Microeconomics (ECON204B)

Allegra Saggese

2024-25 UCSC Economics PhD

Winter 2025

Kristian Miguel Lopez Vargas

Last updated: February 6, 2026

Contents

1 Game theory

MWG: 7, JR: 7.1,7.3

1.1 Structure

A **game** is a situation where 2+ agents face strategic independence in their payoffs, where player i's payoff depend on her actions and the actions of others. Normal (or strategic) form games model environments where players act simultaneously or independently.

In general, a normal form game is characterized by the following elements:

- *Players*: The participants involved in the game.
- *Rules*: These define the order of moves, the available actions at each stage, and what each player knows when they make their decisions.
- *Outcomes*: The result of the game for every possible combination of actions taken by the players.
- *Payoffs*: The preferences of players over the set of possible outcomes. We describe a player's preferences by a utility function that assigns a utility level for each possible outcome.[1]
- *Information Structure*: Describes what each player knows at different stages of the game, which can affect their strategies and choices.

1.2 Definitions

- **Game tree:** A unique, connected path of branches that describes the choices and payoffs of an extended form game. Each branch connects a pair of nodes, and can be identified by the two nodes it connects. Every node in a game represents a history of the game (all points up until the current node).
  - *Initial node:* The initial position where a decision is made, also known as the first move or first point
  - *Decision node:* each point where one, two or more actors make a move
  - *Terminal node:* where a game ends, after a decision is made and the branch of the game leads to a payoff, and not another decision node
- **Subgames:** For use in *sec 3.3*, A subgame is a subset of Γ_E with the following properties:
  - A subgame always begins with an information set containing a single node, it contains all decision (and terminal) nodes that are successors, and it contains only the successive nodes
  - If a decision node, x is in the subgame, then every $x' \in H(x)$ is also in this subgame (i.e. there are no broken information sets)

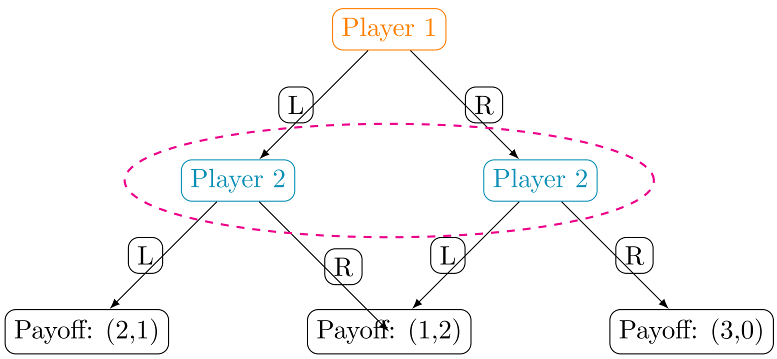

From the above game tree:

1. **Information set:** Where Player 2's information set consists of two nodes. This is because in a simultaneous move game, we assume Player 2 does not yet know what Player 1 has selected.

[1]It is common to refer to the player's utility function as her payoff function and the utility level as her payoff.

2. **Player 1:** Player 1 starts at the initial node, where they make a decision L or R and Player 2 responds. Given this is a simultaneous game, its possible to reverse the players in these positions.

3. **Player 2:** These nodes represent the decisions possible to Player 2. Note that both actions (L, R) are available at both nodes. As it is required the actions from each node in an information set are the same.

- **Finite game:** A game is called finite if each player has a finite strategy set. Conversely, a game is infinite if at least one player has an infinite strategy set.
- **Zero sum game:** A type of game where, when one player wins, the other loses. This is a game of pure conflict.
- **Information set:** A subset of a player's decision nodes. The information set represents the knowledge to a player that they are standing on a node (or set of nodes). The actions available to a player within an information set must be the same, although payoffs may vary.
  - In notation: $\mathcal{H}_i$ is the collection of player i's information sets, each element denoted with H
  - $C(H) \subset A$, which is the set of actions available at information set, H (see actions - 1.2).
- **Perfect recall:** At any point in the game, a player has perfect recall when they are able to remember all previous moves, in the order taken, and therefore know where they are on the tree.
- **Perfect information:** Each player knows exactly what node they are on (i.e. there is a single node in the information set). A game with perfect information will have an information set of a decision node, as opposed to a game of imperfect information.
- **Common knowledge:** Games have common knowledge such that all players know the structure of the game, all players know the other players are aware of the structure, and each player knows this is true of the other players' knowledge.
- **Strategy vs. Action:** A *strategy* is a decision rule, or complete contingent plan that specifies how a player will act in every possible situation, as a best (or not a never best) response. For example, the contingent plan will contain responses to all potential moves by opposing players. While an *action* is any possible available moves or decisions by a player, regardless of the strategic value.
  - A *set of strategies* is defined as n^m where n is the number of decisions and m is the number of nodes within each information set (m must be the same for all nodes in an information set).
  - A is the set of all possible **actions** in a game
  - s_i: $\mathcal{H}_i \rightarrow A$ subject to $s_i(H) \in C(H)$ for all $H \in \mathcal{H}_i$. So s_i is the functional form which represents a strategy for a player.
  - $s = (s_1, ...s_n)$ is a profile of strategy choices for a player, where each strategy, s has an induced outcome.
  - **Example:** with a game of rock-paper-scissor: Each player has the following pure strategies:

    $$S_1 = S_2 = \{\text{Rock (R)}, \text{Paper (P)}, \text{Scissors (S)}\}.$$

    The payoff matrix for Player 1 is as follows:

$P_1 \backslash P_2$	Rock (R)	Paper (P)	Scissors (S)
Rock (R)	$(0,0)$	$(-1,1)$	$(1,-1)$
Paper (P)	$(1,-1)$	$(0,0)$	$(-1,1)$
Scissors (S)	$(-1,1)$	$(1,-1)$	$(0,0)$

  - **Strategies:**
    * Player 1 pure strategy, plays Rock: $s_1 = \text{Rock}$.
    * Player 2 pure strategy plays Scissors: $s_2 = \text{Scissors}$.

* The outcome would be $(s_1, s_2) = (\text{Rock}, \text{Scissors})$, and the payoffs would be:

$$u_1(\text{Rock}, \text{Scissors}) = 1, \quad u_2(\text{Rock}, \text{Scissors}) = -1.$$

* Player 1 could play a mixed strategy:

$$\sigma_1 = (p_R, p_P, p_S) = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right),$$

meaning Player 1 chooses Rock, Paper, and Scissors with equal probabilities.

* Player 2 could play a mixed strategy:

$$\sigma_2 = (q_R, q_P, q_S) = (0.5, 0.25, 0.25),$$

meaning Player 2 chooses Rock with 50% probability, and Paper and Scissors with 25% probability each.

* This would result in the expected payoffs:

$$E[u_1] = \sum_{s_1 \in S_1} \sum_{s_2 \in S_2} \sigma_1(s_1)\sigma_2(s_2)u_1(s_1, s_2),$$

Similarly, Player 2's expected payoff is also:

$$E[u_2] = 0. \text{ due to simplification.}$$

1.3 Representation

1.3.1 Normal form games

Normal form games (NFG) are condensed versions of a game. The expression for the NFG only captures the strategy and associated payoffs. There is no additional information provided in the mathematical expression.

Definition: A normal form game is $(I, \{S_i\}_{i \in I}, \{u_i\}_{i \in I})$, where:

- I is a finite set of players, $I = \{1, 2, \ldots, n\}$.
- $\{S_i\}_{i \in I}$ denotes the collection of strategy sets, where S_i denotes the set of all strategies available to player i (see 1.2). The strategy, again, is the complete contingency plan for a player.
  - A strategy must be **complete**, meaning the strategy will contain actions at information sets that are never reached in a game, but are potential.
  - A strategy will also include an action plan which may become irrelevant based on the player's own strategy (i.e. a strategy will contain actions that are not the best move, and never taken, but nonetheless are included).
- $\{u_i\}_{i \in I}$ the von Neumann-Morgenstern utility levels. The utility represents the collection of payoff functions, where each $u_i : S_1 \times S_2 \times \cdots \times S_n \to \mathbb{R}$ assigns a real value payout to the player i based on the strategy profile.

1.3.2 Extended form games

Extended (or extensive) form games (EFG) are important for depicting games where the rules differ than purely static, simultaneous move games. For example, EFGs can capture dynamic games (see 3), where the order that players move is unique, and the information available to them is not always perfect. EFGs are often depicted by game trees.

Formal representation of an extended form game (EFG) is:

$$\Gamma_E = \{H, A, I, \rho(\cdot), \alpha(\cdot), \mathcal{H}, H(\cdot), \iota(\cdot), p(\cdot), u\}$$

From this formal definition, we can characterize the components:

- $\mathcal{H}$ is a finite set of nodes
  - Where $H(\cdot)$ is the function that assigns the decision nodes, x to an information set, $\mathcal{H}$, such that $H(x) \in \mathcal{H}$
  - $\mathcal{H}$ is a *partition* of X, the entirety of the set of decision nodes
  - It is critical to recall all decision nodes assigned to the same information set must have the same set of choices (see 1.3.2, so $C(x) = C(x')$ if $H(x) = H(x')$
- A is the finite set of actions (1.2)
- I is the set of players
- $p(\cdot)$ is a function, $\mathcal{H} \to \{\mathcal{H} \cup \emptyset$ specifying a single, immediate predecessor of each node, x. So, for every $x \in \mathcal{H}$ $p(x)$is non-empty, except for the case of the initial node, x_0
- **Successor**: $s(\cdot)$ is then the successor node of x, where $s(x) = p^{-1}(x)$. It is required that the sets of $p(x)$ and $s(x)$ are disjoint, to ensure a tree structure.
- α: $\alpha : X\backslash\{x_0\} \to A$ where α is the function of the set of actions that results in the successor nodes. The function then gives (or represents) the actions that lead any non-initial node, x from its predecessor. It satisfies the property $x'.x'' \in s(x)$ and $x' \neq x'' \implies \alpha(x') \neq \alpha(x'')$.
- **Choices**: represented $C(x)$ are the available moves at a decision node, represented as:

$$C(x) = \big\{a \in A : a = \alpha(x') \text{ for all } x' \in s(x)\big\}$$

- $\iota(\cdot)$: This is a function such that $\mathcal{H} \to \{0, 1, ...I\}$ it assigns information sets within $\mathcal{H}$ to a player, who moves at the decision nodes in the set.
- $\rho(\cdot)$: This is a function such that $\mathcal{H}_0 x A \to [0, 1]$, which assigns probabilities to actions at information sets, where nature moves. The function satisfies

$$\rho(H, a) = 0 \text{ if } a \notin C(H) \text{ and } \sum_{a \in C(H)} \rho(H, a) = 1 \text{ for all } H \in \mathcal{H}_0$$

- **Payoff function**: $u = \{u_1(\cdot), u_2(\cdot), ...u_I(\cdot)\}$ is a collection of payoff functions which assign a utility representation to each player at each terminal node that can be reached. Therefore, each player has a utility, u_i such that $T \to \mathbb{R}$, where u_i has a Bernoulli Distribution (see 6.2.3).

For an extended form game to be operate, the following three definitions for finiteness need to hold:

1. Players have a finite number of possible actions, this limits the number of potential nodes in an information set. So we rule out a game with $[a, b] \subset \mathbb{R}$ possible actions.
2. The game must end in a finite number of moves, implying the number of decision nodes is finite.
3. We also assume a finite number of players.

   Nash equilibrium of an extensive form game is the same in definition, but the notion of *strategy* changes. Where a strategy is now defined for *every information set.*

1.4 Problem types

- **Matching pennies:** The Matching Pennies game is a two-player, zero-sum game where both players simultaneously choose either Heads (H) or Tails (T). If their choices match, Player 1 wins; if they differ, Player 2 wins.

	H_2	T_2
H_1	$(1, -1)$	$(-1, 1)$
T_1	$(-1, 1)$	$(1, -1)$

  Where, payoffs are given as *(Player 1, Player 2)*.

- (H_1, T_1) are Player 1's choices.
- (H_2, T_2) are Player 2's choices.

This game has no pure strategy Nash equilibrium (discussed in 2.5), but there exists a *mixed strategy equilibrium.* where each player chooses H or T with equal probability ($p = 0.5$).

- **Prisoner's Dilemma:** The Prisoner's Dilemma is a two-player, non-zero-sum game where each prisoner must decide whether to cooperate (C) or defect (D) without knowing what the other will do. The payoffs follow:
  - Mutual Cooperation: Both prisoners receive a moderate sentence.
  - Mutual Defection: Both get a severe sentence.
  - One defects while the other cooperates: The defector goes free while the cooperator gets the harshest sentence.

	C_2	D_2
C_1	$(-2, -2)$	$(-5, 0)$
D_1	$(0, -5)$	$(-4, -4)$

Where, payoffs are given as (Prisoner 1, Prisoner 2):

  - (C_1, D_1) are Player 1's choices (Cooperate or Defect).
  - (C_2, D_2) are Player 2's choices.

The *dominant strategy* for each player is to defect, D, leading to a Nash equilibrium at (D, D), even though (C, C) would be socially optimal.

- **Hotelling's game**: Hotelling's Location Model describes a competition between two firms choosing locations along a linear market (e.g., a street, a beach, or a town). Consumers are uniformly distributed across this space, and they purchase from the nearest firm. Firms strategically place themselves to capture market share, and the goal is to maximize profits, which depend on how many customers choose each firm.

Let the two firms (Firm A and Firm B) choose their locations x_A and x_B along a unit interval $[0, 1]$. The payoffs represent the proportion of customers each firm captures:

	$x_B = 0$	$x_B = 0.5$	$x_B = 1$
$x_A = 0$	$(0.5, 0.5)$	$(0.75, 0.25)$	$(1, 0)$
$x_A = 0.5$	$(0.25, 0.75)$	$(0.5, 0.5)$	$(0.75, 0.25)$
$x_A = 1$	$(0, 1)$	$(0.25, 0.75)$	$(0.5, 0.5)$

Where payoffs (A's market share, B's market share) are given:

  - Rows represent Firm A's choice x_A.
  - Columns represent Firm B's choice x_B.

In equilibrium, both firms tend to locate at the center $x_A = x_B = 0.5$, known as "principle of minimum differentiation", where neither firm can improve its market share by moving.

If the game is sequential, one firm (Leader) moves first, and the other (Follower) responds.

Sequential Hotelling's Game: Firm A Moves First

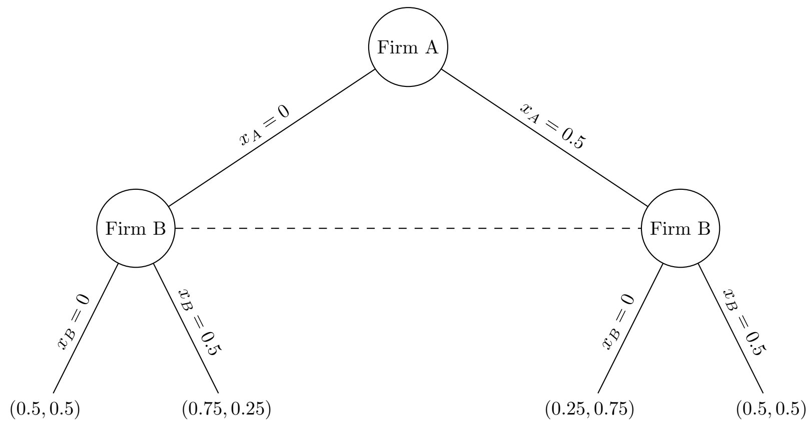

- *Simultaneous Play (Normal Form)*: The Nash equilibrium occurs when both firms locate at $x = 0.5$, as neither can gain an advantage by moving.
- *Sequential Play (Extensive Form)*: The first-mover (Leader) might have an advantage by committing to a position that makes it unattractive for the second firm to differentiate too much.

2 Simultaneous Games

MWG: 7, 8, JR: 7.2, 7.3

2.1 Pure strategy

A pure strategy is a deterministic strategy fro player i, defined on each information set. For ease, it can be helpful to think of a pure strategy as a special case of a mixed strategy, in which the probability distribution of the elements of S_i is degenerate.
A pure strategy set is represented as:

$$S_i = \{s_{1,i}, s_{2,i}, ..., s_{M,i}\} \text{ for } M \text{ strategies}$$

where the M pure strategies are associated with points of a simplex, $\triangle(S_i)$, such that:

$$\triangle(S_i) = \{(\sigma_{1,i}, ...\sigma_{M,i}) \in \mathbb{R}^M : \sigma_{M,i} \geq 0 \text{ for all } m = 1, ...M \text{ and } \sum_{m=1}^{M} \sigma_{mi} = 1\}$$

2.2 Mixed strategy

A mixed strategy for player i is a probability distribution $\sigma_i \in \Delta(S_i)$, where $\Delta(S_i)$ represents the set of all probability distributions over the pure strategies S_i. For a finite pure strategy set S_i, a mixed strategy $\sigma_i : S_i \to [0, 1]$ assigns a probability $\sigma_i(s_i)$ to each pure strategy $s_i \in S_i$, such that:

$$\sum_{s_i \in S_i} \sigma_i(s_i) = 1 \quad \text{and} \quad \sigma_i(s_i) \geq 0 \quad \forall s_i \in S_i.$$

A mixed strategy represents a probabilistic choice, capturing the expected behavior of a player. Unlike pure strategies, which yield deterministic payoffs, the payoff from a mixed strategy is not deterministic; it reflects the expected value over the possible outcomes. A mixed strategy may be used if you receive some private signal, Θ with a $N(0, 1)$ distribution, and then you will form a plan of action based on the realization of the signal. The game, with a mixed strategy, is still represented as:

$$\Gamma_N = [I, \{\triangle(S_i)\}, \{u_i(\cdot)\}]$$

The expected payoff of player i using a mixed strategy σ_i against a pure strategy profile s_{-i} of the other players is given by:

$$u_i(\sigma_i, s_{-i}) = \mathbb{E}_{s_i \sim \sigma_i}[u_i(s_i, s_{-i})] = \sum_{s_i \in S_i} \sigma_i(s_i) u_i(s_i, s_{-i}).$$

Expected payoff of player i when all players are using mixed strategies $\sigma = (\sigma_i, \sigma_{-i})$ is:

$$u_i(\sigma) = \sum_{s_i \in S_i} \sum_{s_{-i} \in S_{-i}} [\sigma_i(s_i)\sigma_{-i}(s_{-i})]u_i(s_i, s_{-i}).$$

Relationship between a mixed and pure strategy: For player i, a mixed strategy $\sigma_i \in \Delta(S_i)$ strictly dominates (*see 2.4*) a pure strategy $s_i' \in S_i$ if, for all $s_{-i} \in S_{-i}$,

$$u_i(\sigma_i, s_{-i}) > u_i(s_i', s_{-i}).$$

- The set of pure strategies is a **proper subset** of the set of mixed strategies: $S_i \subset \Delta(S_i)$.
- Player i's set of possible mixed strategies are represented as a set of points in a simplex (where a simplex is the smallest convex set that contains a given set of points).
- A mixed strategy uses **randomization**, such that the choices, in each information set, are randomized, resulting in a distribution of probabilities across the terminal nodes of the game and an random, induced outcome of the game.

How do we find a mixed strategy equilibrium? If our objective, in a game is to find a mixed strategy equilibrium, we first have to equalize the expected payoffs for each player, making them **indifferent**:

1. Fix Player 1's action and assign the probabilities, $p, (1-p)$ to *Player 2's actions.*[2]
2. Calculate the expected payoffs of Player 1's actions
3. Set both (or all) actions equal to each other, and solve for the probability, p, and other variables, r, s, q if necessary.
4. Repeat steps 1 - 3 for Player 2.

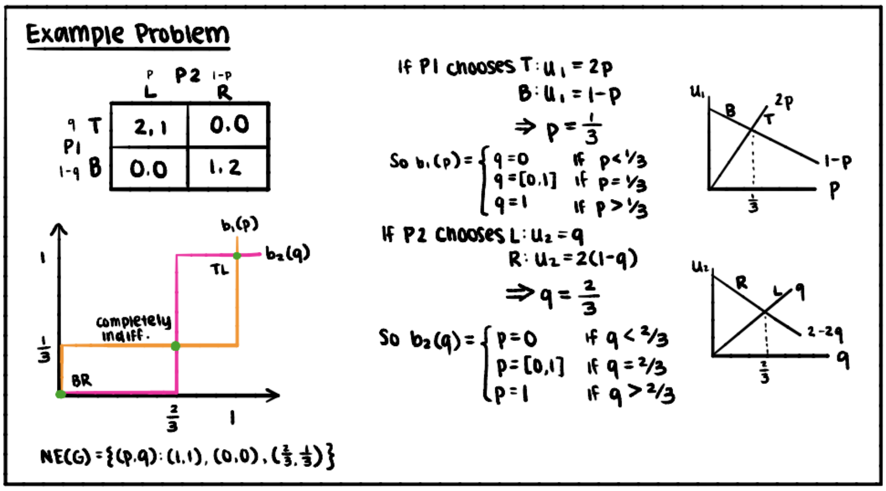

Figure 1: An example of a mixed strategy two-player game

2.3 Rationality

Definition: Player i is said to be rational when playing strategy $s_i \in S_i$ with belief $\mu_i \in \Delta(S_{-i})$, if and only if s_i is a best response to some belief μ_i. That is, a player i is rational if, given their belief μ_i about other players' strategies, they choose a strategy s_i that maximizes their expected payoff. Formally, for every $s_i' \in S_i$ with $s_i' \neq s_i$:

$$\mathbb{E}_{s_{-i} \sim \mu_i}\big[u_i(s_i, s_{-i})\big] \geq \mathbb{E}_{s_{-i} \sim \mu_i}\big[u_i(s_i', s_{-i})\big],$$

What is rationality in a game?

[2]In the case of ¿ 2 potential actions, assign additional variables, q, r, s such that $(1-p) = q + r + s$.

- The **rules of the game** imply rationality when we assume all players are rational, and believe that the other players are also rational, and the players know the others are rational, and so on.
- **A player is rational** if he never plays a strictly dominated strategy.
- In finite games, **the set of rationalizable strategies** is the set of strategies that survive iterated strategic dominance with pure strategies. That is, rationalizable strategies cannot be never-best-responses (NBRs).

Other definitions which define rationalizability at different stategies of the game:

- A *1-rationalizable strategy* for player i is an element of S_i that is a best response to some belief μ_i over other player strategies.
- A *k-rationalizable strategy* for player i is an element of S_i that is a best response to some belief μ_i over $k-1$ rationalizable strategies of other players.
- A *rationalizable strategy* for player i is an element of S_i that is k-rationalizable for every k.

In words, a *1-rationalizable strategy* is one that can be a best response to any belief about what others might play, assuming they are rational. A *k-rationalizable strategy* is optimal when players expect others to choose from their $(k-1)$-rationalizable strategies. Ultimately, a *rationalizable strategy* is one that remains optimal no matter how many rounds of reasoning are applied, i.e., it is consistently a best response to rational expectations of others' rational behavior.

Sequential rationality:

Definition - System of Beliefs:
A system of beliefs μ in an extensive form game Γ_E is a specification of a probability $\mu(x) \in [0,1]$ for each decision node x in Γ_E such that:

$$\sum_{x \in H} \mu(x) = 1$$

for all information sets H.

Definition - Sequential Rationality of Strategies:
Consider an extensive form game Γ_E, an information set H, the player who moves there, i, and i's system of beliefs μ. A strategy profile $\sigma = (\sigma_1, ..., \sigma_I)$ in Γ_E is sequentially rational at H given μ if:

$$E[u_i|H, \mu, \sigma_i, \sigma_{-i}] \geq E[u_i|H, \mu, \tilde{\sigma}_i, \sigma_{-i}]$$

for all $\tilde{\sigma}_i \in \Delta(S_i)$.

2.3.1 Iterated dominance

To find **rationalizable strategies**, we required that the remaining strategies (after iterative removal of the NBRs) were the best responses to some conjectures, not to actual play. We delete *strictly dominated strategies* iteratively until nothing else can be deleted. With each iteration, additional strategies might become dominated (the case of weakly dominated strategies will be dealt with later).

A player must now know not only that her rivals are rational but also that they know that she is, and so on. This is why IDDS relies on the concept of rationality. After iterated deletion of dominated strategies (IDDS), the set of rationalizable strategies can be no larger than the set of strategies surviving IDDS.

Steps to eliminating strictly dominated strategies:

1. Consider only the set of pure strategies
2. Eliminate strictly dominated strategies
3. Consider which mixed strategies are undominated
    - Eliminate mixed strategies which have a pure, dominated strategy within the mixed strategy
    - Look at, and compare, expected payoffs (as some randomization may create new, dominant strategies)

Benefits of IDDS	Limitations of IDDS
The order of elimination is irrelevant.	Won't always reach a solution.
There is no need to know the other player's action.	We must assume common knowledge of rationality and of the game.
It all comes from rationality.	It often leads to inefficient outcomes.

Table 1: Comparison of Benefits and Limitations of Strict Iterative Elimination

Propositions:

1. The set of surviving strategy profiles is invariant to the order of elimination of strictly dominated strategies.
2. If, at the end of this process, only a single profile remains, then it is called the iterated deletion of strictly dominated strategies (IDDS) solution of that game, the game is said to be dominance solvable.

2.3.2 Best response

If a strategy is strictly dominant (see 2.4), then it is also a best response, where the player's payoff is maximized, regardless of the other player's response.

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

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

for all $\sigma_i' \in \Delta(S_i)$.
Strategy σ_i is a best response to σ_{-i} if it is an optimal choice when player i *conjectures* that his opponents will play σ_{-i}.

Never best response: Strategy $sigma_i$ is a never best response if there is no σ_{-i} for which $sigma_i$ is a best response. This implies that $sigma_i$ cannot be justified in rational play.

- Any strategy that is *strictly dominated* is never a best response.
- Never best responses contain the set of strictly dominated strategies. But, recall that strictly dominated strategies is a *smaller* set than the set of NBRs.
- The order of removal of strategies that are NBR does not affect the set of surviving strategies.

2.4 Dominance

Dominance provides a prediction for what moves will be played, based on the most obvious way to compare a player's strategy (ignore randomization over pure strategies at this point). So, dominance, in this case, shows a player's payoff is maximized (in the case of strict dominance) if player's payoff is uniquely highest, regardless of the opponent's move.

Dominance is a key subject for eliminating alternatives when evaluating a game. For example, we use **iterated removal** *(discussed in depth in 2.3.1)* to eliminate *dominated strategies*, resulting in a unique set of strategies.

2.4.1 Strictly, weakly dominant

Any strategy where, regardless of what the opposing player does, there is a best move to take. This means that for all of the opposing player(s)' moves, there exists a strategy that is strictly better, regardless of these strategies. Strictly dominant strategies rarely exist, but we know that when they do, they uniquely maximize player i's payoff, regardless of the rivals' moves.

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

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

for all $s_i' \in S_{-i}$

In a case with **mixed strategies**, we check whether any of player i's mixed strategies does better than his pure strategy, s_i, against every possible pure strategy profile by player j. Player i's pure strategy $s_i \in S_i$ is strictly dominant in a game if and only if, for all $\sigma_i' \in \triangle(S_i)$ such that

$$u_i(\sigma_i', s_{-i}) > u_i(s_i, s_{-i})$$

for all $s_{-i} \in S_{-i}$

2.4.2 Strictly, weakly dominated

Definition: A strategy $s_i \in S_i$ is **strictly dominated** for player i in game $\Gamma_N = [I, \{S_i\}, \{u_i(\cdot)\}]$ if there exists another strategy $s_i' \in S_i$ such that for all $s_{-i} \in S_{-i}$:

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

where strategy s_i' dominates strategy s_i

2.5 Nash Equilibrium

Up to now, we have only imposed the assumption of *common rationality*. But rationality leaves us with too many possible strategies. To refine the set of possible outcomes and identify a solution to the game, additional assumptions are required. The concept of Nash Equilibrium introduces two critical assumptions: **mutual best responses** and **mutual correct expectations**. Together, these conditions create a stable outcome where no player has an incentive to deviate unilaterally.

- Mutual best responses ensure that each player's chosen strategy is optimal given the strategies actually played by others.
- Mutual correct expectations mean that players' beliefs about each other's strategies align perfectly.

Since a nash equilibrium is a situation where each player's chosen strategy is the best possible response to the strategies chosen by others, no player has an incentive to unilaterally change their strategy, as doing so would not improve their payoff. This ensures that the outcome is stable.

Formal definition: A strategy profile $(s_1, \ldots, s_n)$ is a pure strategy Nash equilibrium (NE) if for all players i, and for all $s_i' \in S_i$ with $s_i' \neq s_i$,

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

2.6 Existence of Nash Equilibrium

The following theorem is a **sufficient condition** for the existence of a nash equilibrium.

Theorem: Let $S_1, \ldots, S_n$ be compact and convex subsets of $\mathbb{R}^K$, and let u_i be a continuous and quasiconcave function in s_i for each player i. Under these conditions, a pure strategy Nash equilibrium exists.
Proof:
For each player i, define their best response correspondence as:

$$BR_i(s_{-i}) = \arg\max_{s_i \in S_i} u_i(s_i, s_{-i}).$$

The combined best response correspondence is:

$$BR(s) = \big(BR_1(s_{-1}), \ldots, BR_n(s_{-n})\big).$$

Then, a strategy profile $s = (s_1, \ldots, s_n)$ is a Nash equilibrium if $s \in BR(s)$. Use Kakutani's Fixed Point Theorem to show that such a strategy profile exists. To apply Kakutani's theorem (see 6.2.4), verify the following required properties of $BR_i(s_{-i})$:

- **Non-Empty Valued:** Since u_i is continuous in s_i and S_i is compact, the maximum of $u_i(s_i, s_{-i})$ over $s_i \in S_i$ is attained (by the Weierstrass Theorem, see 6.2.2). Therefore, $BR_i(s_{-i})$ is non-empty for all s_{-i}.
- **Upper Hemicontinuity:** Suppose, toward contradiction, that BR_i is not upper hemicontinuous at s^*_{-i}. Then there exist sequences $s^t_{-i} \to s^*_{-i}$ and $s^t_i \in BR_i(s^t_{-i})$ such that $(s^t_i \to s'_i \notin BR_i(s^*_{-i})$.

  Since $s'_i \notin BR_i(s^*_{-i})$, there exists $s^*_i \in S_i$ such that:

  $$u_i(s^*_i, s^*_{-i}) > u_i(s'_i, s^*_{-i}).$$

  By continuity of u_i:

  $$u_i(s^t_i, s^t_{-i}) \to u_i(s'_i, s^*_{-i}), \quad u_i(s^*_i, s^t_{-i}) \to u_i(s^*_i, s^*_{-i}).$$

  Therefore, for sufficiently large t:

  $$u_i(s^*_i, s^t_{-i}) > u_i(s^t_i, s^t_{-i}),$$

  which implies $s^t_i \notin BR_i(s^t_{-i})$, a contradiction.
- **Convex-Valued:** Suppose $s'_i, s''_i \in BR_i(s_{-i})$, then:

  $$u_i(s'_i, s_{-i}) = u_i(s''_i, s_{-i}) = g^*.$$

  For any $\lambda \in [0, 1]$, the quasiconcavity of u_i implies:

  $$u_i(\lambda s'_i + (1 - \lambda)s''_i, s_{-i}) \geq g^*.$$

  Since g^* is the maximum as $s'_i, s''_i \in BR_i(s_{-i})$, equality holds:

  $$u_i(\lambda s'_i + (1 - \lambda)s''_i, s_{-i}) = g^{\cdot}$$

  Therefore, $\lambda s'_i + (1 - \lambda)s''_i \in BR_i(s_{-i})$, proving that $BR_i(s_{-i})$ is convex.

Having verified all the conditions of Kakutani's Fixed Point Theorem, there exists $s^* \in S$ such that:

$$s^* \in BR(s^*).$$

Such s^* is the pure strategy Nash equilibrium.

Computing Randomized Nash Equilibria: We are given some game, including a given set of players N and, for each $i \in N$, a given set of feasible actions C_i for player i and a given payoff function $u_i : C_1 \times \cdots \times C_n \to \mathbb{R}$ for player i.

The **support** of a randomized equilibrium is, for each player, the set of actions that have positive probability of being chosen in this equilibrium. To find a Nash equilibrium, we can apply the following 5-step method:

1. **Guess a support for all players.** That is, for each player i, let S_i be a subset of i's actions C_i, and let us guess that S_i is the set of actions that player i will use with positive probability.
2. **Consider the smaller game where the action set for each player is reduced to S_i.** Try to find an equilibrium where all of these actions get positive probability.

   To do this, we need to solve a system of equations for some unknown quantities.

   - **Unknowns:** For each player $i \in N$ and each action s_i in i's support S_i, let $\sigma_i(s_i)$ denote i's probability of choosing s_i, and let w_i denote player i's expected payoff in the equilibrium. ($\sigma_i(a_i) = 0$ if $a_i \notin S_i$).
   - **Equations:** For each player i, the sum of these probabilities $\sigma_i(s_i)$ must equal 1.

- For each player i and each action $s_i \in S_i$, player i's expected payoff when choosing s_i while all other players randomize independently according to their σ_j probabilities must equal w_i.

  Let $u_i(\sigma_i, [a_i]) = \mathbb{E}[u_i(a_i|\sigma_{-i})]$ denote player i's expected payoff when they choose action a_i and all other players are expected to randomize independently according to their σ_j probabilities.

  Then the equations can be written as:

$$\sum_{s_i \in S_i} \sigma_i(s_i) = 1, \quad \forall i \in N$$

$$u_i(\sigma_i, [s_i]) = w_i, \quad \forall i \in N, \forall s_i \in S_i.$$

3. **Check if the equations in step 2 have a solution.** If no solution exists, then we guessed the wrong support and must return to step 1 and guess a new support.

4. **Verify that all probabilities are nonnegative.** If any probability is negative, then we must return to step 1 and guess a new support.

   If we have a solution that satisfies all these nonnegativity conditions, then it is a randomized equilibrium of the reduced game where each player can only choose actions in S_i.

5. **Check if the solution is an equilibrium of the original game.** A solution from step (2) that satisfies the condition in (4) is still not necessarily an equilibrium of the original game.

   For each player i and for each action $a_i \in C_i \setminus S_i$, we must check whether choosing a_i would yield a higher payoff than w_i. That is:

$$u_i(\sigma_i, [s_i]) = w_i, \quad \forall s_i \in S_i.$$

   If, for any action $a_i \notin S_i$, we have $u_i(\sigma_i, [a_i]) > w_i$, then we must return to step 1 and guess a new support.

In a finite game, there are only a finite number of possible supports to consider. Thus, an equilibrium $\sigma = (\sigma_i(a_i))_{a_i \in C_i, i \in N}$ with payoffs $w = (w_i)_{i \in N}$ must satisfy:

$$\sum_{a_i \in C_i} \sigma_i(a_i) = 1, \quad \forall i \in N;$$

$\sigma_i(a_i) \geq 0, \quad u_i(\sigma_{-i}, [a_i]) \leq w_i,$ with at least one equality (complementary slackness) for all $a_i \in C_i, \forall i \in N$.

The support for each player i is the set of actions $s_i \in C_i$ for which $\sigma_i(s_i) > 0$, so that $u_i(\sigma_{-i}, [s_i]) = w_i$.

2.7 Problem types

For problems, there are two types of questions: (1) intuition-based problems and (2) solving games. There are a few common simultaneous games we should understand. Problems below cover mixed strategy approaches to solving games, using rationality to find strategies, dominance, and finally nash equilibrium.

1. **From MWG - problem unknown:** *Argue that if a player has two dominant strategies, then for every strategy choice by his opponents, the two strategies yield her equal payoffs.*

   Answer:

   In a normal form game, suppose a player has two dominant strategies, s_i^* and $\tilde{s}_i$. A dominant strategy is one that yields at least as high a payoff as any other strategy, regardless of what the opponents choose. Therefore, for all possible strategies s_{-i} of the other players, we have $u_i(s_i^*, s_{-i}) \geq u_i(s_i', s_{-i})$ for all $s_i' \in S_i$, and similarly, $u_i(\tilde{s}_i, s_{-i}) \geq u_i(s_i', s_{-i})$ for all $s_i' \in S_i$.

   Since both s_i^* and $\tilde{s}_i$ are strategies in S_i, we can directly compare them by substituting $s_i' = \tilde{s}_i$ in the first inequality and $s_i' = s_i^*$ in the second inequality. Substituting $s_i' = \tilde{s}_i$ in the first inequality gives $u_i(s_i^*, s_{-i}) \geq u_i(\tilde{s}_i, s_{-i})$ for all s_{-i}. Substituting $s_i' = s_i^*$ in the second inequality gives $u_i(\tilde{s}_i, s_{-i}) \geq u_i(s_i^*, s_{-i})$ for all s_{-i}.

Combining these two inequalities, we obtain $u_i(s_i^*, s_{-i}) = u_i(\tilde{s}_i, s_{-i})$ for all s_{-i}. This shows that for every possible combination of the opponents' strategies s_{-i}, the payoffs from s_i^* and $\tilde{s}_i$ are identical. Therefore, if a player has two dominant strategies, those strategies yield equal payoffs against any strategies chosen by the opponents.

2. **Midterm 2025 - Question 1:** *Consider a three-player normal-form game with Player 1's strategy set: $\{T, B\}$. Player 2's strategy set: $\{L, R\}$. Player 3's strategy set: $\{X, Y, Z\}$.Two tables of payoffs, dependent on whether Player 1 chooses T or B. Each cell shows (u_1, u_2, u_3) = payoffs for Players 1, 2, and 3. Third player is choosing the rows X, Y, Z, but his payoffs are the last entry of each payoff vector.*

P1 chooses T:

	L	R
X	$(2, 2, 0)$	$(1, 1, 2)$
Y	$(3, 1, 2)$	$(1, 2, 0)$
Z	$(1.5, 0, 1)$	$(7, 0, 1)$

P1 chooses B:

	L	R
X	$(1.5, 2, 0)$	$(0, 1, 2)$
Y	$(2, 1, 2)$	$(0, 2, 0)$
Z	$(1, 0, 1.5)$	$(4, 0, 1.5)$

a. *Does any player have a pure dominant strategy? If true, which?*
b. *Does any player have a mixed strategy? If so, who? What are the strictly dominated strategies?*
c. *Find the set of strategies that survive the iterative deletion of never-best-response.*
d. *Find (or characterize) the Nash equilibrium (or equilibria).*

Answer:

a. For Player 1: The table where $s_t = \{T\}$, all values of Player 1's payoff are higher than where $s_t = \{B\}$, written as $u_1(T) > u_1(B)$. From the table, where highlighted the values all exceed in table T.

	L	R
X	$(2, 2, 0)$	$(1, 1, 2)$
Y	$(3, 1, 2)$	$(1, 2, 0)$
Z	$(1.5, 0, 1)$	$(7, 0, 1)$

	L	R
X	$(1.5, 2, 0)$	$(0, 1, 2)$
Y	$(2, 1, 2)$	$(0, 2, 0)$
Z	$(1, 0, 1.5)$	$(4, 0, 1.5)$

For Player 2: There are no pure dominant strategies.
For Player 3: There are no pure dominant strategies.

b. We eliminate $s_t = \{B\}$ because T strictly dominates B in pure strategies. When any strategy is strictly dominated in pure strategies, it is also strictly dominated in mixed strategies. So in any mixed strategy that assigns $\pi > 0$ to B and $1 - \pi \geq 0$ to T will be dominated by a mixed strategy that assigns $\pi = 0$ to B and $\pi = 1$ to T.

Player 2 and Player 3 have no dominant mixed strategies. Any non-degenerate mixed strategy demonstrates it does not result in a strictly better payoff.
For Player 2:

$$p_x(2) + p_y(1) + p_z(0) = p_x(1) + p_y(2) + p_z(0)$$
$$-p_x(1) - p_y(1) = -p_x(1) - p_y(1)$$
$$p_x = p_y \implies$$
$$(p_x, p_y) = (\frac{1}{2}, \frac{1}{2})$$

So we get:

$$E_\sigma(u_2(\sigma, L)) = \frac{1}{2}(2) + \frac{1}{2}(1) + \frac{0}{2}(0) = 1.5$$
$$E_\sigma(u_2(\sigma, R)) = \frac{1}{2}(1) + \frac{1}{2}(2) + \frac{0}{2}(0) = 1.5$$

Here the expected value is the same for Player 2 for both $s_2 = \{L, R\}$. Therefore, there is no strictly dominant mixed strategy. The expected outcome is dependent on the relative probabilities for the choice of x or y by Player 3.

For Player 3, we can use a similar approach, where we have some probability, π associated with Player 2's choice, $s_2 = \{L, R\}$. We see:

$$p_L(0) = p_R(2) \to p_L = 1, p_R = 0$$

$$p_L(2) = p_R(0) \to p_L = 0, p_R = 1$$

$$p_L(0) = p_R(0) \to (p_L, p_R) \in [0, 1]$$

Similar to Player 2's problem, Player 3 has the same expected value = 2, from either L, R. There is no probability that can be used here to find some mixed strategy solution.

c. From above in part (a) and (b), we can see the first steps of iterated deletion. Next, we can particularly see how elimination of Z by Player 3, given it is never a best response once we are in $s_1 = \{T\}$ world will occur. So the iterated deletion goes in this order:

(1) Remove $s_1 = \{B\}$ because it is as player 1 never chooses it as a best response

(2) Now in the environment for $s_1 = \{T\}$ we can see the following:

$$BR_3(L) = Y \quad BR_3(R) = X \quad \& \quad BR_2(Y) = L \quad BR_2(X) = R$$

Its clear that $s_2 = \{Z\}$ is never a best response strategy for Player 3 $\implies$ eliminate Z.

(3) No more strategies can be eliminated, as for Player 2, both $\{L, R\}$ are sometimes best responses.

The surviving strategies are:

$$\begin{aligned} \text{Player 1} &\to \{T\} \\ \text{Player 2} &\to \{L, R\} \\ \text{Player 3} &\to \{X, Y\} \end{aligned}$$

d. To find (or characterize) the Nash equilibrium (or equilibria) we look at the reduced subgame between Player 2 and Player 3 with the surviving strategies, listed above. Given there is not a purely dominant strategy, we will solve for a **mixed strategy** in the subgame:

	L	R
X	$(2,0)$	$(1,2)$
Y	$(1,2)$	$(2,0)$

Where $p = \pi_2(L)$, the associated probability with Player 2. From this, we see the probability that Player 2 chooses R is $(1-p)$. Next, $q = \pi_3(X)$, so the probability that Player 3 chooses Y is $(1-q)$. With these probabilities, we identify the indifference point, where we take the difference between the payoffs for X and payoffs for Y.

Player 3's indifference

$$U_3(\cdot) = (2+0)(1-p) - (2+0)(p) \to \quad 2(1-p) = 2p \to \quad p = \frac{1}{2}$$

Player 2's indifference

$$U_2(\cdot) = (2+1)(q) - (1+2)(1-q) \to \quad 3q = 3(1-q) \to \quad q = \frac{1}{2}$$

With these probabilities, we now have the mixed strategy nash equilibrium, such that:

$$\text{NE} \equiv \left(s_1 = T; \quad s_2(R) = s_2(L) = \frac{1}{2}; \quad s_3(X) = s_3(Y) = \frac{1}{2} \right)$$

3 Dynamic games

MWG: 9

Dynamic games can be solved in their normal form, but importantly, in dynamic games the **credibility** of the other player's strategy becomes important. The subgame perfect nash equilibria is the strategy, underlined by sequential rationality, is the central way that we can create an equilibrium strategy that specifies optimal behavior from any point in the game onward. We therefore *refine* the concept of nash equilibrium with beliefs, to ensure a satisfactory equilibrium prediction.

3.1 Sequential rationality

Definition: The *principle of sequential rationality* is that a player should specify optimal actions at every point in the game tree.

In dynamic games, sequential rationality is a structure in which players must deploy optimal behaviour at any moment of the game onwards. For dynamic games this is *required*, where, in a dynamic setting, some Nash equilibria *(see 2.5)* are not sensible.[3] In order to rule out the nash equilibria that may contain *empty threats*, we use sequential rationality to ensure that the player's strategy does not respond to these incredulous threats.

3.2 Backward induction

In finite games of perfect information, backward induction implements the requirement of sequential rationality by assigning optimal actions at every decision point, starting from the end of the game and proceeding backward. A joint pure strategy s is a *backward induction strategy* in a finite extensive-form game with perfect information if it is constructed by solving the game from the end of the game tree backward, as follows:

1. Identify all *penultimate nodes*—these are decision nodes where all immediate successor nodes are terminal (end) nodes.
2. At each such node x, let the player who moves at x choose the action that maximizes their payoff among the available actions. Let u^x denote the resulting payoff vector.
3. Replace node x and its successors with the payoff vector u^x, turning x into a new terminal node.
4. Repeat this process on the reduced game tree: identify new penultimate nodes and assign optimal actions, working backward through the tree.
5. Continue until all decision nodes have been assigned actions.

Kuhn's Theorem: If s is a backward induction strategy for a perfect information finite EFG, then s is a Nash equilibrium of the game.

Proof: We show that for every player $i \in N$ and every alternative strategy $s'_i \in S_i$, it holds that

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

Suppose, for contradiction, that there exists a player i and a strategy $s'_i \in S_i$ such that

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

Let $s' = (s'_i, s_{-i})$. Let e and e' be the terminal nodes (outcomes) induced by strategies s and s', respectively. Then

$$u_i(e') > u_i(e).$$

Define the set X of decision nodes x such that, starting from x, player i could obtain a strictly higher payoff by deviating from s_i:

$$X := \{x \in H_i \mid u_i(s'_i, s_{-i} \mid x) > u_i(s_i, s_{-i} \mid x)\}.$$

[3] For example, nash equilibria involve actions at decision nodes that are not reached. Often we consider these nodes as threats. Some threats may be credible, or not credible (empty).

Since $u_i(s'_i, s_{-i}) > u_i(s_i, s_{-i})$, this set is non-empty.
Let $\bar{x} \in X$ be a node with no strict successors in X; that is, $\bar{x}$ is a minimal such node in the game tree. By construction, we have:

$$u_i(s'_i, s_{-i} \mid \bar{x}) > u_i(s_i, s_{-i} \mid \bar{x}).$$

Moreover, since $\bar{x}$ has no strict descendants in X, it follows that:

(i) $\bar{x}$ is a decision node belonging to player i;

(ii) At all nodes strictly following $\bar{x}$, player i's strategy s_i is already optimal given s_{-i}.

Thus, we can define a deviation strategy s''_i for player i that modifies s_i only at $\bar{x}$, taking a different action at $\bar{x}$ but following s_i thereafter. Then

$$u_i(s''_i, s_{-i} \mid \bar{x}) > u_i(s_i, s_{-i} \mid \bar{x}).$$

But s_i was constructed via backward induction. So when the backward induction algorithm reached node $\bar{x}$, it selected the action at $\bar{x}$ that maximized i's payoff given the continuation strategies. Therefore, $u_i(s_i, s_{-i} \mid \bar{x})$ must be the maximum payoff player i could achieve at $\bar{x}$, contradicting the fact that s''_i yields a strictly higher payoff.

This contradiction shows that no profitable deviation s'_i exists for any player i. Hence, the strategy profile s is a Nash equilibrium.
Backward induction algorithm The steps to perform backward induction are the following:

1. Take the last decision nodes in the game tree (one less than the terminal node)
2. Find the payoff vector for the payer who is *playing*, or making a decision at the last decision node. Identify the highest payoff.
3. Take one maximum move, and "save" the other, or second highest payoff move, if there are ¿ 1 maximum moves at that node.
4. Eliminate all moves and the terminal nodes following the node from which you extracted the move. Then, we will assign the payoff of the *max* to the previous node
5. Before ending the process, evaluate against these two cases:
   - *Case 1:* If there are more terminal nodes, repeat steps 1 - 4
   - *Case 2:* End the backward induction process
6. Repeat if there are saved nodes

3.3 Sub-game perfect Nash Equilibrium

The concept of SPNE strengthens NE by requiring that players' strategies form a Nash equilibrium not just in the entire game, but also in every proper subgame *(see 1.2)* This ensures that strategies are sequentially rational—optimal at every possible decision node, whether or not that node is reached during equilibrium play. From this follows the lemma: $SPNE \subset NE$.

Definition: A subgame nash perfect equilibrium (SPNE), is a strategy, $s^* \in X_{i \in N, S_i}$ of an EFG, Γ_E, if $O_n(s^*_i I_n, s^*_{-i} I_n) \succeq_i I_n O_n(s_i I_n, s^*_{-i} I_n)$ for each $i \in N$ and for all histories, $h \in H/Z$. This strategy profile must induce a nash equilibrium in a particular subgame, γ of Γ_E, and it must be considered a nash equilibrium when γ is played in isolation.

Propositions:

Prop 1 If the only subgame is the game as a whole *(there are no proper subgames)* then every nash equilibrium is subgame perfect

Prop 2 Once you find an SPNE, an SPNE is then induced in every subgame of Γ_E.

Steps to finding an SPNE:

1. Locate a single decision node t that is its own information set ($h(t) = \{t\}$), meaning the player moving at t knows exactly where they are.
2. Include all successor nodes $S(t)$, which are reachable from t through any sequence of actions.
3. Ensure that for each successor node t', its entire information set $h(t')$ is fully contained within the subgame. There cannot be any "broken" information sets that partially belong to the subgame.

Zermelo's Theorem: Every finite extensive-form game with perfect information has at least one pure strategy Nash equilibrium that can be found by backward induction. Furthermore, if no player has two terminal nodes with the same payoff, then the backward induction outcome is unique.

Proof. *We proceed by backward induction on the game tree.*
Step 1: Existence. Because the game is finite and of perfect information, the game tree is finite and every decision node belongs to a known player with a finite number of available actions. Starting from the terminal nodes and working backwards:

- At any terminal node, payoffs are given.
- At each decision node, the player chooses the action that leads to the successor node with the highest utility for them.
- This process eventually reaches the root node, since the tree is finite.

By recording each player's optimal choice at every decision node, we construct a pure strategy profile. This profile forms a subgame perfect Nash equilibrium and therefore is a Nash equilibrium.

Step 2: Uniqueness under distinct payoffs. Assume each player has a strict preference over terminal nodes (i.e., no player receives the same payoff at two different terminal nodes). Then:

- At each decision node, the player has a strict preference for exactly one available action.
- There is a unique optimal action at each node during backward induction.
- Therefore, the backward induction path is unique.

Thus, the pure strategy profile determined by backward induction is the unique Nash equilibrium.

□

3.4 Problem types

1. **MWG Chapter 9 - Figure 9.B.4** *Firms I and E are playing a simultaneous-move game after entry. Each can choose whether to fight or accommodate. Identify the sequentially rational - subgame perfect -Nash equilibria and compare to a standard nash equilibria (assuming there was not simultaneity in the second move).*

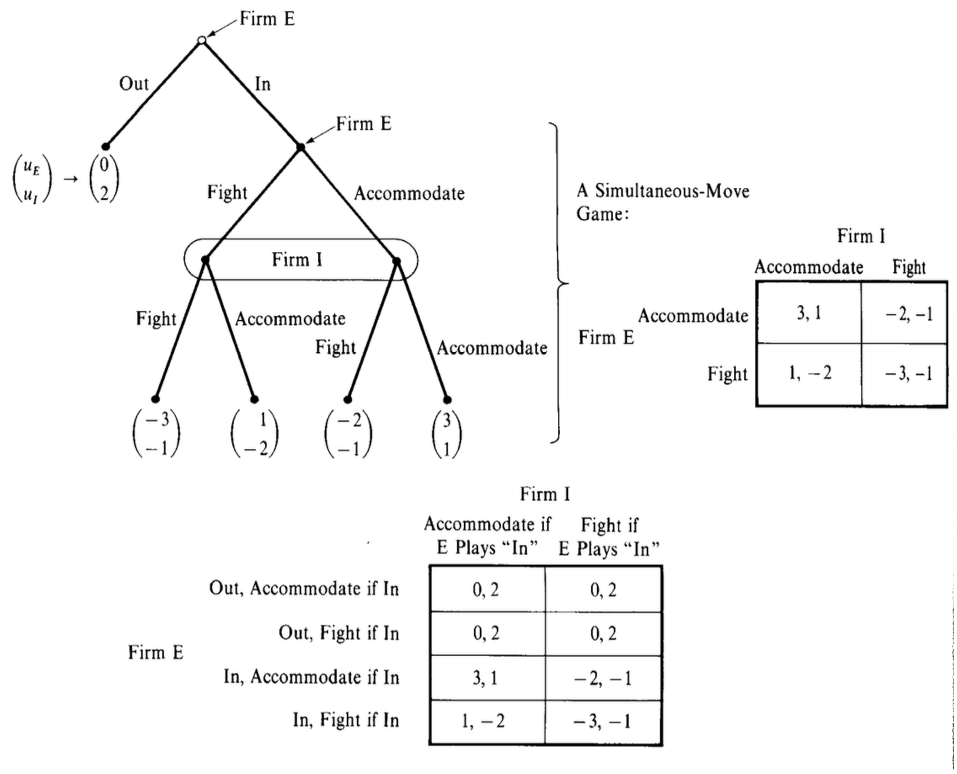

Figure 2: Extensive and normal form depiction of sequential rationality

Answer: First we can start by examining the standard Nash equilibria, (σ_E, σ_I). We first assess the out games:

$$\text{If Firm E} \rightarrow \text{(out, in)} \quad \text{Firm I has } BR_I = \{\text{Acc, Fight}\}$$

For standard nash equilibrium we need to go through the entire set of hypotheticals, such that Firm E has the choice of going *out* first, and his only next available move on the chain is to assume he went *in*, and now subsequently chooses between *accommodate* and *fight*. He chooses *accommodation*, as there is no best response by Firm I that will provide him with a higher payoff. Given this is simply a hypothetical, Firm I plays *fight if in*. The first nash equilibrium is thus:

$$NE_1: \quad (\sigma_E, \sigma_I) = \text{(out, acc if in), (fight if firm E plays in)}$$

Note that Firm I has only one move, and Firm E has two moves. While there is an information set, we see that Firm I is deciding an off-path choice, so it is still valid. Additionally, choosing fight is a credible threat to securing a zero payoff.
The next two nash are the following:

$$NE_2: \quad (\sigma_E, \sigma_I) = \text{(out, fight if in), (fight if firm E plays in)}$$

$$NE_3: \quad (\sigma_E, \sigma_I) = \text{(in, accommodate if in), (accommodate if in)}$$

The second nash equilibrium follows from the first, in that the NE must include all contingency plans for the players that maintain stability. The third nash equilibrium is the final one that results in a payoff that neither firm will **sequential rationality**, and holds under a simultaneous move environment.

2. **Example from outside sources:** *Recall the firm entry game, a simplified version of the game above, from MWG.*

   Recall that the NE of the game, as derived from the normal form, are:

   - (In, Accommodate)
   - (Out, Fight)

   I will now prove that (In, Accommodate) is the unique SPNE of the game. By:

Figure 3: Entry Game

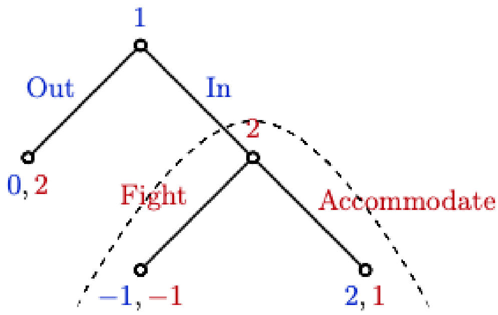

- Identifying the subgames in the game: (1) the entire game, and (2) the subgame beginning at node 2.
  In the subgame starting at node 2, note that Player 2's best response is to "Accommodate," as a payoff of 1 is better than -1. Thus, this subgame simplifies to the payoff $(2, 1)$. Replace node 2 with this payoff in the game tree, reducing it to a single decision for Player 1.
- Next, solve the reduced game for Player 1. Clearly, Player 1 prefers "In," which gives a payoff of 2, over "Out," which yields a payoff of 0.

Therefore, the unique SPNE is (In, Accommodate).

3. **Final Exam 2025:** *Consider the following two-player EFG. Player 1 can choose either "Stop" or "Continue". If she chooses "Stop", the game ends with the payoffs (1, 1) for Player 1, 2 respectively. If she chooses "Continue" then the players simultaneously announce non negative integers* (x, y) *and each player's payoff is the product of the numbers,* $(u_i = xy)$ *for* $i \in \{1, 2\}$.
   (a) Formulate this as an EGF with simultaneous moves to find the subgame perfect equilibria.
   (b) Does the answer change if you bound x,y (i.e. $x, y < M < \infty$*)?*

   Answer (a): S indicates "Stop", C indicates "Continue". Recall xy is the simultaneous announcement of both x and y from each player, where they can choose any number $\in [0, \infty)$.

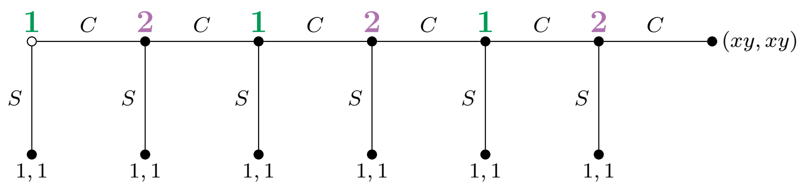

   Now with the game, we have labelled the information sets, such that Player 1 moves first. He has two information sets: the initial node and the information set that occurs when Player 2 chooses "Continue". To find the SPNE, we start by using the **backward induction algorithm** to find the nash equilibrium of the only proper subgame.

   Case 1: Proper subgame from terminal node: The only proper subgame begins at the second to last terminal node. Assume this game goes on for an infinite amount of rounds, and while the figure above depicts Player 2 at the information set before the terminal node, it could be either Player 1 or 2. Because its a simultaneous move game, it is irrelevant.

   Lets assume the following choices are played by Player 2:

   Player$_2$ plays $\rightarrow n \in [1, \infty)$, $\quad BR_1(n) = q = max(n, \infty] \equiv \lim_{n\to\infty} n{*}q \rightarrow \infty$, $\quad$ Payoff $= (n{*}q, n{*}q) > \infty$

   Here, we can see, it is in Player 1's interest to play the highest number possible. Meaning that there is actually no best response to a non-zero play by Player 2. There is no nash equilibrium in this case. The next choice we see is Player 2 plays zero:

   $$\text{Player}_2 \text{ plays } \rightarrow 0 \quad BR_1(0) = n, n \in [0, \infty), \quad \text{Payoff} = (n * 0, n * 0) = (0, 0)$$

Here, any number, played by Player 1 results in a payoff of zero for both. So, this is the only pure nash equilibrium in this subgame, such that:

$$(\sigma_1, \sigma_2) = ((x = n), (Continue, y = 0)), \text{ with payoff } (0,0)$$

Case 2: Subgame = entire game: Now we start at the initial node, the other subgame (which is not proper as it is equivalent to the entire game). When evaluating whether Player 1 wants to C or S, he will choose S, given that it is predicted the other outcome will be less than the payoff at the end of the S node (because he ends up in the alternative nash equilibrium, as described above, with payoff (0,0)). The result is an SPNE that takes the form,

$$(\sigma_1, \sigma_2) = ((Stop, x = 0), (y = 0)), \text{ with payoff } (1,1)$$

Answer (b): In short, **yes**. The introduction of a bound on the values of x, y creates additional opportunities for SPNE, given that we know the value of M. Note, whether we get a new SPNE depends on what value M takes. We have to set up conditions for the value of M. Returning to the subgames from part (a):

Case 1: terminal subgame: This proper subgame now has two nash equilibria, in the case that $M > 1$. When $M < 1$, the value of $\frac{1}{n}$ will converge to zero, and thus is not increasing in the exponential. So, when $M > 1$, there is an actual highest number available to both players. Both players will choose this number, M to maximize their payoffs. The result is the SPNEs:

$$(\sigma_1, \sigma_2) = ((Continue, x = M), (y = M)), \text{ with payoff } (M^2, M^2)$$

The second SPNE in this case is the original one found above in part (a). Now if $M = 1$ (we will ignore the case of $0 < M < 1$), there are three SPNEs, including the first one found in part (a):

$$\begin{aligned} SPNE_1 =& (\sigma_1, \sigma_2) = ((Continue, x = 1), (y = 1)) && \text{with payoff} \equiv (1,1) \\ SPNE_2 =& (\sigma_1, \sigma_2) = ((Stop, x = 1), (y = 1) && \text{with payoff} \equiv (1,1) \end{aligned}$$

4 Incomplete information

MWG: 8.E 12, 14.A-C

4.1 Bayesian games

Until now, we assumed players had *complete information* about the game and each other's payoffs. In reality, players often face *incomplete information*, where they lack full knowledge of others' preferences or strategies. This introduces the need to account for beliefs. In a Bayesian game, a player has beliefs about the other player(s) preferences, and the other players have their own beliefs about his preferences, and so on. This concept is similar to rationalizability.

To model this, each player's preferences are treated as the realization of a **random variable**. While a player observes their own realization, the ex-ante probability distribution of these random variables is common knowledge. This framework transforms a situation of incomplete information into a game of *imperfect information*, where Nature makes the first move by randomly selecting each player's type, and players then observe only their own type.

Definition: A Bayesian game is

$$(I, \{A_i\}_{i\in I}, \{u_i\}_{i\in I}, \{\Theta_i\}_{i\in I}, F)$$

where:

- $I = \{1, 2, \ldots, n\}$ is the finite set of *players*.
- A_i is the set of actions available to *player* i. The set of action profiles is $A = A_1 \times A_2 \times \cdots \times A_n$.
- Θ_i represents the set of possible *types* for player i. A type $\theta_i \in \Theta_i$ is the private information known only to player i. The type space is $\Theta = \Theta_1 \times \Theta_2 \times \cdots \times \Theta_n$.

- F is the common prior *probability distribution* over type profiles Θ. Each player i knows their own type θ_i and holds beliefs about others' types based on the conditional probability distribution derived from F:

$$F(\theta_{-i} \mid \theta_i) = \frac{F(\theta_i, \theta_{-i})}{F(\theta_i)}$$

  The joint probability distribution F(.) is assumed to be common knowledge among the players.

- $u_i : A \times \Theta_i \to \mathbb{R}$ is the payoff function for player i, where $u_i(a, \theta_i)$ depends on the action profile and the type profile.

Feature	**Imperfect Information**	**Incomplete Information**
Definition	A player does not know the exact actions previously taken by others at the time of their move.	A player lacks knowledge about some aspect of the game structure, such as payoffs or types of other players.
Cause of Uncertainty	Unobserved moves within the game tree.	Uncertainty over the "type" of a player or unknown parameters (modeled with beliefs).
Standard Model	Extensive-form games with information sets that include multiple decision nodes.	Bayesian games (or Harsanyi transformation).
Does every player know the structure of the game?	Yes. Players know the structure, payoffs, and possible strategies.	No. At least one player is uncertain about the game structure, usually the payoffs or types.
Typical Solution Concept	Subgame Perfect Equilibrium or Sequential Equilibrium.	Bayesian Nash Equilibrium or Perfect Bayesian Equilibrium.
Example	Player 2 doesn't observe Player 1's move before choosing.	Player 1 doesn't know whether Player 2 is a high-cost or low-cost type.

Table 2: Comparison of Imperfect and Incomplete Information in Game Theory

4.2 Bayesian nash equilibrium concepts

Definition: Bayesian Nash Equilibrium is a strategy profile $\{s_i(\theta_i)\}$ such that each player maximizes their expected utility given their beliefs about the types of other players:

$$E[u_i(s_i(\theta_i), s_{-i}(\theta_{-i})|\theta_i] \geq E[u_i(s'_i, s_{-i}(\theta_{-i})|\theta_i], \quad \forall s'_i \in S_i, \quad \forall \theta_i \in \Theta_i.$$

Building off this definition, we consider the *pure strategy* Bayesian nash equilibrium of (I, A, u, Θ, F), such that is a strategy profile where:

$$s_i(\theta_i) \in argmax_{a_i \in A_i} \mathbb{E}[u_i(a_i, s_{-i}(\theta_{-i}), \theta_i) \mid \theta_i].$$

This definition shifts the perspective to the point after a player observes their own type θ_i but before knowing the types of other players. In this context, each player selects a best response based on their private information and expectations about others' strategies. A BNE is defined as a strategy profile where each player's strategy $s_i(\theta_i)$ maximizes their expected utility, given their own type θ_i and their beliefs about the types of other players θ_{-i}.

Bayesian nash and nash equilibrium with mixed strategies: In complete information games, mixed strategy equilibria involve players intentionally randomizing their actions. In Bayesian games, players instead act deterministically based on private types (signals), which effectively replicate mixed strategies. When types are independent and preferences align, deterministic actions in a BNE correspond to the mixed strategy NE.

Definition: Weak Perfect Bayesian Equilibrium (WPBE)
A profile of strategies and system of beliefs (σ, μ) is a weak perfect Bayesian equilibrium (WPBE) in an extensive form game Γ_E if:

(i) The strategy profile σ is sequentially rational given belief system μ.

(ii) The system of beliefs μ is derived from strategy profile σ through Bayes' rule whenever possible:

$$\mu(x) = \frac{\Pr[x|\sigma]}{\Pr[H|\sigma]}, \quad \forall x \in H$$

Definition: Sequential Equilibrium
A strategy profile and system of beliefs (σ, μ) is a sequential equilibrium of an EFG, Γ_E if:

1. The strategy profile σ is sequentially rational given belief system μ.
2. There exist:
   - A sequence of completely mixed strategies converging to σ, $\{\sigma_k\}_{k=1}^{\infty} \to \sigma$.
   - A sequence of belief systems converging to μ, $\{\mu_k\}_{k=1}^{\infty} \to \mu$.
   - Each belief system μ_k is derived from strategy profile σ_k using Bayes' rule.

What we are looking for in a sequential equilibrium is that there is consistency in the mixed strategies. Specifically, consistency means that as a sequence of strictly mixed strategies σ^t converges to a strategy σ, the corresponding belief systems μ^t, induced by σ^t via Bayes' Rule, converge to the belief system μ. This ensures that the belief system μ is tied to the players' strategy choices σ and reflects their strategic behavior.

Definition (expanded): (σ, μ) is a sequential equilibrium (SE) if it is sequentially rational and consistent. By this definition of SE, the sequential rationality requirement in the SE is identical to that in Weak PBE: players' strategies σ must maximize their expected payoffs at every information set, given the belief system μ . However, the consistency requirement is significantly stronger. Specifically, it ensures that beliefs are consistent both on and off the equilibrium path. This eliminates equilibria that rely on arbitrary or implausible off-path beliefs about what might occur if the game deviates from the equilibrium path.

textbfSteps to finding an SE:

- Use backward induction to eliminate dominated strategies.
- Solve for the belief system to rationalize the moves (starting from the last mover). *This becomes more difficult if there are > 1 moves by each player.*
- Check if μ is consistent using either Baye's Rule or checking for nonzero probabilities.

Proposition: Relationship between WPBE and Nash Equilibrium $\Rightarrow$ A strategy profile σ is a Nash equilibrium of an extensive form game Γ_E if and only if there exists a system of beliefs μ such that:

1. Strategy profile σ is sequentially rational given μ at all info sets H such that $\Pr[H|\sigma] > 0$.
2. The system of beliefs μ is derived from σ through Bayes' rule whenever possible.

Every WPBE is a Nash equilibrium, but not all Nash equilibria are WPBE.

4.3 Asymmetric information

Asymmetric information occurs when different agents possess different information on the game. This results in different relevant strategic opportunities for agents, given strategic opportunities are directly related to the distribution of information. Strategic opportunities arising from asymmetric information typically lead to *market failures.*

A key example of this is documented in Akerlof (1970) where the market for used cars, lemons, is found to contain a significant amount of information asymmetries. Buyers and sellers both have very different information about the product quality, and buyers only know the product quality on average, resulting in a supply of low quality 'lemons' in the market in which buyers cannot differentiate between.

4.3.1 Adverse selection (in the labor market)

In the labor market, here is often a case with many identical firms seeking to hire workers and produce identical output, using constant return to scale (CRS) technology. Technology and labor are the only inputs, and firms seek to profit maximize, under risk neutral probabilities. Firms also take prices as given. In this case, workers differ in their productivity level, θ. Therefore, the set of productivity levels is $0 \leq \underline{\theta} < \overline{\theta} < \infty$, where $[\underline{\theta}, \overline{\theta}] \in \mathbb{R}$. There are some proportion of workers with a productivity that is less than average.

When workers' productivity is unobservable, firms need to develop a notion of competitive equilibrium in this environment. The wage rate is independent of workers' type, θ. This is a **single wage rate case**. Workers only work if the firm offers a wage that is greater than their reservation rate. Note the reservation rate, or $r(\theta)$ is the opportunity cost to a worker, of type θ who accepts employment. The condition is therefore:

$$r(\theta) \leq w \qquad \text{where set of workers who accept the wage:} \quad \Theta(w)\{\theta : r(\theta) \leq w\}$$

So how do we approach this model? In a competitive labor market model with unobservable worker productivity levels, a competitive equilibrium is a wage rate w^* and a set of workers, Θ^* of worker types who will accept the employment, such that:

$$\Theta^* = \{\theta : r(\theta) \leq w^*, \quad w^* = \mathbb{E}[\theta | \theta \in \Theta^*]\}$$

This solution exhibits rational expectations by the firm, such that firms can anticipate the average productivity of workers who accept employment in this equilibrium. If $\Theta^* = \emptyset$ then it is not well defined.

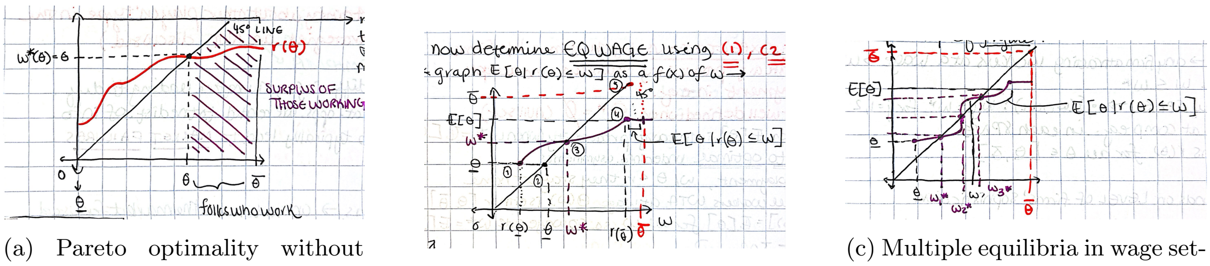

(a) Pareto optimality without asymmetries

(b) Equilibrium wage

(c) Multiple equilibria in wage setting

Figure 4: Descriptive figures for labor market asymmetries

The key takeaways from this model are the following:

- **Competitive equilibrium fails to Pareto optimal** under information asymmetries. This is because some workers will stay home who should be working, and sometimes there are not enough workers to accept the job. Therefore, this type of equilibrium depends on the share of both good and bad workers. The problem is highly dependent on the distribution of quality workers in the economy. When you have a high share of high productivity workers, the average productivity of workforce is much greater than the reservation wage. But if there is a high share of low productivity workers than the firms do not have offer a wage rate that is sufficiently high to induce any employment.
- **Adverse selection** can occur when an informed individual's trading decision depends on her unobservable characteristics in a manner that adversely affects the uninformed agents in the market. This can occur when relatively less productive workers have a willingness to accept any wage. Additionally, when the reservation wage varies with respect to θ this leads to a breakdown in efficiency.

Other issues that may occur include:

- Market unravelling
- Complete market failure
- Multiple equilibria with pareto ranking

Game proof: Let's assume we have the market behaviour that is captured in a two stage game, such that in the first stage, 2 firms simultaneously announce wage offers. In the second stage, workers decide whether to work for a firm, and if so, they select which firm to work for.

- Assume $F(\cdot), r(\cdot)$ is common knowledge
- Assume $r(\cdot)$ is strictly increasing and $r(\theta) \leq \theta$ for $\theta \in [\underline{\theta}, \overline{\theta}]$
- Assume $F(\cdot)$ has an associated density such that $f(\theta) > 0$ for all $\theta \in [\underline{\theta}, \overline{\theta}]$. This implies there exists a subgame nash perfect equilibrium, such that:
  1. $w^* =$ set of competitive equilibrium wages for an adverse selection labor market model
  2. $w^* = max\{w : w \in w^*\}$
- if $w^* > r(\theta)$ and there exists some value, ϵ such that $\epsilon > 0 \implies \mathbb{E}[\theta | r(\theta) \leq w'] > w' \ \forall w' \in (w^* - \epsilon, w^*)$. There is a unique pure strategy SPNE in this case, where workers receive their wage w^* and workers of different types in the set accept employment at firms
- But if $w^* = r(\theta)$, then there are multiple pure strategy SPNEs. In every pure strategy SPNE, each agents payoff will be equal to the payoffs in the highest wage competitive equilibrium.

4.3.2 Signalling (in the labor market)

Signalling is the process where an agent seeks a mechanism to help firms distinguish between the ability of workers, where workers can then perform some action to reveal their type. In signalling problems, we primarily have two types of workers, high skilled and low skilled (θ_h, θ_l). A detailed case of a signalling problem is presented in 5.
Lemma Separating Equilibria in signalling games: In any **separating PBE**, the equilibrium wages satisfy:

$$w^*(e^*(\theta_H)) = \theta_H, \quad w^*(e^*(\theta_L)) = \theta_L.$$

That is, each worker type receives a wage equal to her productivity level.

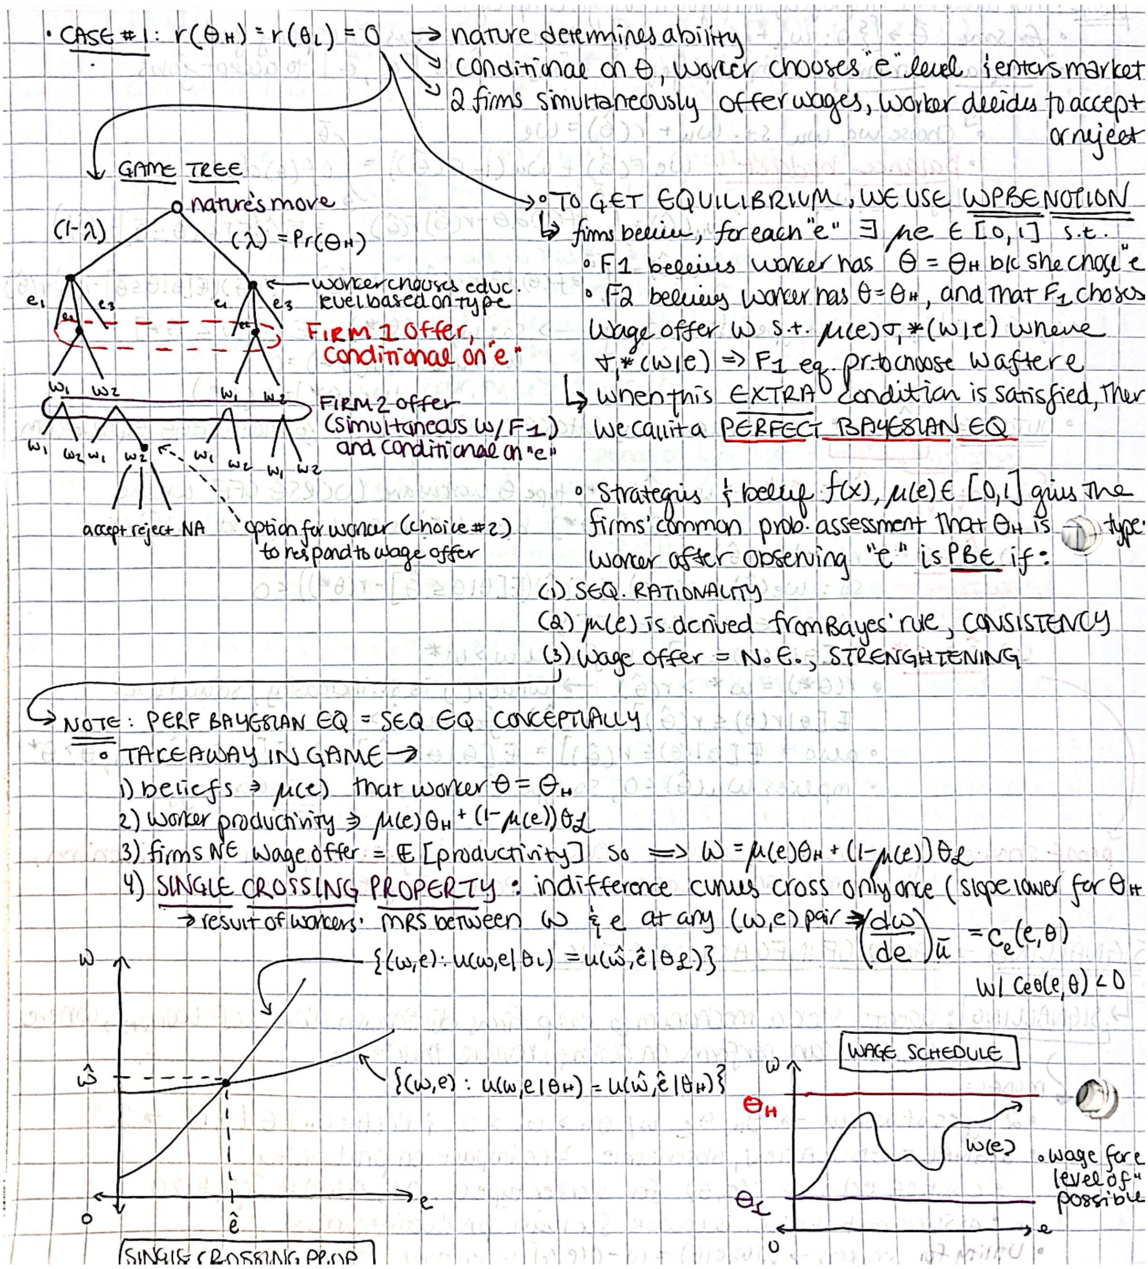

Figure 5: Example of a signalling game and solutions

Definition A pooling equilibrium occurs in a game with incomplete information. All types of players choose the same action, making it impossible for the uninformed agent, normally the firm, to distinguish the different *types* of workers. This is based only on observed behavior.

What is the implication? This leads to the receiver forming a single belief over all types and then responding accordingly. For example, if both θ_h, θ_l choose the same education level, e^*, then the employer cannot distinguish between them. They will then be forced to offer the same wage. For a pooling equilibrium to hold - or to be *credible* - the receiver's beliefs and the responses given to those beliefs need to be **consistent**.

- Must believe all types choose the same action
- No type should have an incentive to deviate from the pooling action

Takeaway: A pooling equilibrium is weakly pareto dominated by a no-signalling outcome.

4.4 Problem types

1. **Midterm 2022, Problem 3:** *Two shy graduate students wish to attend an economics conference. There are two conferences, one on microeconomics (m) and the other on macroeconomics (M). They have become friends recently, so neither knows the preferences of the other. Student $i \in \{1, 2\}$ may prefer macroeconomics ($t_i = t^M$) or microeconomics ($t_i = t^m$). Each can choose to go to either conference ($a_i \in \{M, m\}$). Hence, there are four possible states for their preferences or types:*

$$(t_1, t_2) = (t^m, t^m); (t^m, t^M); (t^M, t^m); (t^M, t^M).$$

Only player i knows their preferences.

It is common knowledge that the probability that player 1 prefers microeconomics is p, and that the probability that player 2 prefers microeconomics is q. Whether student i likes micro over macro is independent of player j's preferences. $(a_1, a_2) \in \{(m, m); (m, M); (M, m); (M, M)\}$ denotes the pair of actions taken by the students (where they go). Student i's utility is 0 if they go alone to a conference (whichever it may be); 2 when they are accompanied by their friend and to their preferred conference, and 1 if they go with their friend to their least-preferred conference.

   (a) *How many pure-strategy profiles are there in this Bayesian game?*

   (b) *Is there at least one Nash equilibrium that holds for all possible values of (p, q)? Justify your answer.*

   (c) *Find the set of values for (p, q) under which $((m, M), (m, M))$ (i.e., each student attends their preferred conference), constitutes a Nash equilibrium.*

Answer (a): Each player has 4 pure strategies, each of which indicate the player's action under both preference types: $(m, m), (m, M), (M, m)$, and (M, M), where the first letter represents the player's action when he prefers microeconomics and the second shows his actions when he prefers macroeconomics. As such, there are 4 x 4 = 16 strategy profiles in this Bayesian game.

Answer (b): The strategy profiles $((m, m), (m, m))$ and $((M, M), (M, M))$ both constitute a nash equilibrum, for any set of beliefs, (p, q). Let's show this by setting a strategy, $s' = (m, m)$, and $s'' = (M, M)$, and showing that a unilateral deviation does not yield an increase in utility:

$$u_i(m; s'_{-i}|t^m) = 2$$
$$u_i(m; s'_{-i}|t^M) = 1 \quad \text{where } u_i(M; s'_{-i}|t^M) = 0$$

From this, we see that the mutual best response is to go to the conference where the other person has determined they will go, $s'_1 = (m, m) \implies s'_2 = (m, m)$. This holds as well in the case of (M, M).

Answer (c):

Find the values for (p, q) under which both players attending their preferred conference is a Nash equilibrium. With symmetry, we only need to solve for one player's type of preference. Choose Player 1. Player 1 prefers m. Player 1's expected utility:

$$EU_1(m) = q2 + (1 - q)0 = 2 \qquad EU_1(M) = q0 + (1 - q)1 = 1 - q$$

Set the two expected utilities equal to each other to find the indifference point:

$$2q = 1 - q \to q = \frac{1}{3}$$

In order to stay in her *preferred strategy*, m, she would need $2q \geq 1 - q \implies q \geq \frac{1}{3}$. By symmetry, Player 2 would choose the macro conference, M, with the same probability, i.e. $p \geq \frac{1}{3}$. By symmetry, we say that for Player 2, they would choose M if $q \leq \frac{2}{3}$. Therefore,

$$q \in \left[\frac{1}{3}, \frac{2}{3}\right] \quad p \in \left[\frac{1}{3}, \frac{2}{3}\right] \quad \to \quad \text{Nash Equilibrium}$$

2. **Outside reference:** *Linear Cournot model with two firms - can we show dominance solvable? Explain.*

Answer:

SET UP: Consider three firms, indexed by $i \in \{1, 2, 3\}$, choosing $q_i \geq 0$. The market price is given by:

$$P(Q) = a - bQ, \quad \text{where } Q = q_1 + q_2 + q_3,$$

with $a, b > 0$. Each firm incurs a constant marginal cost $c \geq 0$.
The profit for firm i is:

$$\pi_i(q_i, q_{-i}) = [a - b(q_i + q_{-i}) - c]q_i,$$

where $q_{-i} = q_j + q_k$ is the total output of the other two firms.

Best Response Function

Firm i chooses q_i to maximize its profit given q_{-i}:

$$\max_{q_i \geq 0} \pi_i(q_i, q_{-i}) = \max_{q_i \geq 0} [a - b(q_i + q_{-i}) - c]q_i.$$

Taking the FOC:

$$\frac{\partial \pi_i}{\partial q_i} = a - 2bq_i - bq_{-i} - c = 0.$$

Solving for q_i, we obtain the best response function:

$$BR_i(q_{-i}) = \frac{a - c}{2b} - \frac{q_{-i}}{2}.$$

Dominance Solvability

Since $q_i \geq 0$, the largest quantity that a firm would choose occurs when the other firms produce nothing, i.e., $q_{-i} = 0$. In this scenario,

$$q_i = \frac{a - c}{2b}.$$

Producing any quantity $q_i > \frac{a-c}{2b}$ would lead to a negative market price, resulting in negative revenue or profit. Therefore, quantities $q_i > \frac{a-c}{2b}$ are strictly dominated and can be eliminated.

With $q_j, q_k \leq \frac{a-c}{2b}$, the maximum total output of the other firms is:

$$q_{-i} \leq 2 \times \frac{a - c}{2b} = \frac{a - c}{b}.$$

Substituting this into the best response function:

$$q_i = \frac{a - c}{2b} - \frac{q_{-i}}{2} \geq \frac{a - c}{2b} - \frac{a - c}{2b} = 0.$$

This implies that the minimum quantity firm i can produce is zero.

As each firm's optimal output q_i depends directly on the total output of the other two firms, $q_{-i} = q_j + q_k$, this variable is not fixed and can take any value within a feasible range, specifically $[0, \frac{a-c}{b}]$. Thus, we cannot further restrict q_{-i} using dominance. Indeed, without knowing the exact split of production between q_j and q_k, or having more detailed assumptions about their behavior, refining the bound for q_{-i} is not possible.

Comment: In the Cournot model with $N = 2$, we can use iterative elimination of strictly dominated strategies to reach a unique outcome. This is because each firm's decision depends only on the quantity produced by the other firm. However, with $N = 3$, each firm's choice depends on the combined output of the other two. This makes it harder to eliminate further strategies after the first step, as we can't precisely determine what the other firms will produce using dominance alone. As a result, the game isn't dominance solvable.

3. **Outside reference:** *Friend and Foe game:*

- "Friend Game" (played with probability p):

	h	t
H	$3,1$	$0,0$
T	$2,1$	$1,0$

- "Foe Game" (played with probability $1-p$):

	h	t
H	$3,0$	$0,1$
T	$2,0$	$1,1$

Applying definition of a pure strategy Bayesian nash equilibrium

The row player (Player 1) chooses between strategies H and T without knowing whether they are playing the "Friend" or "Foe" game. The column player (Player 2), who knows the game type, selects a pure strategy from the set $S_2 = \{hh, ht, th, tt\}$. Each pair in S_2 specifies Player 2's actions in the two games: for example, ht indicates that Player 2 plays h in the "Friend" game and t in the "Foe" game.

Step 1: Compute Normal Form Representation:

- $\{H, hh\}$
  The row payoff in the Friend game is 3, and the column payoff in the Friend game is 1, as the strategy pair (H, h) yields payoffs $(3, 1)$. In the Foe game, the row payoff is 3 and the column payoff is 0, since (H, h) results in $(3, 0)$.
  Thus, the expected row payoff is

$$u_1(H, hh) = 3p + (1-p) \cdot 3 = 3.$$

  The expected column payoff is

$$u_2(H, hh) = p \cdot 1 + (1-p) \cdot 0 = p.$$

- $\{H, ht\}$
  In the Friend game, the row player earns 3 and the column player earns 1 from (H, h). In the Foe game, the row player earns 0 and the column player earns 1 from (H, t). Therefore,

$$u_1(H, ht) = 3p, \quad u_2(H, ht) = 1.$$

- $\{H, th\}$
  In the Friend game, the row player earns 0 and the column player earns 0 from (H, t). In the Foe game, the row player earns 3 and the column player earns 0 from (H, h). Therefore,

$$u_1(H, th) = 3(1-p), \quad u_2(H, th) = 0.$$

- $\{H, tt\}$
  In the Friend game, the row player earns 0 and the column player earns 0 from (H, t). In the Foe game, the row player earns 0 and the column player earns 1 from (H, t). Therefore,

$$u_1(H, tt) = 0, \quad u_2(H, tt) = 1-p.$$

- $\{T, hh\}$
  In the Friend game, the row player earns 2 and the column player earns 1 from (T, h). In the Foe game, the row player earns 2 and the column player earns 0 from (T, h). Therefore,

$$u_1(T, hh) = 2, \quad u_2(T, hh) = p.$$

- $\{T, ht\}$
  In the Friend game, the row player earns 2 and the column player earns 1 from (T, h). In the Foe game, the row player earns 1 and the column player earns 1 from (T, t). Therefore,

$$u_1(T, ht) = 1+p, \quad u_2(T, ht) = 1.$$

- $\{T, th\}$
  In the Friend game, the row player earns 1 and the column player earns 0 from (T, t). In the Foe game, the row player earns 2 and the column player earns 0 from (T, h). Therefore,
  $$u_1(T, th) = 2 - p, \quad u_2(T, th) = 0.$$
- $\{T, tt\}$
  In the Friend game, the row player earns 1 and the column player earns 0 from (T, t). In the Foe game, the row player earns 1 and the column player earns 1 from (T, t). Therefore,
  $$u_1(T, tt) = 1, \quad u_2(T, tt) = 1 - p.$$

The resulting table is:

	hh	ht	th	tt
H	$3, p$	$3p, 1$	$3 - 3p, 0$	$0, 1 - p$
T	$2, p$	$1 + p, 1$	$2 - p, 0$	$1, 1 - p$

Step 2: Find Indifference Condition:

First of all, I will eliminate dominated strategies for Player 2:

- In the "Friend" game, h is strictly better than t, so Player 2 will never play t in this game.
- In the "Foe" game, t is strictly better than h, so Player 2 will never play h in this game.

Thus, Player 2's optimal candidate strategies are:

$$S_2 = \{ht, tt\}.$$

Now, let us turn to Player 1. Player 1 does not know the game type but believes the Friend game occurs with probability p. Their expected payoffs depend on Player 2's strategies. Since Player 2 will never play a dominated strategy, the relevant strategies to assess are Player 2's optimal candidate strategies.

- If Player 2 plays ht:
  $$u_1(H, ht) = 3p + 0(1 - p) = 3p,$$
  $$u_1(T, ht) = 2p + 1(1 - p) = 2p + 1 - p = 1 + p.$$
  Indifference:
  $$u_1(H, ht) = u_1(T, ht),$$
  $$3p = 1 + p \implies p = \frac{1}{2}.$$
  Player 1's preferences are as follows:
  $$\begin{cases} H & \text{if } p > \frac{1}{2}, \\ T & \text{if } p < \frac{1}{2}, \\ \text{indifferent} & \text{if } p = \frac{1}{2}. \end{cases}$$
- If Player 2 plays tt:
  $$u_1(H, tt) = 0p + 0(1 - p) = 0,$$
  $$u_1(T, tt) = 1p + 1(1 - p) = 1.$$
  In this case, Player 1 always prefers T:
  $$u_1(T, tt) = 1 > u_1(H, tt) = 0.$$

Since Player 2's strategy tt results in a deterministic response from Player 1 (T), the critical probability p arises solely from the case where Player 2 plays ht. Thus, the relevant threshold is $p = \frac{1}{2}$.

Step 3: BNE:

If $p > \frac{1}{2}$, Player 1 plays H, and Player 2's best response is ht. If $p < \frac{1}{2}$, Player 1 plays T, and Player 2's best response is tt. If $p = \frac{1}{2}$, Player 1 is indifferent and randomizes between H and T, while Player 2 is indifferent and can randomize between ht and tt.

The resulting BNEs under definition of pure strategy bayesian nash are:

$$\begin{cases} (H, ht) & \text{if } p > \frac{1}{2}, \\ (T, tt) & \text{if } p < \frac{1}{2}, \\ (H, ht) \text{ and } (T, tt) & \text{if } p = \frac{1}{2}. \end{cases}$$

Applying Definition of the strategy profile under BNE

Step 1: Player 2's Optimal Actions:

Player 2 (the column player) knows the game type and chooses the action that maximizes their payoff based on the type.

- In the "Friend" game, h is strictly better than t.
- In the "Foe" game, t is strictly better than h.

Player 2's strategy function is therefore:

$$s_2(\theta_2) = \begin{cases} h & \text{if } \theta_2 = \text{“Friend”}, \\ t & \text{if } \theta_2 = \text{“Foe”}. \end{cases}$$

This corresponds to the strategy $s_2 = ht$.

Step 2: Indifference Condition

Player 1 (the row player) does not know the game type but believes the Friend game occurs with probability p. Their expected payoffs for actions H and T, given $s_2 = ht$, are:

$$u_1(H, ht) = 3p + 0(1 - p) = 3p.$$

$$u_1(T, ht) = 2p + 1(1 - p) = 2p + 1 - p = 1 + p.$$

The indifference condition is:

$$\begin{aligned} u_1(H, ht) &= u_1(T, ht), \\ 3p &= 1 + p, \\ p &= \frac{1}{2} \end{aligned}$$

Then,

$$\begin{cases} H & \text{if } p > \frac{1}{2}, \\ T & \text{if } p < \frac{1}{2}, \\ \text{indifferent} & \text{if } p = \frac{1}{2}. \end{cases}$$

Step 3: BNE:

From step 1 and step 2, the resulting BNEs under definition of a pure strategy bayesian nash are:

$$\begin{cases} (H, ht) & \text{if } p > \frac{1}{2}, \\ (T, ht) & \text{if } p < \frac{1}{2}, \\ (H, ht) \text{ and } (T, ht) & \text{if } p = \frac{1}{2}. \end{cases}$$

Remark: The results from the two definitions differ due to the different assumptions about the informational structure. In the first definition which considers the game from an ex-ante standpoint of expected utilities, does not assumed that Player 2 knows which game they are playing. In contrast, our definition of the pure strategy BNE assumes that Player 2 has full knowledge of the game type.

4. **Outside reference**: *Show that a strategy profile σ is a Nash equilibrium of finite extensive form game G if and only if there exists a system of beliefs μ such that:*

   (i) *The strategy profile σ is sequentially rational given μ at all information sets reached with positive probability under σ.*

(ii) *The system of beliefs μ is derived from σ through Bayes' rule whenever possible.*

What is the difference between these conditions and the conditions that define a weak perfect Bayesian equilibrium?

Answer:

(i) $\implies$ (ii)

If σ is a NE, no player can improve their payoff by unilaterally deviating from their strategy. If condition (i) is violated, there must exist an information set reached with positive probability under σ where a player is not playing a best response to the strategies of the other players. This contradicts the definition of a NE, as the player would have an incentive to deviate.

If condition (ii) is violated, players may have beliefs inconsistent with the strategies being played at on-path information sets. Since NE implicitly requires players to act optimally given the strategies of others, consistent beliefs derived via Bayes' Rule are necessary where possible. Inconsistent beliefs could lead to suboptimal actions and incentives to deviate, breaking the equilibrium.

Therefore, if σ is a NE, both (i) sequential rationality and (ii) beliefs consistency must hold.

(i) $\impliedby$ (ii)

If σ satisfies sequential rationality, then at every information set reached with positive probability under σ , players are choosing their best responses to the strategies of others. This ensures that no player has an incentive to unilaterally deviate at any on-path information set.

Belief consistency ensures that the belief system μ is derived from σ using Bayes' Rule wherever applicable, guaranteeing that players' expectations about the game's state are consistent with the strategies being followed.

Then, if (i) and (ii) hold, σ is a NE.

In a NE, sequential rationality is required only at information sets reached with positive probability under σ (on-path information sets). In contrast, a Weak PBE requires sequential rationality at all information sets, including those off the equilibrium path. Additionally, while a NE requires belief consistency only for on-path information sets, a Weak PBE permits arbitrary beliefs at off-path information sets.

5 Static oligopoly models

MWG: 12

5.1 Bertrand Competition

In the Bertrand game, $N \geq 2$ firms simultaneously set prices, competing to offer the lowest price. Customers purchase from the firm with the lowest price, splitting equally among firms in case of a tie. Each firm incurs no fixed costs and has a marginal cost $MC = c > 0$. The inverse demand function $Q(p)$ is continuous and decreasing.
A firm's profit is defined as:

$$\pi(p) = (p - c)Q(p), \quad p \in [c, p^m]$$

where p^m is the monopoly price, and c is the marginal cost which corresponds to the competitive price. Profits are increasing in prices within the range $p \in [c, p^m]$.

Definition: If the discount factor satisfies $\delta > 1 - \frac{1}{n}$, then for any price $p^* \in [c, p^m]$, there exists a SPNE where all firms set their price at p^*.
Propositions:

- **Subgame Perfect Nash Equilibrium in Bertrand Duopoly**: The strategies described for price-setting in Bertrand duopoly constitute a **SPNE** (see 3.3) if and only if $\delta \geq 1/2$.
- **SPNE in Infinitely Repeated Bertrand Duopoly**: In an infinitely repeated Bertrand duopoly game, when $\delta \geq 1/2$, repeated choice of any price $p \in [c, p_m]$ can be supported as an SPNE outcome path using Nash reversion strategies.

5.2 Cornout Competition

In the Cournot competition, $n \geq 2$ firms simultaneously compete by setting quantities. The stage game is a Cournot game with no fixed costs and constant marginal cost $MC = c \geq 0$. The inverse demand function is $P(Q) = 1 - Q$.
For $n = 2$, the equilibrium quantities under competitive and monopoly outcomes are:

$$q^c = \frac{1-c}{3}, \quad q^m = \frac{1-c}{2}.$$

Each firm's profit is given by:

$$\pi(q) = (P(Q) - c)q, \quad q \in [q^c, q^m],$$

where q^c represents the competitive quantity and q^m the monopoly quantity.

Proposition: There exists an SPNE where firms collude on the monopoly outcome if δ is sufficiently large.

Proof:

Consider the following strategy:
- If no firm deviated, each firm sets $q = \frac{Q^m}{2}$.
- If any firm deviated, all firms set $q = q^c$.

Let's verify whether this strategy is an SPNE by evaluating incentives to deviate:

1. No past deviation(s)
   - By following the equilibrium strategy, each firm produces $\frac{Q^m}{2}$ and earns a per-period payoff of $\frac{\pi(Q^m)}{2}$. The discounted sum of payoffs is:

     $$\sum_{t=0}^{\infty} \delta^t \frac{\pi(Q^m)}{2} = \frac{1}{2}\frac{\pi(Q^m)}{1-\delta}.$$

     Substituting,

     $$\frac{(1-c^2)}{8(1-\delta)}$$

   - A firm deviating from the equilibrium strategy maximizes its one-period payoff by choosing q_i to capture as much profit as possible, given the other firm produces $\frac{Q^m}{2}$. The total payoff from deviation is:

     $$\max_{q_i} \left(P\left(\frac{Q^m}{2} + q_i\right) - c \right) q_i + \sum_{t=1}^{\infty} \delta^t \pi(q^c).$$

     where:
     - $\max_{q_i} \left(P\left(\frac{Q^m}{2} + q_i\right) - c \right) q_i$, represents the firm's profit in the current period. The firm chooses a quantity q_i to maximize its payoff, where $P\left(\frac{Q^m}{2} + q_i\right)$ is the price it can charge, depending on its own production q_i and the quantity produced by the other firm, which is $\frac{Q^m}{2}$. The term $\left(P\left(\frac{Q^m}{2} + q_i\right) - c \right)$ gives the firm's profit margin, and multiplying by q_i gives the total profit for that period.
     - $\sum_{t=1}^{\infty} \delta^t \pi(q^c)$, represents the future discounted payoffs for all periods after the deviation. If the firm deviates, the other firms will revert to the competitive quantity q^c, and the firm's future profit will be $\pi(q^c)$, the payoff in the competitive equilibrium.

     Total payoff from deviating,

     $$\frac{9(1-c)^2}{64} + \frac{\delta}{1-\delta} \cdot \frac{(1-c)^2}{9}$$

   Comparing the two, there are no incentives to deviate for:

   $$\delta \geq \frac{9}{17}$$

2. Past deviation(s)

- After a deviation, the equilibrium strategy prescribes that all firms produce $q = q^c$, earning zero profits in every subsequent period:

$$\pi(q^c) = 0.$$

- Any deviation from the equilibrium strategy leads to a payoff no greater than zero. If the firm sets $q > q^c$, it will earn zero profits because the price will fall below marginal cost. Conversely, if the firm sets $q < q^c$, it will incur negative profits due to losses on production. Therefore, in both cases, the payoff from deviating is at most zero.

There are no incentives to deviate from the equilibrium strategy here.

In conclusion, the condition on δ to sustain the collusive equilibrium is $\delta \geq \frac{9}{17}$.[4]

5.3 Repeated Games

A repeated game is a game in which a stage game (a base or one-shot game) is played multiple times by the same players, possibly infinitely, and where players observe past actions and may condition their future strategies on this history. The game is set up with the following parameters:

1. Stage game: a one period or one shot game, usually presented as a normal form game
2. Set of stages: $t = 1, ...T \leq \infty$ in discrete time
3. Ending rule: when $T \leq \infty$ is known with certainty to all players, and when there is random end, with probability, $p \leq 1$, the game continues
4. Observed history of play, H_{t-1}
5. Strategies $(\sum_i)$ for player i where a pre strategy, S_i is a sequence of functions, $\{s_{it}(H_{t-1})\}_{t=1}^{T}$
6. Π earned each period, and accumulates such that $s = (s_i, ...s_i)$ induces the outcome path, $Q(s)$ as a sequence of actions, with a discounted payoff accruing from that outcome path:

$$v_i(Q) = \sum_{\tau=0}^{T} \delta^\tau \Pi_i(q_{1+\tau}) \quad \forall \quad T \leq \infty$$

Definition: Nash Reversion Strategy in Repeated Games
A strategy profile $s = (s_1, ..., s_I)$ in an infinitely repeated game is one of **Nash reversion** if each player's strategy calls for playing some outcome path Q until someone defects and playing the stage game **Nash equilibrium** $q^* = (q_1^*, ..., q_I^*)$ thereafter.

Lemma: Nash reversion strategy profile that calls for playing the outcome path $Q = \{q_{1,t}, ...q_{l,t}\}_{t=1}^{\infty}$ prior to any deviations is an SPNE if and only if:

$$\hat{\Pi}_j(q_{-j,t}) + d/1 + d\Pi_j(q^*) \leq v_j(Q_t)$$

Where $q_{-j,t}$ denotes player j's pay off from deviation, while q^* denotes the stage game nash equilibrium quantity. This lemma expresses the tradeoff between one period gains and future losses, which serve to tell us why a player would deviate for just one period.

Recall a nash reversion is a type of punishment mechanism that is used in repeated games. It occurs when players start by playing an agreed upon outcome that yields higher payoffs, but when one player deviates, the game *reverts* to a nash equilibrium of the stage game. We call this a reversion because players return to a baseline nash equilibrium of the one shot game as a punishment for deviation.

Propositions:

[4] Comments:
- If the static Cournot equilibrium is unique, cooperation cannot be sustained in SPNE for finitely repeated games.
- Unlike the Bertrand model, Cournot competition allows partial cooperation with lower discount factors. For any $\delta > 0$, profits exceeding $\pi(q^c)$ are achievable.

- **Folk Theorem for Repeated Games**: In an infinitely repeated game, any feasible discounted payoffs that give each player, on a per-period basis, more than the lowest payoff that they could guarantee in a **single-play simultaneous-move component game**, can be sustained as the payoffs of an **SPNE** if players discount the future sufficiently.

- **Sustaining Cooperation in Repeated Games**: If there is potential for joint improvement in payoffs near the stage game **Nash equilibrium**, then for sufficiently high δ, there exists a SPNE where payoffs exceed those in the static **Nash equilibrium**.

- **Credible Punishment in Repeated Games**: If each player's stage game Nash equilibrium payoff strictly exceeds their minimax payoff, then there exists an SPNE with more severe punishments than Nash reversion

Lemma 1: Nash Reversion Strategies and SPNE: A Nash reversion strategy profile that calls for playing outcome path $Q = \{q_{1t}, ..., q_{It}\}_{t=1}^{\infty}$ prior to any deviation is an **SPNE** if and only if:

$$\hat{\pi}_i(q_{-it}) + \frac{d}{1-d}\pi_i(q^*) \leq v_i(Q,t), \quad \forall t, i$$

where $\hat{\pi}_i(\cdot)$ denotes player i's payoff from deviating and q^* is the stage game Nash equilibrium.

6 Appendix

6.1 Fundamental welfare theorems

1. In a perfectly competitive market, any equilibrium allocation of resources is pareto efficient (i.e. no individual can be made better off without making someone else worse off).

2. Any pareto efficient allocation can be achieved by having a competitive market equilibrium, provided that the appropriate redistribution of initial endowments is implemented (i.e. markets make efficient outcomes where the initial distribution of resources can bve adjusted).

6.2 Other theorems

6.2.1 Principal-Agent problem

Definition: A principal–agent problem models a situation where a principal relies on an agent to act on their behalf under asymmetric information, and must design incentives to align the agent's behavior with their own interests. The purpose of this type of problem is that the principal must design an **incentive-compatible contract** or mechanism to induce the agent to act in the principal's interest

6.2.2 Weierstrass Theorem

Theorem: Let $f : S \to \mathbb{R}$ be a continuous real-valued function, where $S \subseteq \mathbb{R}^n$ is a non-empty, compact set. Then f attains its maximum and minimum on S. That is, there exist $x^{,}x_{\in}S$ such that:

$$f(x_{)}) \leq f(x) \leq f(x^{)} \quad \text{for all } x \in S.$$

Definition: Any continuous function defined on a non-empty, compact set has both a maximum and a minimum value. This theorem is necessary for proving the existence of nash equilibria under standard conditions, by ensuring that the utility function has a maximum.

6.2.3 Core mathematical properties:

- **Convex Set**: A set $X \subseteq \mathbb{R}^n$ is **convex** if, for any $x, y \in X$ and $\alpha \in [0,1]$, the combination

$$\alpha x + (1-\alpha)y \in X.$$

- **Quasiconcave Function**: A function $f : \mathbb{R}^n \to \mathbb{R}$ is **quasiconcave** if, for any $x, y \in X$ and $\alpha \in [0,1]$,

$$f(\alpha x + (1-\alpha)y) \geq \min\{f(x), f(y)\}.$$

- **Correspondence**: A **correspondence** $r : X \rightrightarrows Y$ is a set-valued function from X to subsets of Y.
  - It is **convex valued** if, for all $x \in X$, $r(x)$ is convex.
  - It is **upper hemicontinuous** if, for every sequence $x^t \to x$ and $y^t \to y$ with $y^t \in r(x^t)$ for all t, we have $y \in r(x)$.

    (Intuition: A correspondence $r : X \rightrightarrows Y$ is upper hemicontinuous if the set $r(x)$ "shrinks nicely" as x approaches a limit. Specifically, if $x^t \to x$ and $y^t \in r(x^t)$ with $y^t \to y$, then $y \in r(x)$. This ensures that $r(x)$ does not include any "new" elements not approached by $r(x^t)$.)

- **Compact Set**: A set $K \subseteq \mathbb{R}^n$ is **compact** if it is both:
  - **Closed**: K contains all its boundary points.
  - **Bounded**: K fits within some finite region of $\mathbb{R}^n$.

  In other words, K is compact if it is a closed and bounded subset of $\mathbb{R}^n$.

- **Bernoulli distribution:** A discrete random variable X follows a Bernoulli distribution with parameter p, where $0 \leq p \leq 1$, if it takes values in $\{0, 1\}$ with probability mass function (PMF) given by:

$$P(X = x) = \begin{cases} p, & \text{if } x = 1, \\ 1 - p, & \text{if } x = 0. \end{cases}$$

  This is denoted as:

$$X \sim \text{Bern}(p).$$

  The expected value and variance of X are:

$$E[X] = p, \quad \text{Var}(X) = p(1 - p).$$

6.2.4 Kakutani's Fixed Point Theorem

Theorem 2: Let $X \neq \emptyset$ be a compact and convex subset of $\mathbb{R}^K$, and let $r : X \rightrightarrows X$ be a correspondence such that:

- r is non-empty valued,
- r is convex valued,
- r is upper hemicontinuous.

Then, there exists $x \in X$ such that $x \in r(x)$, i.e., r has a fixed point.

Mathematical Formula	Description of Items	Description of the Game
Normal Form (Mixed Game)		
$\Gamma_N = \left[I, \{\Delta(S_i)\}, \{u_i(\cdot)\}\right]$	• I: Set of players • $\Delta(S_i)$: Set of mixed strategies for player i, where S_i is the set of pure strategies • u_i: Utility function of player i	Describes a game where players choose mixed strategies and payoffs are determined based on mixed strategies in normal form.
Normal Form (Pure Game)		
$\Gamma_N = \left[I, \{S_i\}, \{u_i(\cdot)\}\right]$	• I: Set of players • S_i: Set of pure strategies for player i • u_i: Utility function of player i	Describes a game where players choose pure strategies, and payoffs are determined based on these strategies in normal form.
Extensive Form (Mixed Game)		
$\Gamma_E = \left[I, T, P, \{\Delta(S_i)\}, \{u_i(\cdot)\}\right]$	• T: Set of decision nodes • P: Player assignment function for decision nodes • $\Delta(S_i)$: Set of mixed strategies • u_i: Utility function of player i	Describes a game represented as a tree, where players make decisions at each node, and mixed strategies determine probabilities of actions.
Extensive Form (Pure Game)		
$\Gamma_E = \left[I, T, P, \{S_i\}, \{u_i(\cdot)\}\right]$	• T: Set of decision nodes • P: Player assignment function for decision nodes • S_i: Set of pure strategies • u_i: Utility function of player i	Describes a game represented as a tree, where players make decisions at each node, and strategies are pure.

Table 3: Comparison of Mixed and Pure Games in EFG and NFG

Concept	Type	Game Type	Short Description
Rationality	Condition	Perfect, Imperfect, Adverse Selection, Signaling, Repeated	Players choose strategies maximizing expected payoffs given beliefs.
Nash Equilibrium (NE)	Equilibrium Concept	Perfect, Imperfect, Adverse Selection, Signaling, Repeated	No player can gain by unilaterally deviating from equilibrium strategies.
Bayesian Nash Equilibrium (BNE)	Equilibrium Concept	Imperfect info, Adverse Selection, Signaling	NE for incomplete information; strategies maximize expected payoffs given beliefs about others' types.
Subgame Perfect Nash Equilibrium (SPNE)	Equilibrium Refinement	Perfect info (clearly defined subgames), Repeated	NE refinement requiring equilibrium within every subgame (no non-credible threats). In repeated games, credible equilibrium strategies specifying optimal actions after every possible history, often using punishment.
Weak Perfect Bayesian Equilibrium (Weak PBE)	Equilibrium Refinement	Imperfect info, Signaling, Adverse Selection	Equilibrium requiring sequential rationality and weak consistency of beliefs (Bayesian updating where possible).
Weak Sequential Equilibrium (WSE)	Equilibrium Refinement	Imperfect info, Signaling, Adverse Selection, Repeated	Like WPBE but off-path beliefs must arise as limits of totally-mixed trembles; combines sequential rationality with tremble-limit consistency.
Nash Reversion Strategies	Strategy/ Punishment Method	Repeated Games	Punishment reverting permanently to stage-game NE after deviation; sustains cooperation.
Sequential Rationality	Condition	Perfect, Imperfect, Signaling, Adverse Selection, Repeated	Requires optimal actions at every decision point given beliefs. This condition must be met for a sequential equilibrium and WPBE.
Backward Induction	Method	Perfect info, Finite Repeated Games	Solves perfect-info dynamic games backward from end to beginning, finds SPNE.
Generalized Backward Induction	Method	Imperfect info, Signaling, Adverse Selection, Repeated	Extends backward induction with explicit beliefs for imperfect-info games; finds PBE.
Pareto Dominance	Equilibrium Criterion	Perfect, Imperfect, Signaling, Repeated	Outcome where no player can be better off without making another worse off; selects among multiple equilibria.

Table 4: Summary of key concepts on equilibrium, strategies and refinements in game theory.