all* $j \in N$*, is played.*

A Bayesian game 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 a payoff equal to $u_i(t, a)$ when the type profile realized is $t \equiv (t_1, \ldots, t_n)$ and the action profile is $a \equiv (a_1, \ldots, a_n)$.

Figure 9 illustrates a Bayesian game for two players with the type set of Player 1 being $\{t_1, \hat{t}_1\}$ and that of Player 2 being $\{t_2, \hat{t}_2\}$. As the Figure shows, a Bayesian game can be described by a sequence of moves, where the first move is by Nature determining the type vector of players. Figure 9 shows the four possible type vectors. Once the type vectors are realized Players know the actions available to them (but not to others as they do not know the types of others). Hence, there is still uncertainty about the game being played.

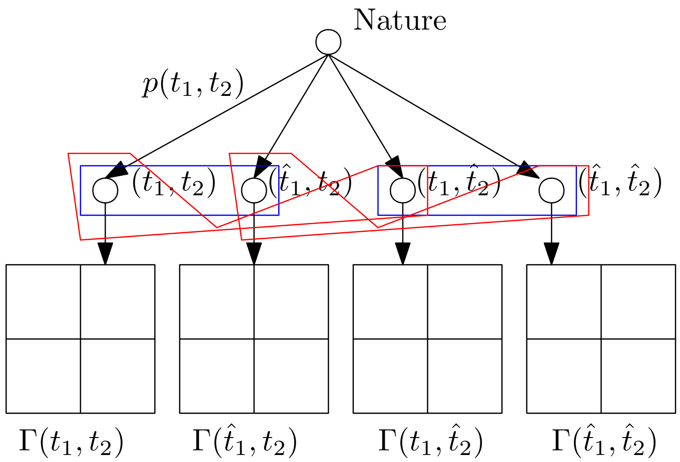

Figure 9: A Bayesian game

In most of the examples, we will make the assumption that for all t_{-i}, t'_{-i},

$$u_i((t_i, t_{-i}), a) = u_i((t_i, t'_{-i}), a) \text{ for all } a, \text{ for all } t_i, \text{ for all } i \in N$$

This is called a **private values** model. It rules out the possibility that a player's utility depends directly on the type of other players. Notice that the action chosen by a Player may depend on his type in the game, and hence, indirectly, Player i's utility will depend on the type of other players (through the actions chosen by other players).

13.1 A simple example: market for lemons

Consider a used car market in which a car is in **good** condition with probability $q \in [0, 1]$ and and in **bad** condition with probability $1 - q$ (call such cars *lemons*). There is one buyer and one seller. The seller knows whether the the car is good or bad but the buyer does not know the quality of the car. There is a market price for every used car, denoted by p. The probability q is common knowledge.

The seller has two possible actions: SELL and NOT SELL. The buyer has two possible actions: BUY and NOT BUY. If the car is **good**, then the buyer enjoys a value of 6 and the seller enjoys a value of 5. If the car is **bad**, the buyer enjoys a value of 4 and the seller has zero value for it. Notice that in both the *states*, it is "efficient" to trade the car.

What is the Bayesian game? First, the Nature informs the seller (and not the buyer) if the car is good or bad. If the car is good, then the game in Table 34 is played. If the car is bad, then the game in Table 35 is played. In this game, at the *interim* stage (i.e., the stage just after the nature moves), the seller knows the complete state of the world and hence, knows which game will be played. On the other hand, the buyer does not know which game will be played.

	Sell	Not sell
Buy	$6-p,p$	0,5
Not buy	0,5	0,5

Table 34: Good car

	Sell	Not sell
Buy	$4-p,p$	0,0
Not buy	0,0	0,0

Table 35: Bad car

In another words, we can interpret this as saying that the seller has two types: good and bad. The buyer has only one type (i.e., no uncertainty from the buyer side). So, depending on the type of the seller, we play two different strategic-form games. The distribution of type is: probability of good type is q and bad type is $(1-q)$.

13.2 Strategy and Utility

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 functions. To remind, a strategy must describe the action to be taken for every possible contingency. Hence, a strategy must describe what action to take for every signal/type that the player receives.

A **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 strategy prescribes one action for every type.

In the example in Section 13.1, the strategy of a seller is a map

$$s_\ell : \{\mathbf{good}, \mathbf{bad}\} \to \{\textsc{sell},\ \textsc{not sell}\}.$$

Since the buyer has no type in this example, its strategy is

$$s_b \in \{\textsc{buy},\ \textsc{not buy}\}.$$

What is the payoff of Player i from a strategy profile $s \equiv (s_1, \ldots, s_n)$? 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 s is

$$U_i(s) := \sum_{t \in T} p(t) u_i(t, (s_1(t_1), \ldots, s_n(t_n))).$$

Here, if type profile t is realized, then action profile $(s_1(t_1), \ldots, s_n(t_n))$ is played according to the strategy profile s. Hence, the payoff realized by Player i at type profile t is just $u_i(t, (s_1(t_1), \ldots, s_n(t_n)))$. Then, $U_i(s)$ computed using expectation from this.

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 $t_{-i} \in T_{-i}$,

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

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

$$U_i((a_i, s_{-i})|t_i) := \sum_{t'_{-i} \in T_{-i}} p_i(t'_{-i}|t_i) u_i((t_i, t'_{-i}), (a_i, s_{-i}(t'_{-i}))).$$

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.

An easy consequence of this definition is the following. Consider Player i and a strategy profile (s_i, s_{-i})

$$\sum_{t_i \in T_i} U_i((s_i(t_i), s_{-i})|t_i)\pi_i(t_i) = \sum_{t_i \in T_i} \pi_i(t_i) \sum_{t_{-i} \in T_{-i}} p_i(t_{-i}|t_i) u_i(t, (s_i(t_i), s_{-i}(t_{-i})))$$
$$= \sum_{t \in T} p(t) u_i(t, (s_i(t_i), s_{-i}(t_{-i}))) = U_i(s). \quad (15)$$

Note: The above expressions are for finite type spaces, but similar expressions (using integrals) can also be written with infinite type spaces.

13.3 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 **22** *A strategy profile s^* is a* **Nash equilibrium** *in a Bayesian game if for each player i and each pure strategy s_i,*

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

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

$$U_i((s_i^*(t_i), s_{-i}^*)|t_i) \geq U_i((a_i, s_{-i}^*)|t_i)$$

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

Consider the example in 13.1. We find prices for which the following strategy

$$s_\ell^*(good) = \text{not sell}, s_\ell^*(bad) = \text{sell}, s_b^* = \text{buy}$$

is a Nash equilibrium and a Bayesian equilibrium if there is equal probability ($\frac{1}{2}$) of being good and bad. In that case, the expected payoff of the buyer from this strategy is

$$U_b(s^*) = \frac{1}{2} \times 0 + \frac{1}{2} \times (4 - p).$$

Expected payoff of the buyer from the strategy where the buyer does not buy is zero. So, for this to be an equilibrium, we will need $4 - p \geq 0$ or $p \leq 4$.

The seller has three more strategies:

- s_ℓ^1: she sells in both states,
- s_ℓ^2: she does not sell in both states,
- s_ℓ^3: she sells in good and does not sell in bad.

The expected payoffs from each of these strategies are:

$$\begin{aligned} U_\ell(s_\ell^*, s_b^*) &= \frac{1}{2}(5 + p) \\ U_\ell(s_\ell^1, s_b^*) &= p \\ U_\ell(s_\ell^2, s_b^*) &= \frac{5}{2} \\ U_\ell(s_\ell^3, s_b^*) &= \frac{p}{2} \end{aligned}$$

Hence, for s^* to be a Nash equilibrium, we will need $\frac{1}{2}(5 + p) \geq p$ or $p \leq 5$. So, s^* is a Nash equilibrium if $p \leq 4$.

For Bayesian equilibrium, the equilibrium condition for buyer remains the same: $p \leq 4$. For the seller, she needs to compute payoffs at interim stages given that the buyer follows s_b^*. Suppose the car is good. In that case, following s^* (not selling) gives her a payoff of 5. By selling she gets a payoff of p. Hence, Bayesian equilibrium will require $5 \geq p$. Suppose the car is bad. In that case, following s^* (selling) gives her a payoff of p. By not selling she gets a payoff of 0. Hence, following s^* is optimal when the car is bad. Overall for s^* to be a Bayesian equilibrium, we will need: $p \leq 4$, which is the same condition as s^* being a Nash equilibrium. As we will show next, this is no coincidence.

The first property that we show is that (with finite type spaces) 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 **13** *Suppose type space of each player is finite. A strategy profile is a Bayesian equilibrium if and only if it is a Nash equilibrium.*

Proof: Consider a strategy profile s^*. Suppose s^* 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

$$U_i((s^*_i(t_i), s^*_{-i})|t_i) \geq U_i((a_i, s^*_{-i})|t_i).$$

For any strategy $s_i : T_i \to \cup_{t_i} A_i(t_i)$ with $s_i(t_i) \in A_i(t_i)$ for all t_i, we know from Equality 15 that

$$U_i(s_i, s^*_{-i}) = \sum_{t_i \in T_i} \pi_i(t_i) U_i((s_i(t_i), s^*_{-i})|t_i) \leq \sum_{t_i \in T_i} \pi_i(t_i) U_i(s^*_i(t_i), s^*_{-i}|t_i) = U_i(s^*_i, s^*_{-i}).$$

Hence, s^* is a Nash equilibrium.

Now, suppose that s^* is a Nash equilibrium. Assume for contradiction that s^* 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

$$U_i((a_i, s^*_{-i})|t_i) > U_i((s^*_i(t_i), s^*_{-i})|t_i). \tag{16}$$

Now, construct a new strategy s_i such that $s_i(t_i) = a_i$ but $s_i(t'_i) = s^*_i(t'_i)$ for all $t'_i \neq t_i$.

Now, observe the following:

$$\begin{aligned}
U_i((s_i, s^*_{-i})) &= \sum_{t'_i \neq t_i} \pi_i(t'_i) U_i((s_i(t'_i), s^*_{-i})|t'_i) + \pi_i(t_i) U_i((s_i(t_i), s^*_{-i})|t_i) \\
&= \sum_{t'_i \neq t_i} \pi_i(t'_i) U_i((s^*_i(t'_i), s^*_{-i})|t'_i) + \pi_i(t_i) U_i((a_i, s^*_{-i})|t_i) \\
&> \sum_{t'_i \neq t_i} \pi_i(t'_i) U_i((s^*_i(t'_i), s^*_{-i})|t'_i) + \pi_i(t_i) U_i((s^*_i(t_i), s^*_{-i})|t_i) \\
&= U_i(s^*_i, s^*_{-i}),
\end{aligned}$$

where the strict inequality followed from Inequality 16 and the fact that $\pi_i(t_i) > 0$ for all i and for all t_i. This contradicts the fact that s^* is a Nash equilibrium. ■

The equivalence result needs type spaces to be finite. In general, we will consider Bayesian games where type space is not finite. In such games a Bayesian equilibrium will continue to imply a Nash equilibrium but the converse need not hold. So, we will use the solution concept Bayesian equilibrium in all the Bayesian games that we analyze.

But do all Bayesian games admit a Bayesian equilibrium? Which Bayesian games admit a Bayesian equilibrium? There is a long literature on this topic, which we will skip. Just like Nash equilibrium, there is a well-behaved class of games that admit a Bayesian equilibrium. One simple way to think of an existence result is to allow for *mixed actions* for every type. Essentially, we enrich the action space but keep the finite nature of type space. That is, after every t_i, the set of actions available to a player is $\Delta A_i(t_i)$, where $A_i(t_i)$ is finite. This will be the analogue of the mixed 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 $t_i \in T_i$, $\sigma_i(t_i) \in \Delta A_i(t_i)$. The utility of player i from such a mixed strategy will be evaluated by taking expectation. A mixed strategy Bayesian equilibrium always exists if action spaces are finite and type spaces are finite - a result which we will not prove.

14 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. The auction that we will study will involve zero payments by the agent who does not win the object. 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 $f(x) \neq 0$ 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.

A random variable that will come handy is the highest of $(n-1)$ values: we denote it by G. In particular, we will be interested to know what is the probability that $(n-1)$ bidders have *value* less than or equal to x: this is precisely

$$\text{Prob}\Big[n-1 \text{ bidders have value less than } x\Big] = G(x) = [F(x)]^{n-1}.$$

First-price auction. The first-price auction is probably the most popular format of 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 $W_j(b)$ for each buyer $j \in N$. Note that $W_j(b) = 1$ if $b_j > \max_{k \neq j} b_k$ and $W_j(b) = 0$ 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

$$W_j(b)\big[x_j - b_j\big]$$

Unlike the Vickrey auction, the first-price auction has no weakly dominant strategy (verify). To see this, note that a bidder who bids her valuation as bid will get a payoff of zero when she wins the object. Hence, her expected payoff from bidding her valuation is zero. Clearly, bidding slightly less than her valuation generates higher expected payoff if others are bidding truthfully. So, bidding your valuation is no longer a weakly dominant strategy. 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 action (bid) when his value of bidder i is x must be $\beta(x)$. In particular, bidder i is comparing his expected payoff by bidding $\beta(x)$ and some other bid b. We argue that bidder i need not worry about bidding more than $\beta(w)$. This is because other bidders, who are following β, will never bid more than $\beta(w)$. Hence, bidding more than $\beta(w)$ ensures that bidder i always wins. If $b > \beta(w)$, for some small enough $\delta > 0$, bidding $b - \delta$ also ensures bidder i wins for sure and pays less. So bidding strictly above $\beta(w)$ cannot be optimal. Hence, any *optimal* bid must be between 0 and $\beta(w)$. For this reason, we will assume that $b \in [0, \beta(w)]$. In other words, to establish the strategy profile β is a Bayesian equilibrium, it is enough to check that each bidder i does not gain by taking action $b \in [0, \beta(w)]$ for each of his type.

An advantage of this simplification is that for any $b \in [0, \beta(w)]$, there is a unique value $z \in [0, w]$ such that $b = \beta(z)$: uniqueness follows because β is strictly increasing. This z need not be the actual value of bidder i. Mathematically, $z = \beta^{-1}(b)$.

Note that if bidder i with value x bids $\beta(x)$, and since everyone else is using $\beta(\cdot)$ strategy, increasing β ensures that the probability of winning for bidder i is equal to the probability that x is greater than the highest value among $n - 1$ other bidders, which in turn is equal to $G(x)$.

Similarly, suppose a bidder bids b and other bidders follow a symmetric strategy β. Then this bidder wins if each other bidder bids less than b. That is i wins if $b > \beta(v_j)$ for each $j \neq i$. Equivalently, bidder i wins if $v_j < \beta^{-1}(b)$ for each $j \neq i$. Probability that bidder

j's value is less than b is given by $F(\beta^{-1}(b))$. By independence, probability that all $n-1$ bidders except i has value less than b is given by

$$G(\beta^{-1}(b)) = \Big[F(\beta^{-1}(b))\Big]^{n-1}.$$

Hence, the notion of Bayesian equilirbium reduces to the following definition.

DEFINITION **24** *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 $v \in [0,w]$*

$$G(v)(v-\beta(v)) \geq G(\beta^{-1}(b))\ (v-b) \qquad \forall\ b \in [0,\beta(w)]$$

THEOREM **14** *A symmetric equilibrium in a first-price auction is given by*

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

REMARK. The interpretation of this bid function is as follows. A bidder with type/value x for the object bids an amount equal to his *conditional* expectation of the highest value of other bidders, where the conditioning is done on the fact that he has the highest value.

Proof: Suppose every bidder except bidder j follows the suggested strategy. The suggested strategy generates non-negative payoff. Let bidder j bid $b \in [0, \beta^I(w)]$. Hence, any bid b can be mapped to a z such that $\beta^I(z) = b$. Then the expected payoff from bidding $\beta^I(z) = b$ when his true value is v is

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

Differentiating this expression with respect to z, we get

$$g(z)(v-z),$$

which is positive for $z < v$, negative for $z > v$, and zero at $z = v$. Hence, the payoff is maximized at $z = v$. Hence, bidding according to $\beta^I(\cdot)$ is a symmetric equilibrium. ■

Uniqueness. We now prove that this is the unique symmetric equilibrium in the first-price auction. Suppose there is a symmetric equilibrium β in the first-price auction. Now, consider any bidder j. Assume that he realizes a true value x, and wants to determine his optimal bid value b. In equilibrium, $b = \beta(x)$. We have already argued that an optimal bid lies between 0 and $\beta(w)$. Hence, any bid b is equivalent to "pretending to be" a buyer of type z and bidding $\beta(z)$. Thus, a buyer of type x gets the following payoff by pretending to be of type z:

$$G(z)(x - \beta(z)),$$

where $G(z)$ is its probability of winning given that others are following β. Clearly, $\beta(0) = 0$. A necessary condition for maximum is the first order condition, which is obtained by differentiating with respect to z.

$$g(z)(x - \beta(z)) = G(z)\beta'(z).$$

If β is an equilibrium, then this first order condition must hold for all x with bid at x equal to $\beta(x)$:

$$\begin{aligned} g(x)(x - \beta(x)) &= G(x)\beta'(x) \qquad \forall\, x \in [0, w] \\ \text{or } xg(x) &= \frac{d(G(x)\beta(x))}{dx} \qquad \forall\, x \in [0, w] \end{aligned}$$

To find a solution to the differential equation (using $\beta(0) = 0$):

$$\int_0^x yg(y)dy = G(x)\beta(x)$$

Hence, the following is the unique symmetric Bayesian equilibrium of the first-price auction:

$$\beta(x) = \frac{1}{G(x)} \int_0^x yg(y)dy.$$

This is the same symmetric equilibrium we had derived in Theorem 14.

15 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 $p_b \leq p_s$, 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 **15** *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 $\frac{11}{12}$ as a best response - this is because any price strictly above $\frac{11}{12}$ can be improved by lowering it a little further, and hence, cannot be a 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 $\pi_b \leq \frac{11}{12}$, we have $\frac{3}{2}\pi_b - \frac{3}{8} \leq 1$.

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 of the unconstrained problem. 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, $\pi_b = \frac{2}{3}v + \frac{1}{12}$. Note that $\pi_b \in [\frac{1}{12}, \frac{9}{12}]$ satisfies our constraint. Hence, it is a best response to p_s^* strategy of the seller.

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 $v > \frac{3}{2}\pi_c - \frac{1}{8} \geq 0$ since $\pi_c \geq \frac{1}{12}$. Further, $\frac{3}{2}\pi_c - \frac{1}{8} \leq 1$ since $\pi_c \leq \frac{3}{4}$. Denote $x_c = \frac{3}{2}\pi_c - \frac{1}{8}$. Hence, the expected payoff of the seller at type c is given by

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

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

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

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

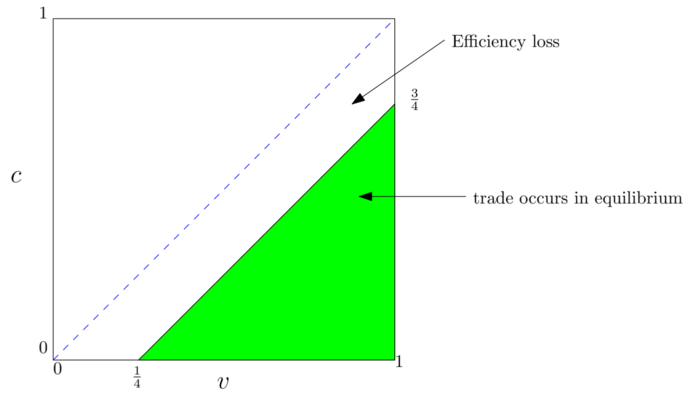

Figure 10: Efficiency loss in bilateral trade with incomplete information

shown to be unique in the class of strategies where players use strategies linear in their type. One 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 $v > c$ - such trades will be possible if there is complete information. Hence, there is some loss in efficiency due to incomplete information. It can be shown that this is a general impossibility. You cannot construct any Bayesian game whose equilibrium will be efficient in this model (more on this in some advanced course). The region of trade in this particular equilibrium and efficiency loss is shown in Figure 10.

16 Repeated Games

16.1 Basic Ideas - The Repeated Prisoner's Dilemma

Consider the Prisoners' Dilemma (PD) game in Table 36. 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	1,1	5,0
R_1	0,5	4,4

Table 36: 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.

We can think of reduced strategic form of this game. In this reduced form, Player $i \in \{1, 2\}$ has a complex strategy. First, she needs to choose an action for Stage 1. Second, she needs to choose an action for every observed action profile of Stage 1 for Stage 2. For instance, if she has observed, (L_1, R_2) being played in Stage 1, her Stage 2 choice of an action can be contingent on that. This leads to a very complex strategy structure of the game in *reduced form*.

Instead of looking at the reduced form, we can also analyze the game *backwards*. In particular, suppose we require that starting at every period, players must play Nash equilibrium of the reduced form from that period onwards. Call this a *subgame perfect equilibrium*. Since the unique Nash equilibrium (weakly dominant strategy 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, we can add $(1, 1)$ to the payoff matrix in the first stage, and then compute a Nash equilibrium of the overall game. 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) in each period (i.e., in period 1 players play (L_1, L_2) and in period 2 players play (L_1, L_2) irrespective of what they observed in period 1).

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 **25** *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 37. There are two Nash equilibria of this game: (L_1, L_2) and (R_1, R_2).

	L_2	M_2	R_2
L_1	1,1	5,0	0,0
M_1	0,5	4,4	0,0
R_1	0,0	0,0	3,3

Table 37: A Game with Multiple Nash Equilibrium

Now, suppose the stage game in Table 37 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 38, 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 38: (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_2	M_2	R_2
L_1	2,2	6,1	1,1
M_1	1,6	7,7	1,1
R_1	1,1	1,1	4,4

Table 38: 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.

16.2 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. The collection of all such histories in period t will be denoted as H^t.

An **infinitely repeated game** of G is defined by $G^\infty \equiv (G, H, \{u_i^*\}_{i \in N})$, where

- $H = \cup_{t=1}^{\infty} H^t$ are the set of all possible histories, with $H^1 \equiv \emptyset$ denoting the null history, H^t denoting the possible histories till period t, and H^∞ denoting all infinite length histories.

- $u_i^* : H^\infty \to \mathbb{R}_+$ 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.

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

$$s_i^t : H^t \to A_i.$$

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. A player may not care about all possible histories in a period but only about some "states" that may occur.

We give an example showing how a strategy in Prisoner's Dilemma can be modelled 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 $q_i^0 := \mathcal{C}$ for each i. Now, $f_i(\mathcal{C}) = C$ and $f_i(\mathcal{D}) = D$. The transition function looks as follows for each i:

$$\tau_i(\mathcal{C}, (C, C)) = \mathcal{C}, \tau_i(\mathcal{X}, (X, Y)) = \mathcal{D} \text{ if } (\mathcal{X}, (X, Y)) \neq (\mathcal{C}, (C, C)).$$

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 $\{w_i^t\}_{t=1}^{\infty}$.

Hence, to properly compare outcomes of two strategy profiles, players need to have preference over infinite utility streams. There are many ways to compare infinite utility streams.

We give some example. Take any two infinite utility streams $v \equiv \{v^t\}_{t=1}^{\infty}$ and $w \equiv \{w^t\}_{t=1}^{\infty}$. We will write $v \succ w$ whenever we want to say v is strictly preferred to w.

- **Utilitarianism.** There is some integer $N \geq 1$ such that $v \succ w$ if and only if $\sum_{t=1}^{N} v^t > \sum_{t=1}^{N} w^t$.
- **Limit of means.** There exists $\epsilon > 0$ such that $v \succ w$ if and only if $\frac{1}{T}\sum_{t=1}^{T} v^t - \frac{1}{T}\sum_{t=1}^{T} w^t > 0$.
- **Overtaking.** There exists $\epsilon > 0$ such that $v \succ w$ (strict relation) if and only if $\sum_{t=1}^{T} v^t - \sum_{t=1}^{T} w^t > \epsilon$ holds for all but finite number of T.

The first criteria is quite simple but it naturally ignores some future payoffs. The other two criteria are quite incomplete. 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 $\{v_i^t\}_{t=1}^{\infty}$. Here, δ can either be interpreted as the discount factor of agents or the probability with which the game continues to the next period. 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 different 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 $\delta > \frac{1}{2}$.

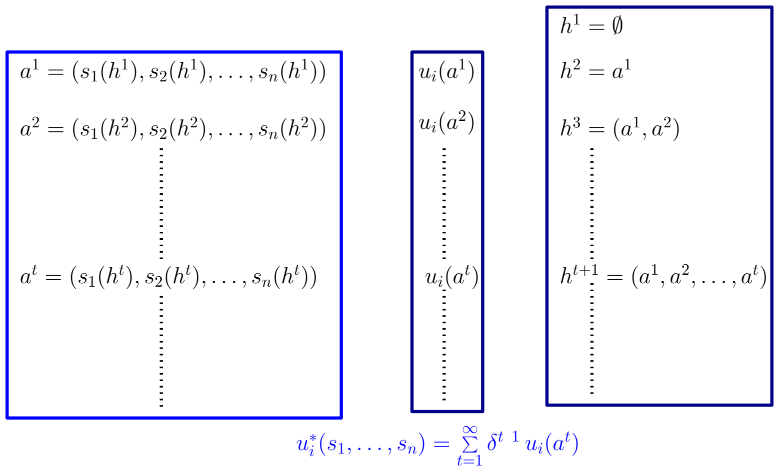

Figure 11: Strategies and discounted payoffs

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 plays $a_i^t \equiv s_i^t(h^t)$. From this, we can compute the utility of Player i as

$$u_i^*(s) := (1-\delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(a^t).$$

Figure 11 gives a pictorial description of how payoff of a strategy profile is calculated with discounting.

Having defined strategies and payoffs, we are now ready to define the equilibrium concepts for repeated games. For this, we refer to the original infinitely repeated game G^∞ and the infinitely repeated game starting at any arbitrary period t and history h^t as G^{∞}_{t,h^t}. For every t and every $h_t \in H^t$, the infinitely repeated game G^{∞}_{t,h^t} is referred to as a **subgame** of G^{∞}. Of course, $G^{\infty} = G^{\infty}_{1,\emptyset}$ and G^{∞} is a subgame of itself.

DEFINITION **26** *A strategy profile $s \equiv (s_1, \ldots, s_n)$ is a* **Nash equilibrium** *of the infinitely repeated game G^{∞} if for every $i \in N$, for every s'_i, we have*

$$u^*_i(s_i, s_{-i}) \geq u^*_i(s'_i, s_{-i}).$$

16.3 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 39. 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 39: Prisoner's Dilemma

PROPOSITION **2** *Suppose $\delta \geq \frac{1}{2}$. Then, there is a Nash equilibrium in the Prisoner's Dilemma game (Table 39), 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. Following L_1 gives Player

1 a payoff of 1. Playing R_1 in some period 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 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. Thus, the specified strategy is a Nash equilibrium. ■

16.4 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. This is the minimum payoff player i can be held to by its opponents (using pure actions), given that he plays best response to the action profile a_{-i}. Let $u_i(\underline{a}_i, \underline{a}_{-i}) = \underline{v}_i$ for player i. Then, we call $\underline{a}^i = (\underline{a}_i, \underline{a}_{-i})$ the **minmax action profile** against player i. Notice that this includes an action for Player i also.

As an example, consider the game in Table 40. Consider Player 1 (the game is symmetric, so calculations for Player 2 is similar). When Player 2 plays L_2, maximum Player 1 can get is 2. When she plays M_2, Player 1 can guarantee 7, and when she plays R_2, Player 1 can guarantee 4. Hence, the minmax payoff of Player 1 is 2. Notice that the max-min payoff is calculated differently. The idea there is that the opponents are punishing a player and what is the best a player can do with such opponents. So, here when Player 1 plays L_1, opponent can punish him by playing R_2, which gives him 1. Similarly, for M_1 and R_1, the payoffs are 1 and 1 respectively too. So, max-min payoff of Player 1 is only 1. Thus, the order of optimization of the player and its opponents are different in max-min and min-max calculations.

	L_2	M_2	R_2
L_1	2,2	6,1	1,1
M_1	1,6	7,7	1,1
R_1	1,1	1,1	4,4

Table 40: Minmax payoffs

The reason minmax values are important is the following lemma.

LEMMA **9** *Let* $s \equiv (s_1, \ldots, s_n)$ *be a Nash equilibrium of* G^∞*. Then, for every* $i \in N$*, we have* $u_i^*(s) \geq \underline{v}_i$.

Proof: Fix some period t and some history h^t. Suppose other players choose actions a_{-i} according to s_{-i}. Let a_i be a best response of Player i to some a_{-i} in the stage game. Then,

$$u_i(a_i, a_{-i}) = \max_{a'_i} u_i(a'_i, a_{-i}) \geq \min_{a'_{-i}} \max_{a'_i} u_i(a'_i, a'_{-i}) = \underline{v}_i.$$

Hence, Player i can guarantee at least $\underline{v}_i$ in the stage game every period.

Then, consider the strategy s'_i which assigns the best response in the stage game to s_{-i} at every period and at every history. We just showed that $u_i^*(s'_i, s_{-i}) \geq \underline{v}_i$. Hence, if s is a Nash equilibrium, $u_i^*(s_i, s_{-i}) \geq u_i^*(s'_i, s_{-i}) \geq \underline{v}_i$. ■

Hence, Player i is guaranteed to get at least $\underline{v}_i$ payoff in any pure action Nash equilibrium of the repeated game. In particular, Lemma 9 asserts that any strategy profile that does not ensure a payoff of $\underline{v}_i$ for Player i has an easy deviation for Player i: strategy where he best responds to other players' actions in each period. So, we have no hope of sustaining an equilibrium where some Player i gets less than $\underline{v}_i$. The message of the folk theorems will be that this condition is *almost* necessary and sufficient for an equilibrium payoff.

DEFINITION **27** *A payoff profile* $v = (v_1, \ldots, v_n)$ *is* **strictly enforceable** *if for every* $i \in N$*, we have* $v_i > \underline{v}_i$.

Figure 12 gives a pictorial description of the strictly enforeable payoffs for the game in Table 40. The folk theorem we prove next says that all pure action profiles in the strictly

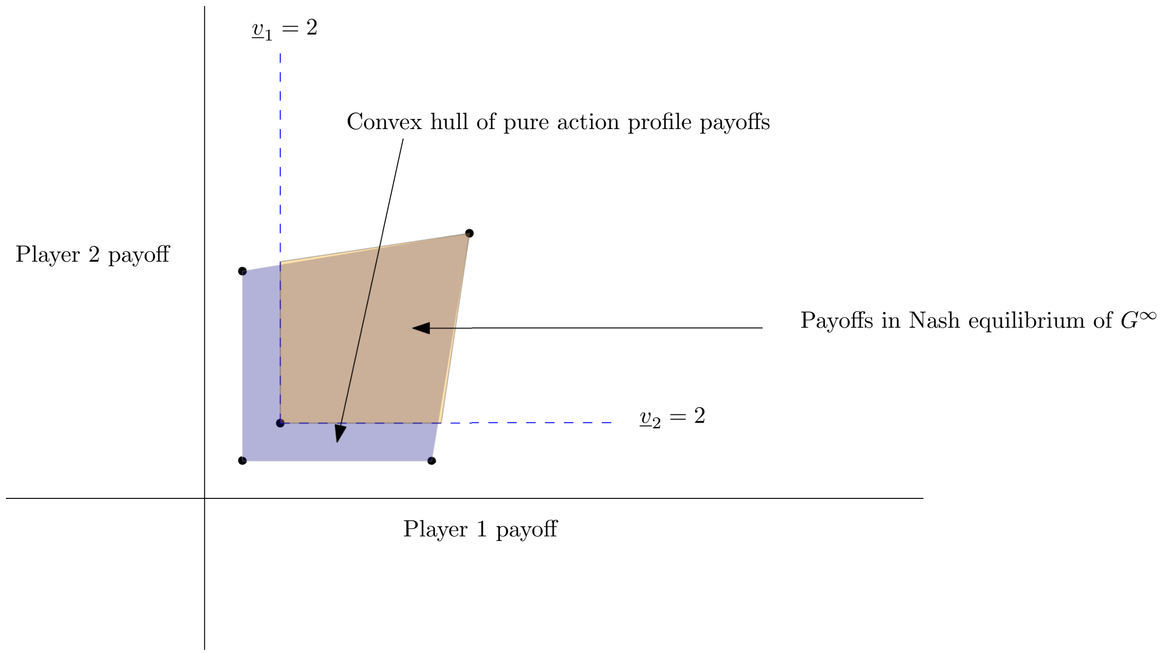

Figure 12: Strictly enforceable payoff profiles

enforceable region can be outcome of Nash equilibrium play in infinite repeatation of a stage game. Hence, a significant portion of the payoff profile can be sustained as equilibrium.

Theorem **16 (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 $\underline{\delta} \in [0,1)$, such that for all $\delta > \underline{\delta}$, 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 $(n+2)$ states: (a) normal state (b) i-punishment state (these are n states), and (c) more-punishment state. The initial state is normal state. In normal state, the strategy recommends playing a_i to each Player i.

Consider Player j. If the state is normal and every player $i \in N$ plays a_i in a period, then the state remains normal in the next period. If the state is normal and a **unique** player $i \in N$ **does not play** a_i (here i can be equal to j), then the state becomes i-punishment. If the state is normal and **more than one player** in N **does not play** a_i, then state becomes more-punishment. Once the state becomes i-punishment for some i, it stays the same irrespective of the actions in subsequent periods. Similarly, once the state becomes more-punishment, it stays so irrespecitve of the actions in subsequent periods.

The strategy for Player j requires him to play a_j in normal state; play the action $\underline{a}_j^i$, corresponding to the minmax action profile against Player i, in i-punishment state, and play some fixed action (does not matter which one) in more-punishment state.

The strategy is shown in Table 44.

Predecessor state	Action profile observed	Current state	Recommended action
Normal	a	Normal	a_j
Normal	(a'_i, a_{-i})	i-punishment	$\underline{a}_j^i$
Normal	$(a'_S, a_{N\setminus S})$ with $\lvert S\rvert > 1$	more-punishment	Any fixed action
i-punishment	a'	i-punishment	$\underline{a}_j^i$
more punishment	a'	more punishment	Any fixed action

Table 41: 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 $v_i \equiv (u_i(a))$ for Player i. Suppose all the other players except i follows the prescribed strategy. Let the **best response** to a_{-i} give Player i a payoff $\bar{v}_i$ in the stage game G:

$$\bar{v}_i = \max_{a'_i \in A_i} u_i(a'_i, a_{-i}).$$

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 less than or equal to $\underline{v}_i$. Hence, if he deviates in period t, his **maximum possible payoff** from deviation is (the original payoff can be less than this):

$$(1-\delta)\big(v_i + \delta v_i + \ldots + \delta^{t-1}\bar{v}_i + \delta^t \underline{v}_i + \delta^{t+1}\underline{v}_i + \ldots\big)$$

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

$$\underline{\delta}_i := \frac{\bar{v}_i - v_i}{\bar{v}_i - \underline{v}_i} = 1 - \frac{v_i - \underline{v}_i}{\bar{v}_i - \underline{v}_i}.$$

and $\underline{\delta} := \max_i \underline{\delta}_i$. Note that by assumption $\bar{v}_i \geq v_i > \underline{v}_i$. Hence, $\underline{\delta}_{\in}[0,1)$ for all i and $\underline{\delta} \in [0,1)$. This proves the claim. ■

The exact version of folk theorems involve use of mixed actions by players in each period. We do not discuss strategies that involve mixed actions by players in each period. But informally, a mixed strategy can be "approximated" by cycles of pure actions. For instance in prisoner's dilemma, a mixed strategy $\frac{1}{2}C + \frac{1}{2}D$, can be "implemented" by allowing players to play C in odd periods and D in even periods: as $\delta \to 1$, the payoff from this strategy will be exactly the payoff from $\frac{1}{2}C + \frac{1}{2}D$.

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 (or even punishment for some long periods) may not be a reasonable threat. For instance, consider the game in Table 42. The mixmax payoffs of the players are $(0,1)$. Theorem 16 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. In fact L is a dominant strategy for Column player. 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 42: A Stage game

16.5 Subgame perfect equilibrium

To correct the problem identified in Nash equilibrium, a stricter notion of equilibrium has been proposed in repeated games. This notion demands equilibrium play not only at the starting point (i.e., in the reduced form) but also in every *subgame*. For every period t and every history $h_t \in H^t$ in period t, the infinitely repeated game starting at that history is referred to as a **subgame** of G^∞, and it is denoted as G^∞_{t,h^t}. Of course, $G^\infty = G^\infty_{1,\emptyset}$ and G^∞ is a subgame of itself.

DEFINITION **28** *A strategy profile s is a* **subgame perfect equilibrium** *of G^∞ if its restriction from any period t and any history h^t is a Nash equilibrium of the subgame G^∞_{t,h^t}.*

Clearly, a subgame perfect equilibrium is a Nash equilibrium.

Go back to the example in Table 42. If the game reaches a history where Player 1 has played B, Player 2 cannot play R forever in a subgame perfect equilibrium. This is because, L is a strictly dominant strategy. So, in every period t and at every history h^t, Player 2 can do better by playing L. Notice that this does not rule out more complicated strategy. For instance, Player 2 plays R for some periods and then switches back to L. So, trigger strategies cannot sustain (T, L) in in this game as subgame perfect equilibrium. Contrast this with the Nash equilibrium outcome where (T, L) in every period could be sustained through a trigger strategy.

However, a different type of trigger (*reversion to Nash*) can sustain subgame perefection. This is illustrated in the following where we redo Proposition 2 for subgame perfect equilibrium.

PROPOSITION **3** *Suppose $\delta \geq \frac{1}{2}$. Then, there is a subgame perfect equilibrium in the Prisoner's Dilemma game (Table 39), 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 *reversion to Nash 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. ■

16.6 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 strategies starting from that history. The following is quite immediate.

Proposition 4 *Suppose a^* is a Nash equilibrium of the stage game G and consider a strategy profile s^* where every player i plays a_i^* in every history of every period. Then s^* is a subgame perfect equilibrium of G^∞.*

Proof: This follows from the one-shot deviation principle which we prove in the next section. But we give a proof which does not require the use of one-shot deviation principle.

So, we need to establish that s^* is a Nash equilibrium of every subgame. If this strategy is not a Nash equilibrium of some subgame, then there is some history h^t (subgame) at which a Player i has a deviation. According to the original strategy profile s^*, the action profile generated is:

$$(a_i^*, a_{-i}^*), (a_i^*, a_{-i}^*), (a_i^*, a_{-i}^*), \ldots \tag{17}$$

By deviating, suppose the action profile stream generated is

$$(c_i^1, a_{-i}^*), (c_i^2, a_{-i}^*), (c_i^3, a_{-i}^*), \ldots \tag{18}$$

Note that Player i may have deviated at many histories leading to this action profile stream. Actions of players other than i do not change since they are following s^* which requires them to play action in a^*. Since a_i^* is a best response to a_{-i}^* action profile stream (18) cannot generates more payoff than the action profile stream (17) in G^∞. Hence, no deviation is possible for Player i. Thus, s^* is a Nash equilibrium of G^∞. ■

Consider the battle of sexes game in Table 43.

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

Table 43: Nash equilibrium in the Battle of the Sexes game

There are two Nash equilibria of the game: (X, x) and (Y, y). Notice that the two equilbria give different payoffs to players. Consider the following strategy: play (Y, y) in first period; in every history where (Y, y) is played in first period, play (Y, y); in every history where (Y, y) is not played in first period, play (X, x).

Note that in every history, the strategy recommends playing some Nash equilibrium of the stage game. We claim that this strategy profile is **not** a Nash equilibrium of the infinitely repeated game if $\delta \geq \frac{1}{2}$. Consider the first period. Suppose Player 1 deviates plays X in first period and then plays according to the recommended strategy (i.e., Player 1 changes her action only in the first period, and follows the recommended strategy from thereon). Then, her payoff stream is $(0, 2, 2, 2, \ldots)$. Hence, her expected payoff is 2δ. Her payoff from the recommended strategy is 1. So, if $\delta \geq \frac{1}{2}$, then the deviation is profitable. So, the recommended strategy is not a Nash equilibrium of G^{∞}. This shows that even though playing a **fixed** Nash equilibrium of the stage game is a subgame perfect equilibrium (hence, a Nash equilibrium also) of the infinitely repeated game, playing different Nash equilibria of the stage game at different histories will not necessarily be a Nash equilibrium of the infinitely repeated game.

Now, denote by v_i^* the worst payoff of Player i over all Nash equilibria action profiles in G. Also, denote the corresponding Nash equilibrium profile as $a^{*,i}$. We are now ready to state a mild version of the perfect folk theorem.

THEOREM **17 (Pure Perfect Folk Theorem with Nash Reversion)** *Suppose a is any action profile in stage game G 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. It is described by $(n+2)$ states: (a) normal state (b) i-punishment state (n such states), and (c) more-punishment state. The initial state is normal state. In normal state, the strategy recommends playing a_i to each Player i.

Consider Player j. If the state is normal and every player $i \in N$ plays a_i, then the state remains normal. If the state is normal and a **unique** player $i \in N$ **does not play** a_i (here i can be equal to j), then the state becomes i-punishment. If the state is normal and **more than one player** in N **does not play** a_i, then state becomes more-punishment. Once the state becomes i-punishment, it remains so irrespective of the actions taken by all the players. Similarly, if the state becomes more-punishment, it remaisn so irrespective of the actions taken by other players.

The strategy for Player j requires him to play a_j in normal state, play the action $a_j^{*,i}$ corresponding to the **worst Nash equilibrium profile** of Player i (giving Player i a maximum payoff of v_i^*) for i-punishment state, and play an action corresponding some **fixed Nash equilibrium** (does not matter which one) in more-punishment state.

The strategy is shown in Table 44.

In any history which is either a i-punishment state or a more-punishment state, the strategy recommends playing a fixed Nash equilibrium of the stage game. By Proposition 4,

Predecessor state	Action profile observed	Current state	Recommended action
Normal	a	Normal	a_j
Normal	(a'_i, a_{-i})	i-punishment	Worst Nash for Player i - $a_j^{*,i}$
Normal	$(a'_S, a_{N\setminus S})$ with $\lvert S\rvert > 1$	more-punishment	Any fixed Nash action
i-punishment	a'	i-punishment	Worst Nash for Player i - $a_j^{*,i}$
more punishment	a'	more punishment	Any fixed Nash action

Table 44: Trigger strategy for perfect folk theorem

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 $u_i(a)$. Player i can deviate in many periods in this subgame. It is without loss of generality that she deviates in the first period itself (it not, we look at the subgame starting from the period where she deviates). Denote the action profile associated with v_i^* payoff of Player i as b^* (it is a Nash equilibrium of G).

By deviating, the action profile stream that gets generated is the following:

$$(c_i^1, a_{-i}), (c_i^2, b_{-i}^*), (c_i^3, b_{-i}^*), \dots \tag{19}$$

The action profile stream in (19) is dominated by the following action profile stream (for Player i) since b^* is a Nash equilibrium of G:

$$(c_i^1, a_{-i}), (b_i^*, b_{-i}^*), (b_i^*, b_{-i}^*), \dots \tag{20}$$

The maximum payoff from the action profile stream in (20) is

$$(1-\delta)\Big[\max_{a'_i \in A_i} u_i(a'_i, a_{-i}) + \delta v_i^* + \delta^2 v_i^* + \dots\Big] = (1-\delta)d_i(a_{-i}) + \delta v_i^*,$$

where $d_i(a_{-i}) := \max_{a'_i \in A_i} u_i(a'_i, a_{-i})$. Note that the maximum of the expression $\max_{a'_i \in A_i} u_i(a'_i, a_{-i})$ need not occur at a_i since (a_i, a_{-i}) is not necessarily a Nash equilibrium of the stage game G. In particular, a_i need not be a best response to a_{-i} in the stage game G. Hence, his payoff from deviation is

$$(1-\delta)d_i(a_{-i}) + \delta v_i^*.$$

Hence, to be a subgame perfect equilibrium, we will need that $u_i(a) \geq (1-\delta)d_i(a_{-i}) + \delta v_i^*$. This is true if

$$\delta > \frac{d_i(a_{-i}) - u_i(a)}{d_i(a_{-i}) - v_i^*} = \underline{\delta}_i.$$

Note that $d_i(a_{-i}) \geq u_i(a) > v_i^*$ ensures that $\underline{\delta}_i \in [0,1)$. In other words, for $\delta \in (\underline{\delta}, 1)$, the recommended strategy is a subgame perfect equilibrium, where $\underline{\delta} = \max_i \underline{\delta}_i$. This completes the proof. ■

	L	C	R
T	$2,2$	$0,3$	$0,4$
M	$3,0$	$1,1$	$\frac{1}{2},0$
B	$4,0$	$0,\frac{1}{2}$	$-1,-1$

Table 45: A Stage game

Consider the game in Table 45. There is a unique pure strategy Nash equilibrium of this stage game: (M,C) giving a payoff of $(1,1)$. The minmax payoff profile is $(\frac{1}{2},\frac{1}{2})$. The minmax action profile for Player 1 (row player) is (M,R) and for Player 2 (column player) is (B,C). But these action profiles cannot be used to punish in a trigger strategy.

The Nash reversion strategy punishes by playing the unique Nash equilibrium of the stage game (T,L). So, we can support (T,L) in subgame perfect equilibrium of the infinitely repeated game (even though (T,L) is not a Nash equilibrium of the stage game). To see this consider the strategy profile which recommends playing (T,L) to both the players if the history consists of (T,L) and else recommends playing (M,C). To see that it can sustained in subgame perfect equilibrium, we need to show that at every history, the strategies induce a Nash equilibrium. In any history which does not consist of only (T,L) play in the past, our strategy recommends playing the Nash equilibrium of the stage game. By Proposition 4, this is a Nash equilibrium of this subgame. Hence, we only need to worry about histories where only (T,L) is played in the past. In such a history, consider Player 1 (since the game is symmetric, similar argument works for Player 2). If Player 2 follows this strategy, by playing this strategy, Player 1 gets a payoff of 2. By deviating in some period t, she gets a maximum of 4 in period t, but then, Player 2 shifts to playing C. So, the best she can get from period $(t+1)$ onwards is 1. So, any deviation results in the following best case payoff stream: $4, 1, 1, \ldots$. The discounted payoff of this stream is $4(1-\delta) + \delta = 4 - 3\delta$. For this to be an equilibrium, we need $2 \geq 4 - 3\delta$ or $\delta \geq \frac{2}{3}$. Hence, for all $\delta \geq \frac{2}{3}$, we can ensure a subgame perfect equilibrium of the infinitely repeated game where (T, L) is played in every period *along the equilibrium path*.

If we had just required a Nash equilibrium of the infinitely repeated game, the calculations would have been different since we could punish the deviating player by the minmax action profile. So, deviation would have generated a payoff stream of: $4, \frac{1}{2}, \frac{1}{2}, \ldots$, which has discounted payoff of $(1-\delta)4 + \frac{\delta}{2}$. So, equilibrium requirement is $2 \geq (1-\delta)4 + \frac{\delta}{2}$, which gives $\delta \geq \frac{4}{7}$.

16.7 The One-Shot Deviation Principle

The exact version of perfect Nash theorem is stronger than the Nash reversion theorem we stated. It says (under mild conditions on the stage game) that every strictly enforceable payoff vector can be sustained as a subgame perfect equilibrium. Of course, the strategy profile required is more sophisticated. Such sophisticated strategies are difficult to verify if they are a subgame perfect equilibrium. Fortunately, there is a simplification possible.

The *one-shot deviation principle* is a useful tool in the repeated games setting. It says that to verify if a certain strategy is a best response to others strategies at every period and every history, we need to verify optimality of this strategy at every period and every history by deviation at that history only. To understand this, we remind the notion of payoff path of a strategy profile. Given a strategy profile s, each agent i plays action $s_i^1(\emptyset)$ in period 1. Denote this action profile as a^1. Now, inductively define action profile a^t given that we have defined $(a^1, \ldots, a^{t-1})$. The action profiles $(a^1, \ldots, a^{t-1})$ is exactly the history h^t that is relevant in period t for computing payoffs since it will be reached if strategy profile s is played. So, $a_i^t := s_i^t(h^t)$ for each $i \in N$. Now, the payoff path is $(a^1, a^2, \ldots)$. It is clear that if agents play strategy profile s, then repeated game will see action profiles $(a^1, a^2, \ldots)$ and the payoff to each agent $i \in N$ is thus

$$u_i^*(s) = (1-\delta)\sum_{t=1}^{\infty} \delta^{t-1} u_i(a^t).$$

Using the same idea, we can also think of a payoff path of a strategy profile s starting from a history h^t in period t. This means we pretend as if the game has reached history h^t

and the infinite sequence of action profiles that will be played from history h^t due to strategy s. This is the infinite sequence of action profiles that will generate payoffs to agents once we start the game at h^t.

Once we understand the notion of a payoff path, the one-shot deviation principle is not hard to understand. We start from a useful lemma, which basically says that if there is a unilateral deviation from a strategy, then there is a unilateral deviation from a strategy which differs at finite histories.

Lemma **10** *Suppose for some player $i \in N$ and some strategy profile s_{-i} of other players, we have $u_i^*(s_i', s_{-i}) > u_i^*(s_i, s_{-i})$ for some s_i and s_i'. Then, there exists a strategy s_i'' which differs from s_i at finite histories such that $u_i^*(s_i'', s_{-i}) > u_i^*(s_i, s_{-i})$.*

Proof: If s_i and s_i' differ from each other at finite histories, we are done (s_i'' can be chosen to be s_i'). Suppose s_i and s_i' differ from each other at infinite histories. Without loss of generality, assume that s_i and s_i' differ from each other from period 1 - if they do not differ from period 1, consider the restrictions of these strategies from the first period onwards when they differ from each other, and treat it as period 1. Let $(a^1, a^2, \ldots)$ be the action profiles played in (s_i', s_{-i}) on the payoff path. Hence,

$$u_i^*(s_i', s_{-i}) = (1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}u_i(a^t).$$

Let M be the largest difference in payoff of player i across any pair of action profiles in the stage game G. So, $M = \max_{x,y \in A}[u_i(x) - u_i(y)]$. Further, let $\gamma := u_i^*(s_i', s_{-i}) - u_i^*(s_i, s_{-i})$. By definition $\gamma > 0$. Let $\hat{t} > 1$ be such that $M\delta^{\hat{t}-1} < \frac{\gamma}{2}$. Clearly, such a $\hat{t}$ can be found since $\delta < 1$. Let $(b^{\hat{t}}, b^{\hat{t}+1}, \ldots)$ be the action profiles due to strategy profile (s_i, s_{-i}) on its payoff path starting from history $h^{\hat{t}} \equiv (a^1, \ldots, a^{\hat{t}-1})$ onwards.

Hence, we can now write

$$
\begin{aligned}
\gamma = u_i^*(s_i', s_{-i}) - u_i^*(s) &= (1-\delta)\sum_{t=1}^{\hat{t}-1}\delta^{t-1}u_i(a^t) + (1-\delta)\sum_{t=\hat{t}}^{\infty}\delta^{t-1}u_i(a^t) - u_i^*(s) \\
&= (1-\delta)\sum_{t=1}^{\hat{t}-1}\delta^{t-1}u_i(a^t) + (1-\delta)\sum_{t=\hat{t}}^{\infty}\delta^{t-1}\big[u_i(a^t) - u_i(b^t)\big] \\
&+ (1-\delta)\sum_{t=\hat{t}}^{\infty}\delta^{t-1}u_i(b^t) - u_i^*(s) \\
&< (1-\delta)\sum_{t=1}^{\hat{t}-1}\delta^{t-1}u_i(a^t) + (1-\delta)\sum_{t=\hat{t}}^{\infty}\delta^{t-1}u_i(b^t) - u_i^*(s) + (1-\delta)\sum_{t=\hat{t}}^{\infty}\delta^{t-1}M \\
&= (1-\delta)\sum_{t=1}^{\hat{t}-1}\delta^{t-1}u_i(a^t) + (1-\delta)\sum_{t=\hat{t}}^{\infty}\delta^{t-1}u_i(b^t) - u_i^*(s) + \delta^{\hat{t}-1}M \\
&< (1-\delta)\sum_{t=1}^{\hat{t}-1}\delta^{t-1}u_i(a^t) + (1-\delta)\sum_{t=\hat{t}}^{\infty}\delta^{t-1}u_i(b^t) - u_i^*(s) + \frac{\gamma}{2}.
\end{aligned}
$$

Hence, we get

$$
\frac{\gamma}{2} < (1-\delta)\sum_{t=1}^{\hat{t}-1}\delta^{t-1}u_i(a^t) + (1-\delta)\sum_{t=\hat{t}}^{\infty}\delta^{t-1}u_i(b^t) - u_i^*(s). \tag{21}
$$

Now, we can construct a new strategy s_i'' as follows: s_i'' and s_i' coincide from period 1 to period $\hat{t}-1$. Once history $(a^1, \ldots, a^{\hat{t}-1})$ is reached, i plays according to s_i along the payoff path. Hence, the payoff path of (s_i'', s_{-i}) is $(a^1, \ldots, a^{\hat{t}-1}, b^{\hat{t}}, b^{\hat{t}+1}, \ldots)$. So, s_i'' and s_i differ at finite histories. This gives a payoff equal to:

$$
u_i^*(s_i'', s_{-i}) = (1-\delta)\sum_{t=1}^{\hat{t}-1}\delta^{t-1}u_i(a^t) + (1-\delta)\sum_{t=\hat{t}}^{\infty}\delta^{t-1}u_i(b^t) \tag{22}
$$

Using Inequality (21) and Inequality (22), we get

$$
\frac{\gamma}{2} < u_i^*(s_i'', s_{-i}) - u_i^*(s).
$$

Since $\gamma > 0$, we get that $u_i^*(s_i'', s_{-i}) > u_i^*(s_i, s_{-i})$. ■

We will now prove the one-shot deviation principle. Given any strategy s_i, we consider the restriction of this strategy from an arbitrary t-period history h^t and denote it as $s_i^{h^t}$. We denote the restriction of strategy profile s from history h^t as s^{h^t}. Notice that $s_i^{h^t}$ defines a feasible strategy of Player i in subgame G^{∞}_{t,h^t}.

DEFINITION **29** *A pair of strategies s_i and $\bar{s}_i$ of Player i are one-shot deviation at history h^t if $s_i(h^t) \neq \bar{s}_i(h^t)$ and for all $(t', h^{t'}) \neq (t, h^t)$, we have $s_i(h^{t'}) = \bar{s}_i(h^{t'})$.*

So, a one-shot deviation strategy just changes action at exactly *one* history and leaves action at every other history unchanged. We denote the payoff Player i starting from history h^t from strategy profile s as $u_i^{h^t}(s)$. This is the payoff that Player i will receive if she reaches history h^t and players play according to s from h^t onwards. Before we define one-shot deviation optimality, observe that for Player i at history h^t, by following s_i from this period onwards at history h^t gives a payoff equal to $u_i^{h^t}(s_i, s_{-i})$, where s_{-i} are strategies of other players. But this can be decomposed into two parts: (i) utility from the stage game in this period and discounted utility in history h^{t+1} onwards, where h^{t+1} consists of h^t and the action profile in period t. Formally:

$$u_i^{h^t}(s_i, s_{-i}) = u_i(s_i(h^t), s_{-i}(h^t)) + \delta u_i^{h^{t+1}}(s_i, s_{-i})$$

This simple equation reflects the essense of one-shot deviation. One shot deviation only changes the stage game payoff at history h^t and the history from period $(t+1)$ onwards (by having a different action profile in history h^t); everything else remains the same.

DEFINITION **30** *A strategy s_i is* **one-shot deviation (OSD)** *optimal for s_{-i} for Player i if for every history h^t and every s_i' which is one-shot deviation at history h^t, we have*

$$u_i^{h^t}(s_i, s_{-i}) \geq u_i^{h^t}(s_i', s_{-i}).$$

The following theorem is the one-shot deviation principle.

THEOREM **18 (One-shot deviation principle)** *A strategy profile s is a subgame perfect equilibrium of G^{∞} if and only if for every $i \in N$, s_i is OSD optimal for s_{-i}.*

Proof: If s is a subgame perfect equilibrium, then it is clearly OSD optimal. We do the other direction. Suppose s is not a subgame perfect equilibrium. Then, there is another strategy s_i' of some Player i such that in some subgame G^{∞}_{t,h^t}, strategy s_i' (restricted to this subgame) does better than s_i (restricted to this subgame). Without loss of generality assume $t = 1$ and $h^t = \emptyset$ - this essentially means that we pretend as if the game started from this history. Hence, we have $u_i^*(s_i', s_{-i}) > u_i^*(s_i, s_{-i})$. By Lemma 10, we can assume that s_i' and s_i differ from each other at finite histories. We also choose s_i' such that over all strategies which differ from s_i and result in successful deviation (i.e., increased payoff in G^{∞} from s_i), s_i' **minimizes** the number of periods for which such a strategy differs from s_i. Let $\hat{t}$ be the last period where there is a history $h^{\hat{t}}$ such that s_i' and s_i differ from each other. By our assumption, take any other strategy s_i'' which results in increased payoff from s_i and which differs from s_i at finite periods. Let $\hat{t}'$ be the last period where there is a history $h^{\hat{t}'}$ such that s_i'' and s_i differ from each other. Then, it must be that $\hat{t}' \geq \hat{t}$. Further, $\hat{t} > 1$.

At every history $h^{\hat{t}}$ in period $\hat{t}$, strategies s_i and s_i' differ by a single period. So, OSD optimality implies that if we change s_i' from $h^{\hat{t}}$ onwards to s_i, i.e., we set the action of Player i at history $h^{\hat{t}}$ to $s_i(h^{\hat{t}})$, we get a new strategy which is also better than s_i from history $h^{\hat{t}}$ onwards. Repeating this for all histories in period $\hat{t}$, we have a new strategy that differs from s_i till $\hat{t} - 1$ only and which gives more payoff than s_i. This is a contradiction to our assumption that s_i' is minimal. ■

We can now illustrate the exact version of this for $n = 2$ using an example. The stage game is shown in Table 46.

	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 46: A Stage game

Notice that the minmax payoff vector is $(0, 0)$. The unique pure Nash equilibrium is (T, L). Using Theorem 17 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 **19** *Suppose $\delta \geq \frac{1}{2}$. Then, there is a subgame perfect equilibrium of the infinitely repeated game of the stage game in Table 46 such that (M, C) is played in every period on* *equilibrium path.*

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. Note that (B, R) is not a minmax action profile for any player. To punish player 1, Player 2 must play R but Player 1 will best repond with T. So, (T, R) is the minmax punishment of Player 1. Similarly, (B, L) is the minmax punishment profile of Player 2. However, here recommended strategy profile for both the punishments will be (B, R) giving the players strictly lower payoff (-1) than their minmax payoffs (0).

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. These kind of strategies are called **carrot and stick** strategies.

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:

$$(1-\delta)\big(2+\delta(-1)+\delta^2+\delta^3+\dots\big)=(1-\delta)(1-2\delta)+1.$$

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

$$(1-\delta)\big(-1+\delta+\delta^2+\dots\big)=1-2(1-\delta).$$

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

$$(1-\delta)\big(0+\delta(-1)+\delta^2+\delta^3+\dots\big)=1-(1+2\delta)(1-\delta).$$

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 general version 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. Using such strategies, the general folk theorem says that all strictly enforceable payoff vectors can be achieved in a subgame perfect equilibrium of the infinitely repeated game (for sufficiently large δ and under some assumptions on the payoffs of the stage game).

The above example also highlights the role of one-shot deviation principle. Without using the one-shot deviation principle, it will be possible to prove the subgame perfect equilibrium of the carrot-and-stick strategy profile, but it will be really cumbersome. The one-shot deviation principle simplifies the arguments significantly.

16.8 An application: tacit collusion

Consider two firms $\{1, 2\}$ in a Bertrand competition setting. We have already studied the stage game corresponding to this model. In the stage game firms set prices $p_1, p_2 \in [0, 1]$. Given the prices (p_1, p_2), the demand for firm i is:

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

Given this demand function, the utility of firm i at prices (p_1, p_2) is given by

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

There is a **Nash** equilibrium of this game where firms set prices $p_1^* = p_2^* = \frac{1}{3}$ - verify this. In fact, iterative elimination of never best-responses lead to this unique outcome. Notice that the equilibrium utilties of the firms are given by:

$$u_i(p_1^*, p_2^*) = \frac{1}{3}(1 - \frac{2}{3} + \frac{1}{3}) = \frac{2}{9}.$$

Now, suppose this game is repeated infinitely with a common discount factor $\delta \in (0, 1)$. We will show that the if firms are *sufficiently patient* (i.e., δ is sufficiently high), then firms can sustain price $\bar{p}_1 = \bar{p}_2 = \frac{1}{2}$. We can sustain this using the Nash reversion trigger strategies. So, in normal state, firms choose $\bar{p}_1 = \bar{p}_2 = \frac{1}{2}$. If any firm deviates, we go to punishment state, where firms play $p_1^* = p_2^* = \frac{1}{3}$. We do not need to worry about punishment state histories, since we play Nash equilibrium there. In normal state history, playing $\bar{p}_i$ gives firm i a constant payoff of $\frac{1}{2}(1 - 1 + \frac{1}{2}) = \frac{1}{4}$. By one-shot deviation principle, we only check deviation in one period. If firm i deviates, then his maximum payoff can be computed as follows. If it sets a price of x, then given that other firm is setting a price of $\frac{1}{2}$, its utility is

$$x(1 - 2x + \frac{1}{2}) = x(\frac{3}{2} - 2x).$$

This is maximized at $x = \frac{3}{8}$ giving the firm a payoff of $\frac{9}{32}$ in that period. However, in susbsequent period, the firm gets only $\frac{2}{9}$ - Nash equilibrium payoff. Hence, the total payoff is:

$$(1-\delta)(\frac{9}{32}+\delta\frac{2}{9}+\delta^2\frac{2}{9}+\ldots) = (1-\delta)\frac{17}{288}+\frac{2}{9}$$

So, for subgame perfect equilibrium, we need

$$\begin{aligned} \frac{1}{4} &\geq (1-\delta)\frac{17}{288}+\frac{2}{9} \\ \Leftrightarrow \frac{1}{36} &\geq (1-\delta)\frac{17}{288} \\ \Leftrightarrow \delta &\geq \frac{9}{17}. \end{aligned}$$

Hence, firms can maintain a price of $\bar{p}_1 = \bar{p}_2 = \frac{1}{2}$ in the subgame perfect equilibrium with payoff $\frac{1}{4} > \frac{2}{9} = u_i(p_1^*, p_2^*)$ (i.e., achieving a payoff greater than stage game Nash equilibrium payoff). So, even though firms do not communicate, a form of **tacit** collusion emerges in equilibrium of repeated play. This is an empirically observed phenomenon. Consider the interpretation that δ is the probability with which firms will compete in the next period over this product. In markets, where established firms are competing over products that are used for long horizons, we see that prices are higher than stage game equilibrium. However, when firms compete over products with a deadline (say, Christmas tree sale), then prices are lower (closer to one-stage game Nash equilibrium prices).

17 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.

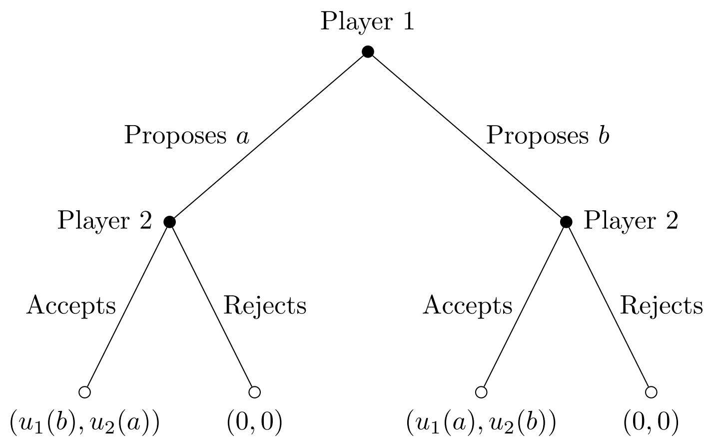

Figure 13: Extensive form game with perfect information

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 13 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

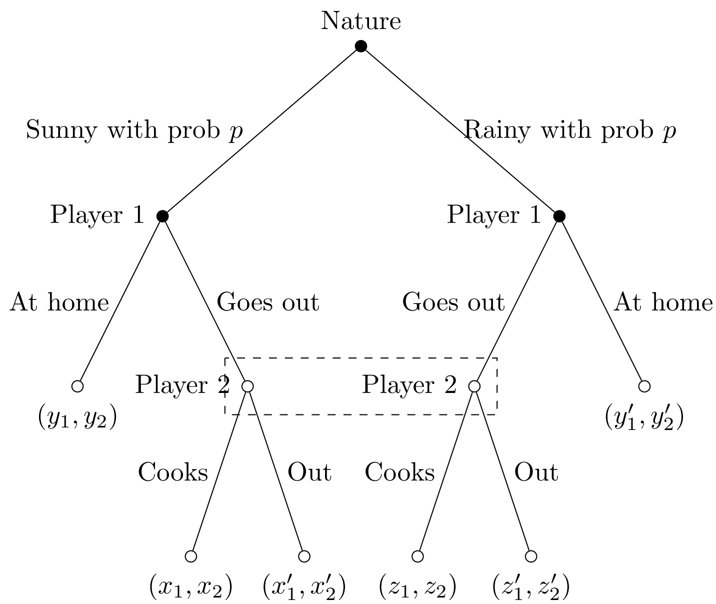

Figure 14: Extensive form game with information sets

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 14 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.

18 Extensive Form 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 15. 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, whereas v_1 has two children: $\{v_3, v_4\}$.

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

DEFINITION **31** *An* **extensive form game of perfect information** *is*

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

where

- N *is the set of players*
- (V, E, r) *is a rooted tree called the* **game tree**, *where*

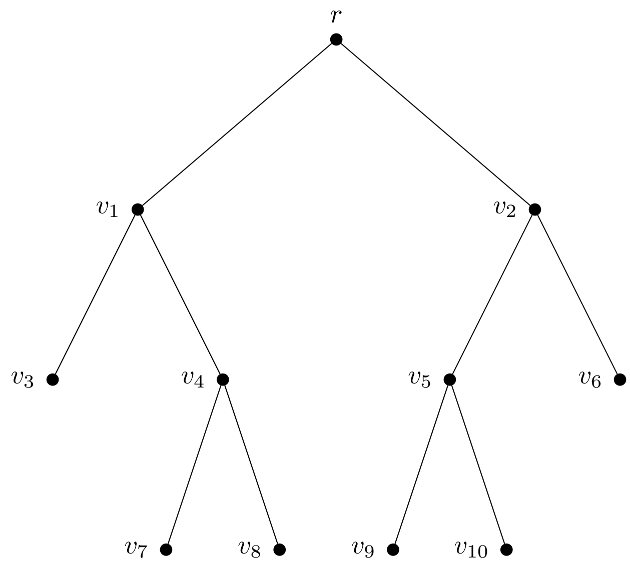

Figure 15: An example of a rooted tree

- *Each non-leaf vertex $x \in V$ specifies a player, called the decision maker at x, in N who will take an action at this vertex.*
- *Each leaf or terminal vertex $x \in V$ is a* **payoff** *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.*
- *Root vertex r specififes the first player in N to take an action.*

- *$A(x)$ is the set of actions available at vertex x (they identify the set of edges from x which lie on the path from x to all the leaf nodes). Note that if x is a leaf vertex, then $A(x)$ is an empty set.*
- *$\{V_i\}_{i\in N}$ is a partitioning of the set of decision vertices. Hence, V_i represents the set of decision vertices where Player i takes action.*
- *For every player $i \in N$, $u_i(x)$ assigns a payoff for every terminal 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 13. 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.

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

$$s_i : V_i \to \cup_{x \in V_i} A(x) \text{ such that } s_i(x) \in A(x) \; \forall \; x \in V_i.$$

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 16, 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.

19 Equilibrium for Extensive Form Games

We discuss equilibrium concepts 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

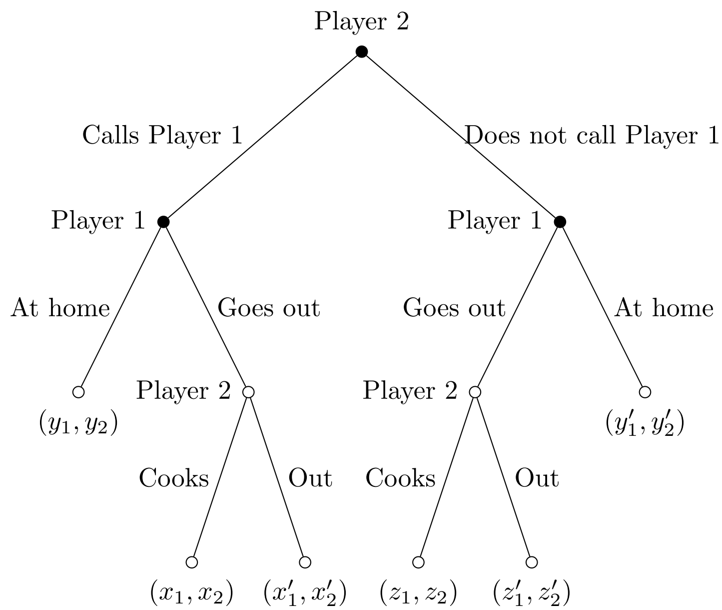

Figure 16: Extensive form game without single move property

form games. Representing an extensive form game as a strategic form game is quite easy: for every player i and every 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 $u_i(x_s)$.

DEFINITION **32** *A strategy profile* $s \equiv (s_1, \ldots, s_n)$ *is a Nash equilibrium of* Γ *if for all* $i \in N$ *and for all* s'_i

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

This definition just says that consider the reduced-form strategic form game and consider the Nash equilibrium of that game. In other words, it ignores all the extensive form (sequential) play of players actions in the game. Hence, Nash equilibrium is not the correct solution concept for extensive form games. We illustrate this with an example.

Consider the game in Figure 17. The reduced strategic form representation of this game

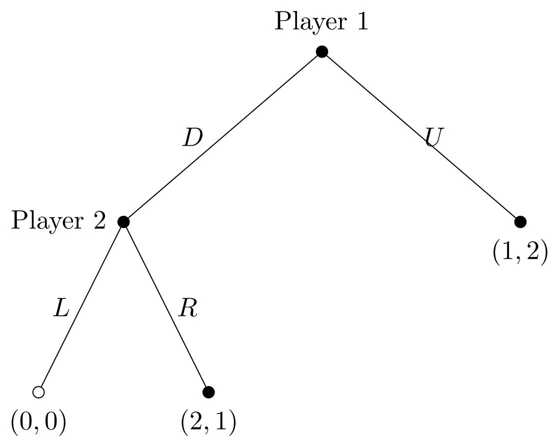

Figure 17: Nash equilibrium

is shown in Table 48. 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 47: Reduced strategic form of the game in Figure 17

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* decision vertex, and *not just* decision vertices reached on equilibrium path.

19.1 Subgame Perfect Equilibrium

We now discuss a *refinement* to Nash equilibrium for extensive form game. This is the single-most important solution concept for extensive form games. it enforces and formalizes the idea of credibility by using the notion of subgames.

The **subgame** of an extensive form game of perfect information

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

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

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

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. If a subgame of Γ starts at x, we will denote the utility function of each player i in this subgame $\Gamma(x)$ by u_i^x.

Note that a game is a subgame of itself. So, every game has a subgame. Game in Figure 16 has many subgames: there are two subgames starting with Player 1's two decision nodes; there are three subgames starting with Player 2's three decision nodes. In general, the number of subgames in a game equals the number of decision nodes in the game.

DEFINITION **33** *A strategy profile s is a* **subgame perfect equilibrium (SPE)** *of the extensive form game Γ if for every subgame of Γ the strategy profile s 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. We document this as a fact below.

FACT **2** *Every subgame perfect equilibrium is a Nash equilibrium.*

The game in Figure 17 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.

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 the one-shot deviation property. This idea is very similar to the one-shot deviation principle in the repeated games setting. Two strategies s_i and s'_i are one-shot deviation at decision vertex x if $s_i(x) \neq s'_i(x)$ but $s_i(y) = s_i(y)$ for all $y \neq x$ and $y \in V_i$.

DEFINITION **34** *A strategy s_i of Player i is* **one-shot deviation (OSD) optimal** *for s_{-i} if for every decision vertex $y \in V_i$ of Player i and each strategy s'_i of Player i which is a one-shot deviation from s_i at y, we have*

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

A fundamental result is that these notions are the same. The result below allows for a decision vertex to have arbitrary (possibly infinite) number of actions - hence, the game tree may have infinite number of vertices. However, it restricts itself to games having finite number of *stages*. To understand the notion of stage, let $L(x)$ denote the length of the longest path from a decision vertex x to any terminal vertex reachable from x. We will say a game Γ has **finite number of stages** if $L(x)$ is finite for each decision vertex x.

Before we go into the theorem, for every strategy profile s and every subgame $\Gamma(y)$ we clarify the notion of a **payoff path**. It consists of sequence of nodes starting with y as follows: $(y = y_0, y_1, \dots, y_k)$, where y_k is a payoff vertex, and for every $j \in \{0, 1, \dots, k-1\}$, the following holds: y_{j+1} is the decision vertex obtained from y_j when the deciding agent, say i (i.e., $y_j \in V_i$), takes action $s_i(y_j)$. Every strategy profile induces a payoff path from every decision vertex, leading to a payoff vertex. This gives its payoff.

The proof strategy for OSD optimality equivalent to subgame perfect equilibrium is much simpler than the repeated game version because of finite number of stages.

THEOREM **20** *Let Γ be an extensive form game of perfect information and finite number of stages. Then the following are equivalent.*

1. *s is a subgame perfect equilibrium.*

2. *For every $i \in N$, s_i is OSD optimal for s_{-i}.*

Proof: The implication (1) $\Rightarrow$ (2) is immediate from definitions. This is because OSD optimality only requires optimality over one-shot deviation strategies but subgame perfect equilibrium requires it over a larger set of strategies.

(2) $\Rightarrow$ (1). Suppose s is a strategy profile satisfying (2) but not a subgame perfect equilibrium. Then, there is a subgame $\Gamma(y)$, which is rooted at y, such that for some $i \in N$, there is another strategy s'_i such that $u_i^y(x_{(s'_i,s_{-i})}) > u_i^y(x_{(s_i,s_{-i})})$. Let $P' \equiv (y'_0 = y, y'_1, \ldots, y'_k)$ be the payoff path of (s'_i, s_{-i}). Let $P \equiv (y_0 = y, y_1, \ldots, y_\ell)$ be the payoff path of (s_i, s_{-i}). Let y_j be the first vertex in P' and P where these paths are different. It is without loss of generality to look at the subgame $\Gamma(y_j)$. Hence, we assume (without loss of generality) $j = 0$.

First, we restore the actions of Player i to the action according to s_i at all decision vertices which do not lie in P'. In other words, consider a strategy s''_i such that $s''_i(z) = s_i(z)$ if $z \notin P' \cap V_i$ and $s''_i(z) = s'_i(z)$ if $z \in P' \cap V_i$. Hence, the payoff path of (s''_i, s_{-i}) and (s'_i, s_{-i}) are the same: P'. Thus, $u_i^y(x_{(s''_i,s_{-i})}) > u_i^y(x_{(s_i,s_{-i})})$.

Now, let $y_j \in V_i$ be the last vertex on the path P' where $s''_i(y_j) \neq s_i(y_j)$. Then, in the subgame $\Gamma(y_j)$, strategy s''_i and s_i are one-shot deviations at y_j. Hence, setting $s''_i(y_j) = s_i(y_j)$ improves the payoff in the subgame $\Gamma(y_j)$ due to OSD optimality. We repeat this procedure backwards from y_j along the path P' till we restore the entire path P' to actions recommended by s_i and improving the payoff in each restoration. This will give us $u_i^y(x_{(s_i,s_{-i})}) > u_i^y(x_{(s_i,s_{-i})})$, which is a contradiction. ■

Finally, an easy method to compute a strategy profile satisfying one-shot deviation principle 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 decision maker 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 **35** *A strategy profile that survives the above procedure is said to be a strategy profile* **surviving the backward induction procedure**.

An easy corollary of Theorem 20 is the following.

Corollary 2 *A strategy profile is a subgame perfect equilibrium if and only if it survives the backwards induction procedure.*

In the game in Figure 17, 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 18. 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 18.

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.

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 both the players reach 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.

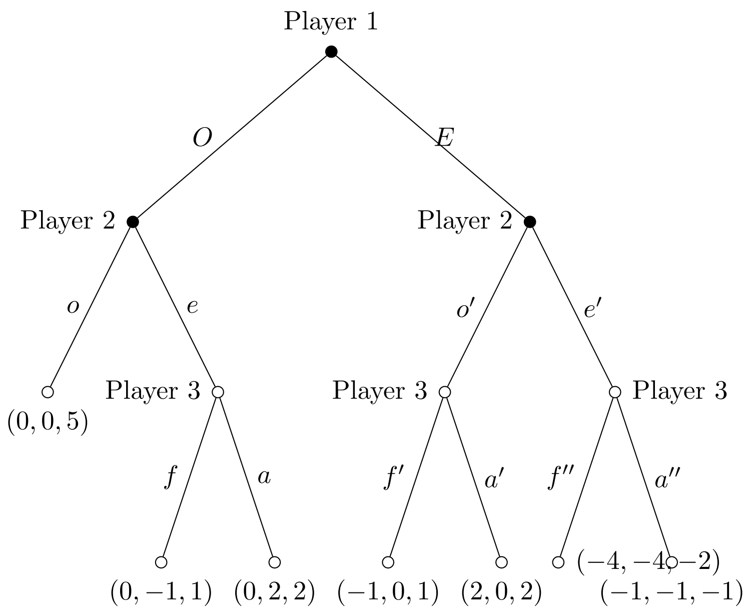

Figure 18: Backward induction

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.

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 20 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*.

20 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 decision vertices. There are two natural ways to define *randomized* strategies in this environment. The first one says that we define a probability distribution over the set of all pure strategies. This is the notion of a mixed strategy. Formally, a **mixed strategy** of Player i is $\sigma_i \in \Delta \prod_{x \in V_i} A(x)$.

Consider the game in Figure 19. Player 1 has two pure strategies - we roughly write it as $\{x, y\}$ to denote that in his only decision vertex, 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 decision vertex he plays A and in the other decision vertex, 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

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

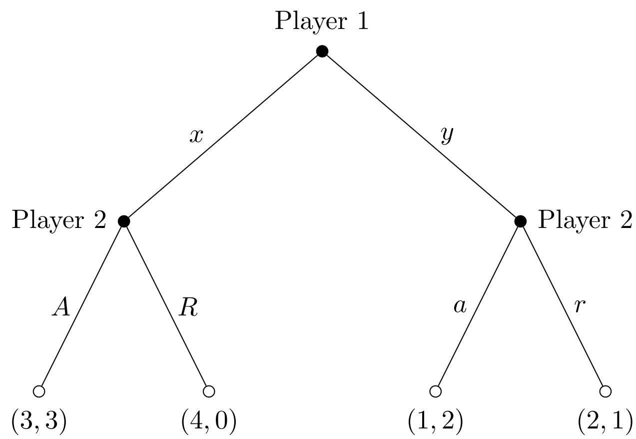

Figure 19: Extensive form game with perfect information

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

In the game in Figure 19, 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 decision vertices - when a player reaches his decision vertex, he randomizes over the actions at that decision vertex 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 19. 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 decision vertices 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 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? Remember, there is a unique path from the root vertex to x in Γ. Then, $\rho(x; \sigma)$ is the multiplication of playing each of the actions along this path (which can be computed from σ).

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. Consider the behavior strategy where Player 2 plays $\frac{1}{2}A + \frac{1}{2}R$ and $\frac{1}{3}a + \frac{2}{3}r$ in her two decision vertices. 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. In particular, since the left two terminal vertices are reached with probability $\frac{1}{4}$ each, the candidate behavior strategy for Player 2 at her left decision vertex is just $\frac{1}{2}A + \frac{1}{2}R$, which is just computed using conditional probabilties: the probability of reaching that decision vertext is $\frac{1}{2}$.

Definition **36** *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.*

Formally, Harold Kuhn established the following theorem.

Theorem **21** *In every extensive game of perfect information, every mixed strategy of a player is outcome equivalent to a 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 20. 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. We need to construct behavior strategies which is outcome equivalent to this. Consider the decision vertex 2 of Player 2. A natural candidate of his behavior strategy is the *conditional* probability of agent 2 playing ℓ (and r can be computed similarly) given that this decision vertex is reached:

$$\frac{\sigma_2(L\ell)p_u}{\sigma_2(L\ell)p_u + \sigma_2(Lr)p_u} = \frac{\sigma_2(L\ell)}{\sigma_2(L\ell) + \sigma_2(Lr)}.$$

Similarly, the candidate behavior strategy for the first decision node of Player 2 is to play L with probability

$$(\sigma_2(L\ell) + \sigma_2(Lr))$$

and R with the remaining probability. These candidates for behavior strategy generates the following probability of reaching the decision vertex with payoff $(4,1)$:

$$(\sigma_2(L\ell) + \sigma_2(Lr))p_u \frac{\sigma_2(L\ell)}{\sigma_2(L\ell) + \sigma_2(Lr)} = p_u \sigma_2(L\ell),$$

which is also the probability of reaching this decision vertex by strategy profile (p_u, σ_2).

Doing the calculations reveal that 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 decision vertex. 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 decision vertex by the following conditional probability:

$$\frac{\sigma_2(L\ell)}{\sigma_2(L\ell) + \sigma_2(Lr)} = \frac{1}{2}.$$

The proof of Kuhn's theorem formalizes this and shows that such computations are always possible.

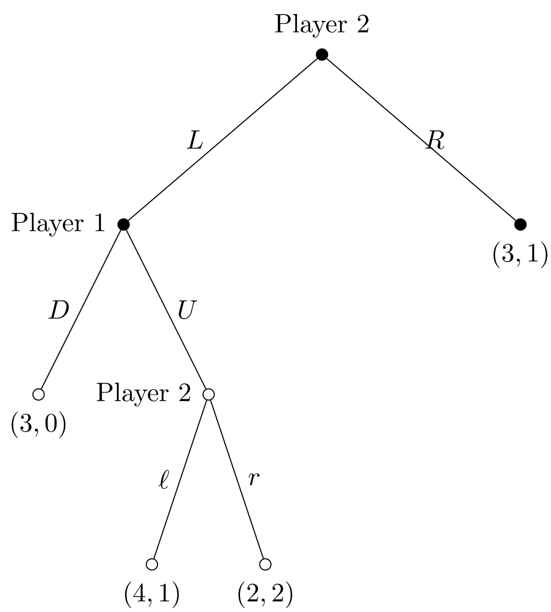

Figure 20: Extensive form game: illustration of Kuhn's theorem

Because of this result, we will only talk about behavior strategies from now onwards. The equivalence between one-shot deviation property and subgame perfect equilibrium (Theorem 20) continues to hold even with behavior strategies since we allowed for infinite action sets in Theorem 20. However, conceptually, a behavior strategy in an extensive form game is a complicated object - after all, players observe others playing a pure action and not the randomization. One way to think of it is that though players choose pure actions, the randomization device they use is public - this is referred to as *public randomization*. This issue is bypassed by the backward induction procedure because it is based on beliefs down a decision vertex.

Indifference. 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 21.

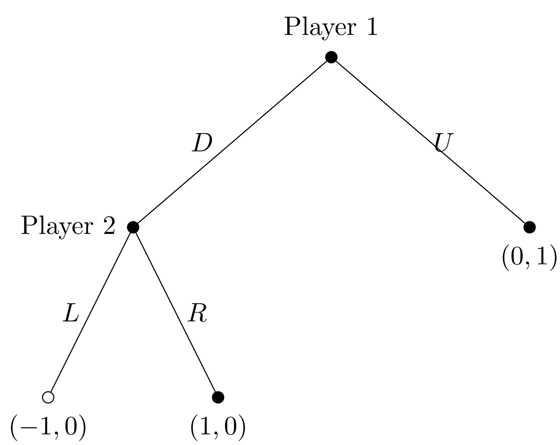

Figure 21: Backward induction with indifference

In the game in Figure 21, 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, L) and (D, R) 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.

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 application illustrates this point clearly.

20.1 Alternative Offers Bargaining

We now visit an 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 **22** *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 $(\frac{1}{1+\delta}, \frac{\delta}{(1+\delta)})$. ■

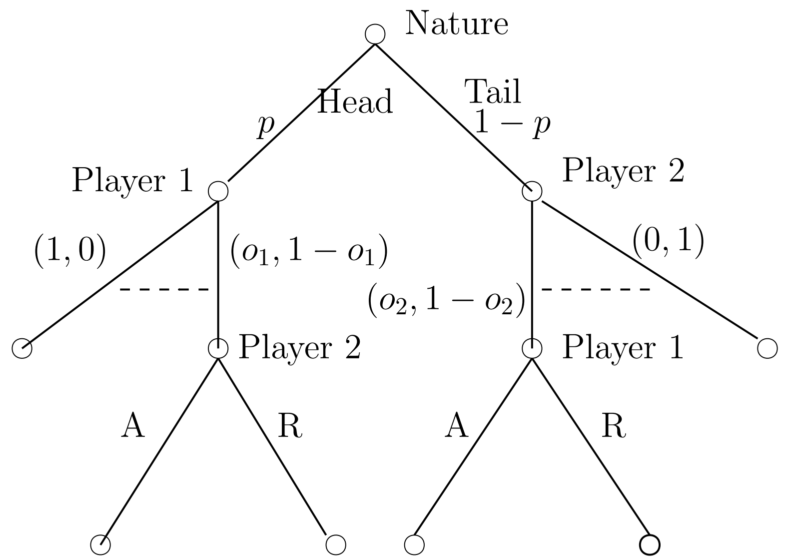

Figure 22: Nature moves in perfect information game

20.2 Nature moves

In many perfect information games, a decision vertex is owned by the *Nature*. The probability with which nature makes its moves is public. One can think of Nature's moves as a behavior strategy. Nature is non-strategic. We can denote Nature as Player 0 and the set of decision vertices where Nature is a decision maker as V_0. The probability distribution over actions at decision vertex $x \in V_0$ is given by $p_x \in \Delta A(x)$. In games with perfect information, this $\{p_x\}_{x \in V_0}$ is known to every player and its realization is observed by all the players. Everything we discussed so far can be modified by accommondating Nature's moves.

As an example, consider the bargaining game for one period. In the first period, a coin is tossed (whose probability of heads is p and tails is $(1-p)$). If the coin comes head, Player 1 proposes $(o_1, 1-o_1)$ and Player 2 accepts or rejects. If the coin comes tail, Player 2 proposes $(o_2, 1-o_2)$ and Player 1 accepts or rejects. Figure 22 shows the corresponding game tree.

21 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. Consider the following examples given below.

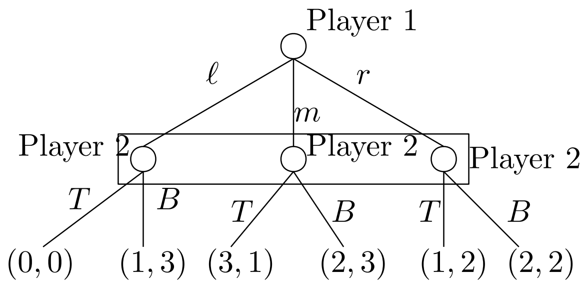

Figure 23: Strategic form game as an extensive form game

1. Strategic form games. Every strategic form game can be represented as an extensive form game of imperfect information. To see this consider a strategic form game of two players: $N = \{1, 2\}$. In the strategic form game, each player $i \in N$ chooses an action from his strategy S_i simultaneously. So, think of an extensive form game, where one of the players, say 1, moves first. However, the action of Player 1 is not observed by Player 2. This can be depicted by an extensive form game. Suppose $S_1 = \{\ell, m, r\}$ and $S_2 = \{T, B\}$. Then, the game is shown in Figure 23.

   Notice that when Player 2 takes her action in Figure 23, she does not know which decision vertex she is in - so her three decision vertices are bundled in one *information set*.

2. Bayesian games. In Bayesian games, there is a clear sequential nature of play. First, Nature draws the type of each player, but informs them *privately*. Hence, the action of Nature's move is observed to corresponding players only. We illustrate this with a simpler version of bilateral trading, where seller's cost c is known to both the buyer and the seller. At the beginning, buyer's value v is drawn from $\{v_L, v_H\}$ with probabilities π_L and π_H respectively. Then, the seller announces one of two prices: p_L and p_H. The buyer observes the prices and chooses either to accept or reject the offer. If accepted, trade happens at the announced price of the seller. Else, no trade happens.

   This game is shown as a game of imperfect information in Figure 24. Here, the imperfect information is generated by the private nature of Nature's move. Since the seller does not know the type of the buyer, he does not know which decision vertex he is in when he announces a price. Hence, his decision vertices are bundled in one information set.

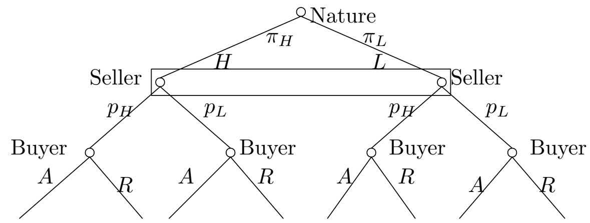

Figure 24: Bilateral trading (one-sided asymmetry) as an extensive form game

The idea of an information set is formalized below.

DEFINITION **37** *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. Another important specification is that we allow for moves by a player, who we denote by 0, called *Nature*. So, there will be a subset of decision vertices V_0, where Nature takes some actions. The probability of these actions are specified and known to all players in the game - Nature is not strategic.

Formally, an extensive form game of imperfect information can be defined similar to a game of perfect information with some minor modifications given as follows.

DEFINITION **38** *An* **extensive form game of imperfect information** *is*

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

where

- $\{U_i^j\}_{i \in N}^j$ *is a partition of V_i for each Player $i \in N$,*
- $A(U_i^j)$ *specifies the actions available at each information set U_i^j for Player i,*
- p_x *specifies a probability distribution at each of Nature's decision vertex $x \in V_0$ over his set of actions $A(x)$.*

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. Sugame starting with a decision vertex x will be denoted by $\Gamma(x)$, but it will have to obey an important restriction: if $y \in U_i^j$ belongs to $\Gamma(x)$, then all $y' \in U_i^j$ also belongs to $\Gamma(x)$. Hence, only decision vertices in information sets having *only* that decision vertex can be a root node of a subgame. For instance, in Figure 26, there is only one subgame of the whole game, that is the entire game itself. 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 $\Gamma(x)$.

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 14, 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.

21.1 Perfect Recall

Consider the following game in Figure 25. 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.

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

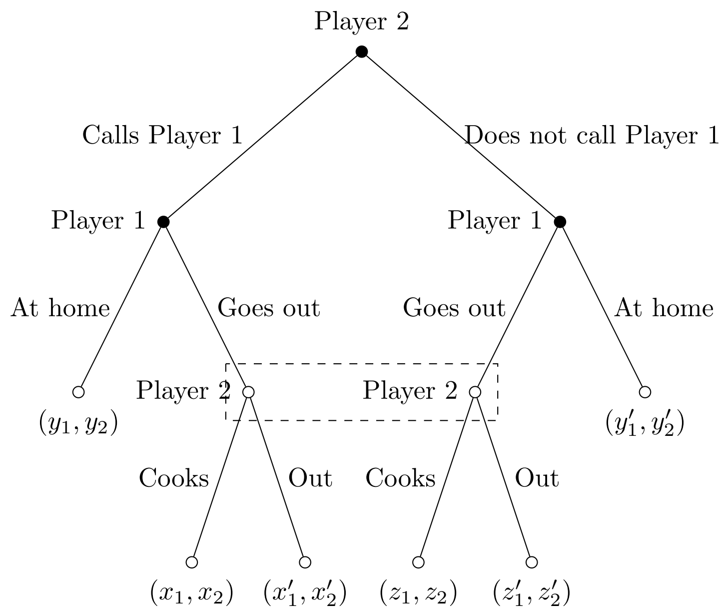

Figure 25: Extensive form game without perfect recall

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.

22 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 26. 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 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.

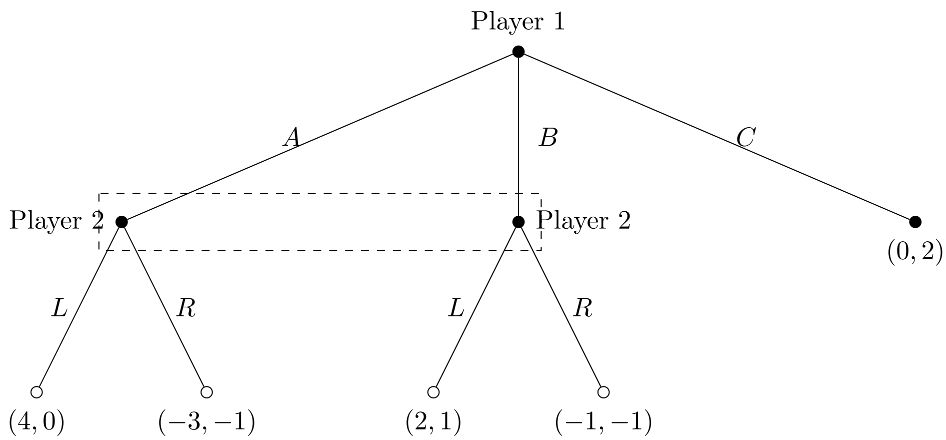

Figure 26: Imperfect Information

22.1 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 26. It is shown in Table 48.

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

Table 48: Reduced strategic form of the game in Figure 17

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 $\mu_i^j : U_i^j \to [0,1]$ for each j such that $\sum_{x \in U_i^j} \mu_i^j(x) = 1$ for all j. We write the collection of beliefs of Player i as μ_i: this specifies a probability distribution for each of his information sets.

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 **39** *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 $x \in U_i^j$,*

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

The next requirement of an equilibrium in extensive form game is sequential rationality. Sequential rationality is a one-shot-deviation optimality condition. It says that if all the players are following σ_{-i} and Player i is following σ_i at all information sets except U_i^j, then playing σ_i at U_i^j is optimal, and this must hold for every U_i^j. The definition captures a one-shot deviation principle which we do not discuss (but is true in this model also). In other words, Player i could have deviated from σ_i in all possible ways but we restrict her to deviate to only those strategies which differs from σ_i at only one information set. To define sequential rationality, we use the notation $u_i(\sigma|x)$ to denote the payoff of Player i by playing strategy σ in the "subgame" starting at decision vertex x. [6]

Definition **40** *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, where σ'_i differs from σ_i only at information set U_i^j, we have*

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

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

[6] We are abusing terminology a bit here. The decision vertex x may belong to a larger information set, and hence, the game starting at decision vertex is not a true subgame of this game.

A consequence of sequential rationality in finite stage extensive form games is that we can do backward induction. An equilibrium here in an imperfect information extensive form game 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 **41** *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 26, for every belief of Player 2, L is a strictly dominant action. Given this, Player 1 must play A irrespective of his beliefs. Hence, the unique PBE of this game is $(A, L, \mu_2(A) = 1)$. In general, a PBE does not allow players to play a strictly dominated action, while a Nash equilibrium does not preclude this off equilibrium path. A fact that we do not prove here but state is: every PBE is a Nash equilibrium.

22.2 Sequential 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 27. 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.

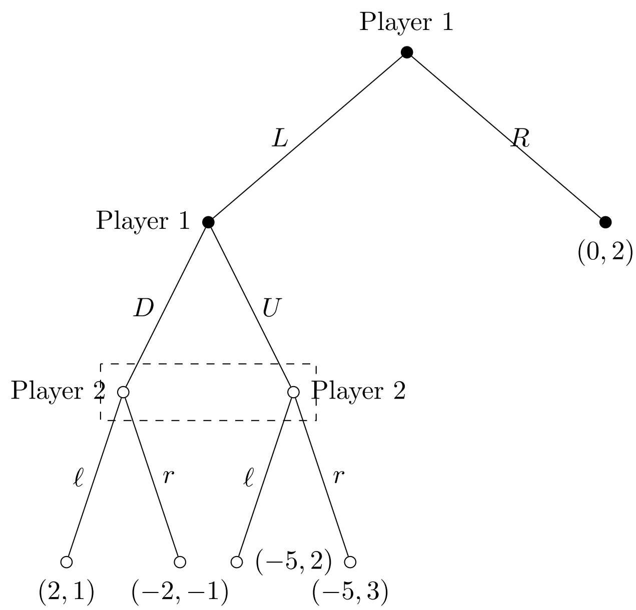

Figure 27: Problems with PBE

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 **42** *An assessment* (σ, μ) *is a* **sequential equilibrium** *if for every player* i

1. μ_i *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_i^k *are the unique Bayesian beliefs for* σ^k*, then* $\lim_k \mu_i^k = \mu_i$.

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

The new 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.

The following proposition says that every sequential equilibrium is also a perfect Bayesian equilibrium.

PROPOSITION **5** *If μ is consistent given σ, it is Bayesian given σ. Hence, every sequential equilibrium is also a perfect Bayesian equilibrium.*

Proof: VERY INFORMAL. For this, we pick an information set U_i^j of Player i which is reached with positive probability in σ. Bayesian belief says that for every $x \in U_i^j$,

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

Any perturbation σ^ϵ, will generate a belief μ^ϵ, which is computed by computing $P_{\sigma^\epsilon}(x)$ and $P_{\sigma^\epsilon}(U_i^j)$. As the perturbations approach zero, $P_{\sigma^\epsilon}(x)$ and $P_{\sigma^\epsilon}(U_i^j)$ approach $P_\sigma(x)$ and $P_\sigma(U_i^j)$ respectively - this happens because σ^ϵ approaches σ and the linear way in which probabilities are computed. So, as long as $P_\sigma(U_i^j)$ is non-zero, these limits give you $\mu_i^j(x)$. ■

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 **23** *Every sequential equilibrium is a subgame perfect equilibrium. Every completely mixed strategy Nash equilibrium is a sequential equilibrium.*

The second part of Theorem 23 follows trivially by taking the sequence of strategies same as the equilibrium strategy.

Let us now revisit the game in Figure 27. 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 (which we assumed throughout), then these equilibria always exist. Summarizing the various notions of equilibrium in extensive form games of imperfect information:

$$\text{Set of sequential equilibria} \subseteq \text{Set of subgame perfect equilibria} \subseteq \text{Set of Nash equilibria}$$

$$\text{Set of sequential equilibria} \subseteq \text{Set of perfect Bayesian equilibria} \subseteq \text{Set of Nash equilibria.}$$

As Example in Figure 27 showed, a perfect Bayesian equilibrium need not be a subgame perfect equilibrium.

22.3 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 28. Agent 1 has two types - High or Low, their probabilities are as shown in Figure 28. 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 28.

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

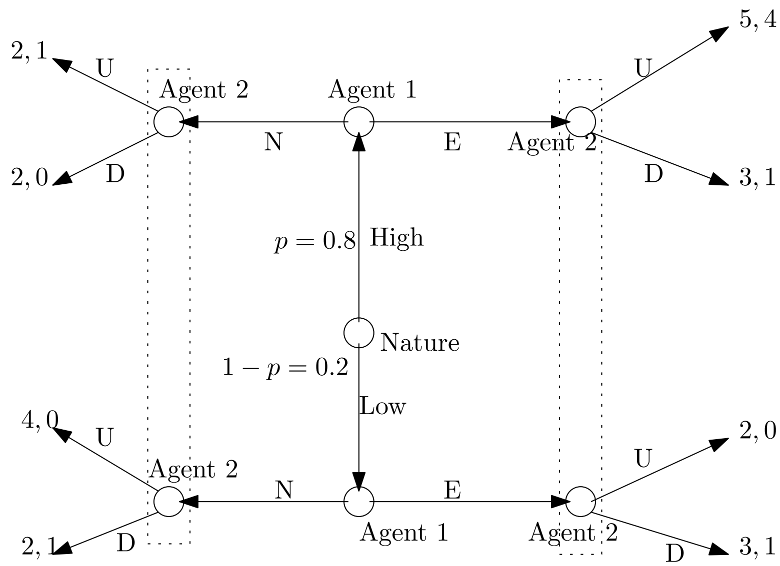

Figure 28: Signaling game

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 $\mu_L \in (0,1)$. 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. For $\mu_L = 0$, we can choose $\epsilon = (\epsilon')^2$ and this gives $\mu_L = \frac{p\epsilon}{p\epsilon + (1-p)\epsilon'} = \frac{1}{1+\frac{1}{4\epsilon'}}$. This converges to zero as $\epsilon' \to 0$.

- **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.

22.4 Example: Two period sale

This example illustrates perfect Bayesian equilibrium in a game with infinite set of actions. A seller has an object to sell. She has zero value for the object, which is common knowledge. There is a single buyer, whose value for the object is uniformly distributed in $[0,1]$ - this is common knowledge. However, the value for the buyer is known privately to him.

There are two periods. In the first period, the seller posts a price p_1 and the buyer chooses one of the actions: BUY (B) or WAIT (W). If the buyer chooses B, then the game ends with the seller getting a payoff of p_1 and the buyer of type v getting a payoff of $v - p_1$. If the buyer chooses W, then the game proceeds to period 2, where the seller posts a price p_2. The buyer can again choose one of the two actions: BUY (B) or WAIT (W). The game now ends. If the buyer chooses W, then both the players get a payoff of zero. But if the buyer chooses B, then the seller gets a payoff δp_2, whereas the buyer of type v gets a payoff of $\delta(v - p_2)$, where $\delta \in (0,1)$ is a common discount factor. The game is shown in Figure 29.

What is a strategy for a player in this game?

SELLER. The seller has exactly two information sets, one corresponding to each period. At each information set, he posts a price. So, his strategy in period 1 is $p_1 \in \mathbb{R}_+$ and in period 2 it is $p_2 \in \mathbb{R}_+$ having posted a price p_1 (note, in period 2, the seller has infinite number of information sets - one corresponding to each choice of p_1).

BUYER. The buyer has two sets of decision vertices: corresponding to period 1 and period 2. In period 1, his decision vertex depends on (a) his own type and (b) the price the seller

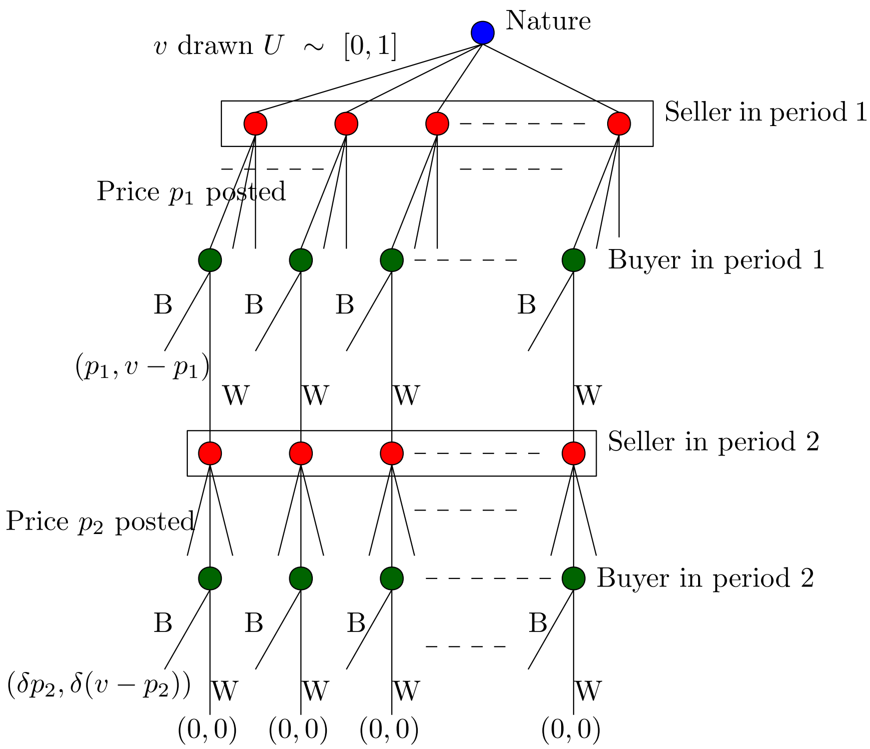

Figure 29: Sale across two periods

posts. So, a period 1 decision vertex can be described by (v, p_1) and for every (v, p_1), the buyer either chooses B or W. In period 2, the decision vertex is characterized by (v, p_1, p_2), and then the buyer either chooses B or W.

We now solve for a perfect Bayesian equilibrium of this game in steps.

Bayesian rationality of seller in period 1. This just requires that his beliefs must be same as nature probabilities: probability that seller is at a decision vertex corresponding to a buyer of type less than or equal to x (note: continuous distribution of types) is x (due to uniform distribution).

Sequential rationality of buyer in period 2. Sequential rationality in period 2 for buyer implies that a buyer of type v must buy if $v > p_2$ and wait if $v < p_2$.

Sequential rationality of buyer in period 1. Consider a buyer in period 1 who has value v and sees price p_1. Given strategy of the seller, Bayesian rational belief of seller, and his own sequentially rational action in period 2, if he finds sequentially rational to BUY, then every type $v' > v$ must also find it sequentially rational to BUY at price p_1. To see this, fix the strategy of the seller as p_1 and p_2 given p_1. The payoff of a buyer of type p_1 by buying today is $v - p_1$ and waiting for next period is $\max(\delta(v - p_2), 0)$. If $v - p_1 > \max(\delta(v - p_2), 0)$, then for all $v' > v$, we also have $v' - p_1 > \max(\delta(v' - p_2), 0)$. Similarly, if $v - p_1 < \max(\delta(v - p_2), 0)$, then for all $v' < v$, we also have $v' - p_1 < \max(\delta(v' - p_2), 0)$. This suggests a **cutoff-action** to be optimal in period 1 for the buyer. For every price p_1, there is a cutoff value $v(p_1)$ such that all buyer types above it BUY and all buyer types below it WAIT.

BAYESIAN RATIONALITY OF SELLER IN PERIOD 2. Given the strategy of the buyer, the seller in period 2 knows that only a buyer with value $v < v(p_1)$ will be active in period 2. Hence, the conditional probability of being at decision vertex where value of buyer is less than or equal to x (we compute cdf because there are infinite number of decision vertices in this information set) is given by (using conditional uniform distribution):

$$\frac{x}{v(p_1)}.$$

Further, such a buyer chooses BUY in period 2 if $v > p_2$. Hence, expected payoff of seller by setting a price p_2 in period 2 given a price p_1 in period 1 is given by

$$p_2 \frac{v(p_1) - p_2}{v(p_1)}.$$

SEQUENTIAL RATIONALITY OF SELLER IN PERIOD 2. Sequential rationality of seller in period 2 who has already posted a price p_1 is to maximize her expected payoff given his Bayesian rational beliefs. This leads to maximizing

$$p_2 \frac{v(p_1) - p_2}{v(p_1)},$$

over all $p_2 \in \mathbb{R}_+$. The maximum of this expression happens at

$$p_2 = \frac{1}{2} v(p_1).$$

SEQUENTIAL RATIONALITY OF BUYER IN PERIOD 1 (AGAIN). Having computed p_2 as a function of $v(p_1)$, we can now be more precise about buyer's action in period 1. We know that the **cutoff** type will be indifferent between BUY and WAIT in period 1. Hence,

$$v(p_1) - p_1 = \delta(v(p_1) - p_2) = \delta\big(v(p_1) - \frac{1}{2}v(p_1)\big) = \frac{\delta}{2}v(p_1).$$

This gives us

$$v(p_1) = \frac{1}{1 - \frac{\delta}{2}}p_1.$$

SEQUENTIAL RATIONALITY OF SELLER IN PERIOD 1. Finally, sequential rationality of the seller must require that the seller must maximize her expected payoff (given the strategy of buyer and his beliefs in period 2) in period 1. His expected payoff by posting a price p_1 is (denoting $1 - \frac{1}{2}\delta = K$ below):

$$\begin{aligned}(1 - v(p_1))p_1 + v(p_1)\Big[\delta p_2 \frac{(v(p_1) - p_2)}{v(p_1)}\Big] &= (1 - v(p_1))p_1 + \delta\frac{1}{2}v(p_1)\frac{1}{2}v(p_1) \\ &= (1 - \frac{1}{K}p_1)p_1 + \frac{\delta}{4}\frac{1}{K^2}(p_1)^2.\end{aligned}$$

Taking the first order condition with respect to p_1 and setting it equal to zero, we get

$$1 - \frac{2}{K}p_1 + \frac{\delta}{4K^2}2p_1 = 0.$$

This gives us

$$p_1 = \frac{2K^2}{4K - \delta} = \frac{1}{2}\frac{(1 - \frac{\delta}{2})^2}{(1 - \frac{3\delta}{4})}.$$

This also gives us the complete specification of the equilibrium:

$$p_1 = \frac{1}{2}\frac{(1 - \frac{\delta}{2})^2}{(1 - \frac{3\delta}{4})}; v(p_1) = \frac{1}{2}\frac{(1 - \frac{\delta}{2})}{(1 - \frac{3\delta}{4})}; p_2 = \frac{1}{4}\frac{(1 - \frac{\delta}{2})}{(1 - \frac{3\delta}{4})};$$

supplemented by beliefs for seller 1: in period 1 information set, her belief is the same as Nature's probability; in period 2 information set, her belief of being at a vertex corresponding to buyer type less than or equal to x and price p_1 is 0 if $x > v(p_1)$ and $\frac{v(p_1) - x}{v(p_1)}$ if $x < v(p_1)$.

For a class of $\delta \in (0, 1)$, $v(p_1) > 0$ if $p_1 > 0$. So, if $p_1 > 0$, then every buyer with value less than $v(p_1)$ will reach the the second period information set. If $p_1 = 0$, then

$v(p_1) = 0$, then the only buyer who reaches the second period information set are zero value buyer. But if we consider a perturbation of this strategy, where buyer with value value v BUYS with probability $1 - \epsilon$ and WAITS with probability ϵ if $v \geq v(p)$ and WAITS with probability $1 - \epsilon$ and BUYS with probability ϵ otherwise, then we reach **all** information sets with positive measure probability. In particular, even at $p_1 = 0$, we may have some buyers with probability greater than zero with positive probability. The limit of the beliefs induced by these strategies can be shown to be the beliefs induced by cutoff strategies. This in turn will show sequential equilibrium.

A Supermodular Games

The concavity assumption made in Theorem 3 does not hold in many games. We now discuss a class of games where we introduce 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.

A.1 Lattices and complete lattices

Lattices are abstract mathematical objects that are useful to define supermodular games. The starting point of a lattice is an abstract set X - for most part of this section, you can assume X is a subset of $\mathbb{R}^k$ for some k. We are given **partial order** (an reflexive, anti-symmetric, and transitive but not necessarily complete) $\leq$ on X. If $X \subseteq \mathbb{R}^k$, then $\leq$ is the usual relation: $x \leq y$ if $x_i \leq y_i$ for each $i \in \{1, \ldots, k\}$.

Given any subset $S \subseteq X$, an **upper bound** of S is an element $x \in X$ such that $y \leq x$ for all $y \in S$. Similarly, a **lower bound** of S is an element $x \in X$ such that $x \leq y$ for all $y \in S$. The **least upper bound** of S is an element $x \in X$ such that for every upper bound x' of S, we have $x \leq x'$. Similarly, the **greatest lower bound** of S is an element $x \in X$ such that for every lower bound x' of S, we have $x' \leq x$. Many a times an upper bound and a lower bound of a subset may not exist. If S is a *rectangle* in $\mathbb{R}^2$, then the two corner points define the greatest lower bound and least upper bound for the entire rectangle. Figure 30 illustrates this. Let $S = \{(x_1, y_1), (x_2, y_2)\}$ as shown and $X = \mathbb{R}^2$. Then, the least upper bound is $(\max(x_1, x_2), \max(y_1, y_2))$ and the greatest lower bound is $(\min(x_1, x_2), \min(y_1, y_2))$.

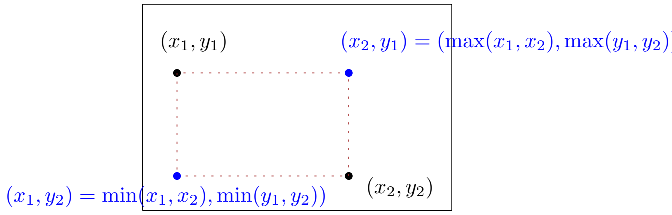

Figure 30: A lattice

Definition 43 *The pair $(X, \leq)$ is a* **lattice** *if for every $\{x, y\} \subseteq X$ the greatest lower bound and the least upper bound of $\{x, y\}$ exist in X.*

As discussed, when $X = \mathbb{R}^k$, then the $\max(x, y)$ and $\min(x, y)$ provide the least upper bound and greatest lower bound respectively, and since they lie in X, it is a lattice. Here are some more examples of lattices which are subsets of $\mathbb{R}^2$:

$$X := \{(0, 0), (1, 0), (0, 1), (1, 1)\}$$
$$X := \{(0, 0), (1, 0), (0, 1), (100, 100)\}$$

In the first lattice, the least upper bounds and greatest upper bounds are the min and max points. But in the second lattice, the least upper bound of $\{(1, 0), (0, 1)\}$ is $(100, 100)$. If we remove $(1, 1)$ from the first X, then the resulting set has no upper bound for $\{(1, 0), (0, 1)\}$, and hence, is not a lattice.

We will be interested in complete lattices.

Definition 44 *The pair $(X, \leq)$ is a* **complete lattice** *if for every $S \subseteq X$, the greatest lower bound and the least upper bound of S exist in X.*

Not every lattice is complete. For instance, $X = (0, 1)$ is a lattice since for any pair of points $x, y \in X$, we can find the greatest lower bound and least upper bound, but not for the entire set X. Clearly, if $X \subseteq \mathbb{R}^k$ and X is a complete lattice, then it has to be compact.

We will be interested in functions on a lattice $(X, \leq)$. Let $f : X \to X$. We say f is **monotone** if for every $x, y \in X$ with $x \leq y$, we have $f(x) \leq f(y)$. As we know, non-monotone functions need not have a fixed point. For instance, let $X = [0, 1]$ and define $f : X \to X$ as follows.

$$f(x) = \begin{cases} 1 & \text{for } x \in [0, \frac{1}{2}) \\ 0 & \text{for } x \in [\frac{1}{2}, 1] \end{cases}$$

This f has no fixed points. It seems to be that this is caused by discontinuity of f. But consider the following f:

$$f(x) = \begin{cases} 0 & \text{for } x \in [0, \frac{1}{2}) \\ 1 & \text{for } x \in [\frac{1}{2}, 1] \end{cases}$$

This is also a discontinuous function but now has a fixed point at 0. The difference is monotonicity. Tarski's fixed point theorem formalizes this.

THEOREM **24 (Tarski's Fixed Point Theorem)** *Let $(X, \leq)$ be a complete lattice. If $f : X \to X$ is monotone, then f has a fixed point. Moreover, if P is the set of fixed points of f, then there is $\underline{x}, \bar{x} \in P$ such that for all $x \in P$,*

$$\underline{x} \leq x \leq \bar{x}.$$

Proof: Define the following set:

$$\bar{X} := \{x \in X : x \leq f(x)\}.$$

Since $(X, \leq)$ is a complete lattice, there is a greatest lower bound x_* of X. Hence, $f(x_*) \geq x_*$, which implies that $x_* \in \bar{X}$. So, $\bar{X}$ is non-empty. Since $(X, \leq)$ is a complete lattice, there is a least upper bound $\bar{x} \in X$ of $\bar{X}$.

We argue that $\bar{x} \in \bar{X}$. For every $x \in \bar{X}$, we know $x \leq \bar{x}$. Monotonicity implies $f(x) \leq f(\bar{x})$. But $x \in \bar{X}$ implies $x \leq f(x) \leq f(\bar{x})$. Hence, $f(\bar{x})$ is an upper bound on $\bar{X}$. Since $\bar{x}$ is the least upper bound, we get $\bar{x} \leq f(\bar{x})$, which implies that $\bar{x} \in \bar{X}$.

We now show that $\bar{x}$ is a fixed point of f. To do so, first note that for every $x \in \bar{X}$, we have $x \leq f(x)$, and monotonicity implies that $f(x) \leq f(f(x))$. Hence, $x \in \bar{X}$ implies $f(x) \in \bar{X}$. This implies that $f(\bar{x}) \in \bar{X}$. But if $\bar{x} < f(\bar{x})$, then $\bar{x}$ is not an upper bound for $\bar{X}$. Hence, it must be that $\bar{x} = f(\bar{x})$.

We can also define the following set:

$$\underline{X} := \{x \in X : f(x) \leq x\}.$$

The set $\underline{X}$ is non-empty because the the least upper bound x^* of X must satisfy $f(x^*) \leq x^*$. We now consider the greatest lower bound of $\underline{X}$, and denote it as $\underline{x}$. For every $x \in \underline{X}$, we have $\underline{x} \leq x$. Hence, $f(\underline{x}) \leq f(x)$. But $x \in \underline{X}$ implies that $f(\underline{x}) \leq f(x) \leq x$. Hence, $f(\underline{x})$ is a lower bound for $\underline{X}$. Since $\underline{x}$ is the greatest lower bound, we must have $f(\underline{(x)}) \leq \underline{x}$, which implies that $\underline{x} \in \underline{X}$.

As earlier, for every $x \in \underline{X}$, we have $f(x) \leq x$. Hence, $f(f(x)) \leq f(x)$ for all $x \in \underline{X}$. As a result, for all $x \in \underline{X}$, we have $f(x) \in \underline{X}$. But, $f(\underline{x}) \leq \underline{x}$. If $f(\underline{x}) < \underline{x}$, then $\underline{x}$ is not a lower bound for $\underline{X}$ since $f(\underline{x}) \in \underline{X}$. Hence, $f(\underline{x}) = \underline{x}$.

Thus, we have discovered two fixed points of f: $\bar{x} \in \bar{X}$ and $\underline{x} \in \underline{X}$. Pick an arbitrary fixed point $x^* \in X$ of f: $x^* = f(x^*)$. By definition, $x^* \in \bar{X} \cap \underline{X}$. Since $\bar{x} \in \bar{X}$ is the least upper bound of $\bar{X}$ and $\underline{x} \in \underline{X}$ is the greatest lower bound of $\underline{X}$, we have,

$$\underline{x} \leq x^* \leq \bar{x}.$$

■

A.2 Supermodularity and comparative statics

We will now focus attention on lattices defined on subsets of Eucliean spaces. Hence, the underlying relation is the standard $\leq$ (or $\geq$). We will avoid explicit mention of this. Our first definition is a definition of increasing differences across two lattices.

Definition **45** *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. For any lattice, denote the least upper bound of a pair of points x and y as $x \vee y$ and the greatest lower bound of a pair of points x and y as $x \wedge y$.

Definition **46** *Let $X \subseteq \mathbb{R}^K$ be a lattice. A function $f : X \subseteq \mathbb{R}$ 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 **25 (Topkis Monotone Comparative Statics)** *Let $X \subseteq \mathbb{R}^K$ be a complete lattice and $T \subseteq \mathbb{R}^L$ be a lattice. Suppose $f : X \times Y \to \mathbb{R}$ 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 complete lattice.*

2. *Let* $\bar{x}^*(t)$ *and* $\underline{x}^*(t)$ *be the least upper bound and the greatest lower bound of* $x^*(t)$ *at each* t. *Then,* $\bar{x}^*(t), \underline{x}^*(t) \in x^*(t)$.

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

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

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

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

Proof: We skip (1) and (2)'s proof. For (3), pick 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 $f(x \vee x', t) \leq f(x \wedge x', t)$.

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 **47** *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 (and hence, complete) 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 3). 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 **26** *Every supermodular game has a pure strategy Nash equilibrium.*

Proof: Pick any strategy profile s. For every $i \in N$ and $B_i(s_{-i}) = \{s_i : u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}) \;\forall\; s'_i\}$. Since S_i and S_{-i} are complete lattices, by Theorem 25, $B_i(s_{-i})$ is a non-empty complete lattice. Now, we define $\bar{B}_i(s_{-i})$ as the lowest upper bound of $B_i(s_{-i})$ - note that this is a strategy in $B_i(S_{-i})$ due to Theorem 25. Now, we can define for every strategy profile s,

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

Hence, $\bar{B} : S_1 \times \ldots \times S_n \rightarrow S_1 \times \ldots \times S_n$. By Theorem 25, 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 complete lattice $S_1 \times \ldots \times S_n$. By Theorem 24, a fixed point of $\bar{B}$ exists, and it must be a Nash equilibrium. ■

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 26, 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 3. 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.

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

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

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

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

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 $s_{-i} \in S_{-i}^{k-1}$, 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 $\bar{s}'_i > s_i^k > \bar{s}_i$. 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 2). 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 $s^0 \equiv (s_1^0, \ldots, s_n^0)$ 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 this idea 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 $S_i^0 = [0,1]$. 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 $S_i^1 = [0, \frac{1}{2}]$ and $s_i^1 = \frac{1}{2}$ 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 $S_i^2 = [0, \frac{3}{8}]$ and $s_i^2 = \frac{3}{8}$. 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

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

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, $S_i^1 = [\frac{1}{4}, 1]$ and $s_i^1 = \frac{1}{4}$ 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

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

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.

B 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 **48** *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 $\sum_{a' \in A} \lambda_{a'} = 1$ and for every $i \in N$*

$$v_i = \sum_{a'} \lambda_{a'} u_i(a').$$

Let V denote the set of all feasible profiles. 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* (not necesssarily correlated equilibrium). Correlated strategies require a public randomization device. So, achieving payoffs in $Conv(V)$ requires public randomization. This requires mixed/correlated strategies. A mixed 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 49: A Stage game

Consider the game in Table 49. We draw its feasible payoff vector in Figure 31. The minmax values of both the players are also shown in Figure 31. 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

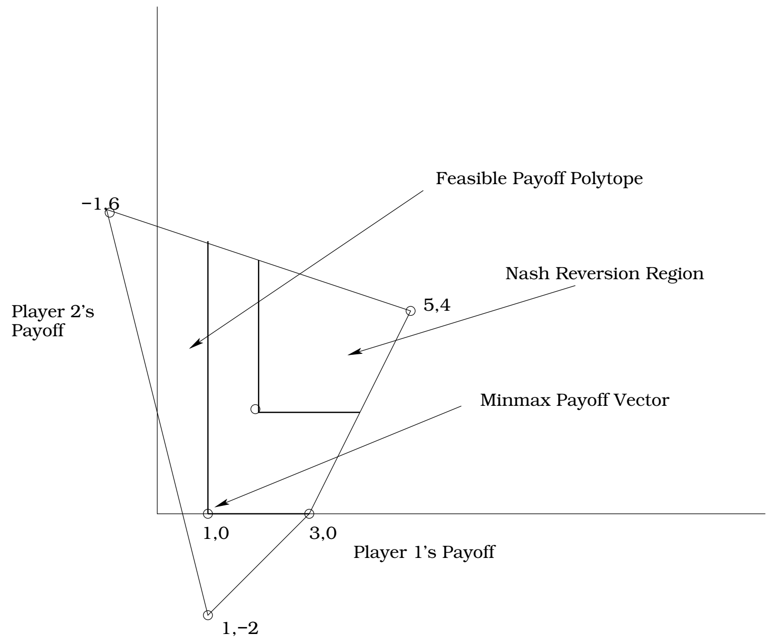

Figure 31: Feasible Payoff Vectors and Minmax Values

(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 16 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 and Maskin.

Theorem **27** *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.*

Proof: We will only give a partial proof here. We will assume that $Conv(V)$ has dimension n. Take a payoff vector v which is strictly enforceable such that there exists a pure action profile a with $u_i(a) = v_i$ for all $i \in N$ (a pure action profile generates the payoff in the stage game - this can be dispensed away if we allow for correlated action profiles with publicly observable randomization). We will show that v can be a payoff profile in a subgame perfect equilibrium of the infinitely repeated game.

Take a payoff profile v' in the interior of the set of strictly enforceable payoff vectors and $\epsilon > 0$ but sufficiently small such that for every $i \in N$,

$\underline{v}_i < v'_i < v_i$

$v^i := (v'_1 + \epsilon, \ldots, v'_{i-1} + \epsilon, v'_i, v'_{i+1} + \epsilon, \ldots, v'_n + \epsilon)$ is a strictly enforceable payoff vector.

The full dimension assumption ensures that such v' and ϵ can be chosen. The payoff vector v^i is a ϵ-reward for everyone except Player i with respect to strictly enforceable payoff vector v'. Assume that there is a pure action profile $\bar{a}^i$ which generates the payoff v^i (this can be dispensed away if we allow for correlated action profiles with publicly observable randomization).

Now, choose a T such that

$$T > \frac{\max_{a' \in A} u_i(a') - \min_{a' \in A} u_i(a')}{v'_i - \underline{v}_i}. \tag{23}$$

The strategy profile is a little complicated to describe with many states. As in earlier proofs, at every state, if more than one player deviates from recommendation, we do not need to worry - this will never happen in equilibrium - we can just reset the game by coming to the normal state. So, we ignore those transitions from here on.

- Normal state. This is the initial state. It recommends playing a_j to Player j. If Player i does not play a_i in normal state, then we enter punishment$-i$ state. Else, we stay in normal state.
- Punishment$-i$ state. In this state, each Player j plays the minmax action a^i_j of Player j. Once we enter punishment$-i$ and everyone has been playing a^i, we stay in this state for exactly T periods, and then enter post-punishment$-i$ state. If some player j does not play a^i_j in punishment$-i$ state, we enter punishment$-j$ state.
- post-punishment$-i$ state. In this state, each player j plays $\bar{a}^i_j$. If Player j does not play $\bar{a}^i_j$ in this state, we enter punishment$-j$ state. Else, we stay in this state.

It is a matter of routine checking using one-shot deviation principle and Inequality (23 that the strategy is a subgame perfect equilibrium (though calculations are messy, the idea here is that we reward players to punish another player by giving them ϵ extra, and that allows them to not deviate from punishment). ■