

INTRODUCTION

What is Game Theory¹

Game Theory is a particularly useful methodological tool used to analyze and predict the outcome of specific types of situations, namely situations of strategic interactions or strategic interdependencies. These situations involve the interaction of multiple decision makers (agents or players) where (i) each agent -acting rationally- is trying to maximize her well-being, and (ii) each agent's action has consequences on other agents' well-being. This interdependence causes each player to consider the other player's possible decisions when deciding her own strategy/action². Agents might try to resolve the situation by acting individually (non-cooperatively), or they might try to coordinate their actions (cooperatively). We will only consider non-cooperative game theory in this class. The goal is to use Game Theory to *ex ante* identify the (possible) outcome of the game, *i.e.*, the actions/strategies chosen by all players and the resulting payoffs for each one of them. We will be calling this predicted outcome a **solution of the game**.

For instance, consider the market dynamics between two neighboring hot-dog vendors, Ms. A and Mr. B., who are located next to each other in a popular place. Their close proximity and the similarity of their products result in the situation meeting the aforementioned criteria (i) and (ii): Ms. A's decision over her own price will affect her own profits but, by changing her market share, it will also affect the profits of Mr. B. Obviously, a similar argument can be made for Mr. B (*e.g.*, Bertrand type competition). This illustrates a game set up!

On the other hand, when there are many firms in a market, all producing an identical product, there is no case of strategic interactions: every single firm realizes that it has no power to affect the market price and, hence, it cannot affect the profits of any other firm (*e.g.*, Perfect Competition). This example should warn you that having multiple agents involved in a situation does not necessarily imply that this situation is characterized by strategic interdependency, and, in such a case, Game Theory is not an appropriate analysis tool. This is not a game!

¹ Game Theory is fun but not exactly related to frivolous activities.

² In such situations traditional optimization fails.

Why is Game Theory important to us? First, we should understand that trying to analyze situations involving strategic interdependencies is a complicated task. The degree of complexity increases the more the elements of interactions one must consider. In the example with the hot-dog sellers, Mrs. A and Mr. B, it is understandable that their individual profits will depend on the price choices both make. If we assume that each can choose two prices, say “High” or “Low”, there will be four possible outcomes³: one where both choose a low price, one where both choose a high price, and two outcomes where one of them chooses low while the other chooses high price. Therefore, the analysis of this situation involves some sort of comparison of the four different outcomes⁴. By changing the price options of the two competitors from two to three (for example, setting a “low”, “medium”, or “high” price), the number of potential outcomes to be considered in the analysis becomes nine while if the price options for each seller are four the potential outcomes become sixteen. This exponential growth in the number of potential outcomes underscores the challenge of comprehensively evaluating strategic interactions without a formal framework. Without Game Theory, analysts risk drowning in the myriad possibilities, rendering effective decision-making difficult. Thus, Game Theory offers a systematic methodology to disentangle the complexities of strategic interdependencies, enabling a more efficient and insightful analysis of such scenarios.

Moreover, beyond managing complexity, Game Theory provides invaluable insights into strategic decision-making processes, offering a roadmap for understanding and addressing similar situations in diverse contexts. By dissecting the dynamics of strategic interactions and identifying optimal strategies and potential outcomes, Game Theory equips decision-makers with a toolkit to anticipate and respond effectively to competitive pressures. This systematic approach not only enhances decision-making capabilities but also fosters a deeper understanding of the underlying mechanisms governing strategic behavior. Consequently, Game Theory serves as a powerful tool for both theoretical analysis and practical applications, empowering individuals and organizations to navigate the intricate landscape of strategic interactions with confidence and clarity.

Applications of Game theory in Economics

Game Theory, as a branch of applied mathematics, has been widely used in economics, the field within social sciences where quantitative methods play a very crucial role. Indicatively, the following are some examples where Game Theory is used in economics:

³ Two players, each having to choose between two available actions implies $2 \times 2 = 4$ potential outcomes. In general, for a case where two players are involved with one player having k and the other player having m actions, the number of potential outcomes will be $k \times m$.

⁴ The “rules” for making such comparisons of the alternative outcomes will be introduced later.

- In different types of oligopolistic competition one firm's price or quantity choices affect another firm's profits through a shared market demand. The goal is to identify the firms' outputs, price(s), and profits that will be realized in equilibrium (*i.e.*, the solution of the game).

- In all kinds of auctions, one bidder's bid affects another bidder's profit. The goal is to identify the bid of each player (and who gets the auctioned item) and the corresponding payoffs in equilibrium. Understanding how bidders behave in auctions can help, for example, in the design of appropriate auctioning mechanisms to secure the highest revenue for the seller.

- In bilateral and multilateral bargaining processes (negotiations) between firms and trade unions. The goal is to identify the agreed wage, the number of hirings, the welfare of the trade union (or of its members) and the profits of the firms that will be realized in equilibrium.

- In cases of pollution externalities where, by definition, one agent's action affects another agent's well-being. Identifying the way polluting agents behave strategically in such situations can help design appropriate policies to increase efficiency.

- In contract theory where contracts must be designed in an incentive compatible fashion, that is, the group of buyers targeted by the seller should be the one attracted by the terms of the contract (*e.g.*, discounts in car insurance contracts should be awarded to the prudent drivers).

- A situation described as a *Moral Hazard* problem occurs when the actions of one party may change to the detriment of another after a financial transaction has taken place (informational asymmetry, principal-agent problem). Situations describing a moral hazard are characterized by inefficiency. Game theory can analyze the incentives for post-agreement behavior of the agents in the absence of complete and perfect information, thus helping to regain efficiency losses.

- Game theory plays an important role in *market design*, which is the process of designing and implementing market mechanisms that allocate goods and services efficiently. Game theory provides a framework for understanding the strategic interactions between market participants and for designing market mechanisms that induce desirable outcomes depending on the type of the market (*e.g.*, one-sided/two-sided, one-to-one, many-to-one matchings).

- In the case of *Coalition Formation*, players can form groups to coordinate their actions. By understanding the strategic interactions between players and the benefits of forming coalitions, game theorists can propose policies and ways of action that promote cooperation, stability, and efficiency (*e.g.*, international environmental agreements against the climate change, economic integration, coalitional governments when there is no clear majority) or deter coalitions that are socially harmful (*e.g.*, cartels).

Representation of Games

When analyzing situations of strategic interactions (henceforth *games*) we do it not from a specific player's point of view but from that of a bystander. While pieces of information are not available to some player(s) this information is readily available to the bystander. As such, and to be able to analyze it, we must clearly define the following fundamental elements of a game:

1) WHO? → all the agents (players) whose decisions must be taken into consideration.

Not all agents present in a situation of strategic interactions should be considered as "players" of the game. According to the description of what Game Theory is, an agent involved in a case of strategic interdependencies is a "player" if (i) she must choose an action/strategy among many available actions/strategies, and (ii) her well-being is affected by the realization of the game (outcome of the game). For example, there is a situation involving three people whose well-being is affected by the outcome of the game. One of them has only one option (say, to turn her car left). Then this person is irrelevant to the game, and she is not considered a player. Furthermore, imagine a scenario where an individual has the option to select from a multitude of actions, yet regardless of the chosen action, their well-being remains constant or non-existent. While this individual's actions hold significance for the ultimate outcome of the game, they are not classified as a player. In essence, this person could be effectively replaced by a machine that randomly selects one action⁵.

2) WHAT? → the actions and the information available to them.

The set of actions available to each player should be known to all players, *i.e.*, there are no concealed choices for any participant. However, not all information regarding the game should necessarily be available to all players: players might or might not be able to observe the actions taken by players preceding them and/or the payoffs of the other players in any possible outcome of the game. We will discuss later the concepts of complete (and incomplete) and perfect (and imperfect) information.

3) WHEN? → the "protocol" according to which players choose their actions leading to an outcome.

In many games players choose simultaneously (*e.g.*, the game "rock-paper-scissors") while in other games players choose their actions/strategies in a sequential order. Therefore, there are two ways to represent situations of strategic interactions:

- strategic (or normal) form games, where players choose "simultaneously," without knowledge of the choices made by others at the same time, and

⁵ In Bayesian games it is often the case that Nature chooses the type of a player (

- extensive (or sequential) form games, where players choose their actions sequentially according to a specific order (*i.e.*, a “protocol”) leading to a sequential unfolding of the game.

4) HOW MUCH? → their preferences over all possible outcomes.

The payoffs of all players in every possible outcome, as well as the preferences of the players over their own outcomes, should be known to us to properly analyze the game.

In what follows we will be examining four different categories of games. The classification of the non-cooperative games into these four categories relates to specific traits along two dimensions. First, whether a game is static or extensive and, second, according to its information structure. Regarding the latter we should distinguish between

- **complete and incomplete information (*knowledge about the game*)**: complete information refers to a situation where all players have full knowledge of the game's structure, including the rules, the available strategies, and the payoffs associated with each strategy, but some information may be unknown to some players about other players' moves. Incomplete information then refers to a case where some information about the game's structure, including the rules, the available strategies, and the payoffs associated with each strategy, is missing.
- **perfect and imperfect information (*knowledge about the current state of play*)**: perfect information refers to a situation where, if there is a sequence of choices made by the players, each player can observe the choices made by the players preceding her. Imperfect information then refers to the case where some previous moves of players are not visible to the ones that follow.

How do we “solve” a game?

As mentioned before, the goal is to use Game Theory to *ex ante* identify the (possible) outcome of the game, *i.e.*, the actions/strategies chosen by all players and the resulting payoffs for each one of them. This requires understanding the way each player of the game will choose. However, every player of the game understands that the final outcome depends not only on her choice(s) but also on all other players' choices, hence every player should consider how all other players will play! Analyzing a game requires that we make assumptions on how the players deal with this complexity. While there are various assumptions one could consider, they all must lay their foundation on *rationality*: players are rational individuals striving to maximize their own welfare. We will investigate this matter further later on.

PART I: STRATEGIC FORM GAMES OF COMPLETE INFORMATION

As we have discussed, a strategic situation involves several agents where each agent must take at least one action while having specific preferences over the set of potential outcomes (every player being able to compare the payoffs that correspond to her among all possible outcomes of the game). The payoff each agent receives at the end depends not only on her own actions but on the actions of all other players as well.

We first consider strategic or normal form games⁶ of complete information, that is games where

- the players simultaneously choose actions (static game)
- each player in the game is aware of the sequence, strategies, and payoffs throughout gameplay (complete information)

It should be noted here that players need not make decisions simultaneously, that is the games we consider in this part need not be “static”. The games we analyze here can evolve over time having players choosing in a sequential order. However, we assume imperfect information, that is, no player knows any choice of any other player that preceded her. We then define a strategic form game as

Definition: A strategic form game has a finite set of “players”, N , and each player $i \in N$ has a non-empty actions/strategies set $\{A_i\}_{i \in N}$. For each $i \in N$ there is a preference relation \succsim_i on the set $A = \times_{j \in N} A_j$. Hence, a game can be stated as $\langle N, \{A_i\}_{i \in N}, \{\succsim_i\}_{i \in N} \rangle$

Some important notes regarding the definition above:

- In normal form games the terms *action* and *strategy* do not differ. This is no longer true in extensive form games that we will analyze later. For the moment, we should only say that a strategy $s_i \in \{S_i\}_{i \in N}$ of a player $i \in N$ is a function that assigns an action to each point (node or information set) where the specific player must make a choice.
- We define the preference relation of a player i not on the set of his actions alone but rather on the set $A = \times_{j \in N} A_j$, the product of all players’ action sets. The reason is because the choices of others affect the payoffs of player i .
- If A_i is finite for all the players $i \in N$ then the game is finite.
- A utility function $u_i: A \rightarrow \mathcal{R}$, such that $u_i(a) \geq u_i(b)$ if and only if $a \succsim_i b$, may represent preferences. In such a case the game can be stated as $\langle N, \{A_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$.

⁶ Also called static games.

- If a strategic form game consists of $N \leq 3$ players, then it is analytically convenient to be represented using payoff tables (matrices), provided that the game is finite and that the number of actions available to each player is not very large.

In the following examples we will first try to identify the important elements of the game, *i.e.*, the players, the set of actions, the information structure, the order of the play, and the players' preferences over the outcomes.

Example 1 – “The prisoner’s dilemma”: Two members of a criminal gang are arrested and imprisoned. Each prisoner is in solitary confinement with no means of communicating with the other. The prosecutors lack sufficient evidence to convict the pair on the principal charge, but they have enough to convict both on a lesser charge. Simultaneously, the prosecutors offer each prisoner a bargain. Each prisoner is given the opportunity either to betray the other by testifying that the other committed the crime, or to cooperate with the other by remaining silent. As we have seen in many police movies and TV series, the prisoners are warned that the other might confess, and it will be wise for them to hurry. Although it does not have to be the case, we assume here that these two suspects are totally individualistic, *i.e.*, each one cares about his own wellbeing. The possible outcomes are:

- If prisoners 1 and 2 each betray the other, each of them serves six years in prison.
- If prisoner 1 betrays prisoner 2 but prisoner 2 remains silent, prisoner 1 will be set free and prisoner 2 will serve nine years in prison (and vice versa)
- If prisoners 1 and 2 both remain silent, they will both serve only one year in prison (on the lesser charge).

In the scenario above, we first identify the players, namely prisoners 1 and 2. Note that although the story includes prosecutors, they are not players in this set up since (i) they do not choose among available actions of theirs, *i.e.*, their strategy set is a singleton containing only one action (lead the suspects to separate cells and offer each one a specific deal, and (ii) there is no prosecutor’s payoff associated with the outcome of this game. It is clear from the story that each player must choose between two actions, to confess or not.

Furthermore, both prisoners understand the situation they are involved in (*i.e.*, their own and the other person’s action sets), and what is at stake (*i.e.*, how many years of jail time for each one of them in every possible outcome), hence it is a complete information game. However, being in separate interrogation rooms, a prisoner cannot know the decision the other prisoner has taken, hence it is a game of imperfect information.

Finally, there are four possible outcomes clearly described in the story and the payoffs each player receives in each outcome is well defined. Assuming prisoners are completely selfish, each prefers less own jail time to more.

Based on the above, the described typical “prisoner’s dilemma” game can be represented in the table below

		Player 2	
		Lie (L ₂)	Confess (C ₂)
Player 1	Lie (L ₁)	-1, -1	-9, 0
	Confess (C ₁)	0, -9	-6, -6

Each player has been placed along one dimension (vertically or horizontally) and each has two actions to choose from. In the payoff matrix above, this set up has created four different cells each referring to one of the potential outcomes.

In summary, we define all the elements of the game, that is:

- $\mathcal{N} = \{1, 2\}$
- $A_1 = A_2 = \{C_i, L_i\}$ where $i \in \mathcal{N}$
- $u_1(C_1, L_2) = u_2(L_1, C_2) = 0$
- $u_1(L_1, C_2) = u_2(C_1, L_2) = -9$
- $u_1(C_1, C_2) = u_2(C_1, C_2) = -6$
- $u_1(L_1, L_2) = u_2(L_1, L_2) = -1$

Example 2: Consider a duopoly where firms produce an identical product at zero per unit cost and simultaneously choose quantities (*i.e.*, à la Cournot competition). Let the inverse demand be $p = 1 - q_1 - q_2$.

Clearly, in this game the players are the two firms competing in the market. Since the control variable of a firm in this game is its quantity, the strategy set includes all possible quantities that can be produced⁷. This implies that the strategy set of each firm is infinite, hence the game is not a finite game⁸. Finally, regarding payoffs, we can assume here that each firm cares exclusively about its profits. In the absence of any costs, as in this example, profits will equal the revenue, *i.e.*, the product of price by the quantity. We can define all parts of the game, that is:

- $\mathcal{N} = \{1, 2\}$
- $A_1 = A_2 = [0, \infty)$ or $A_i = [0, \bar{q}_i]$, $i = 1, 2$ with $\bar{q}_i \in \mathcal{R}_+$
- $\Pi_i(q_i, q_{-i}) = (1 - q_i - q_{-i})q_i$

Dominant and Dominated Strategies

In some situations of strategic interdependencies, it is possible that the best interest of a player is to choose the same action/strategy. We call such an action/strategy a **dominant strategy**. In layman terms, a dominant strategy is a strategy that provides a player with the highest payoff for any given set of strategies chosen by the other players. In other words, if a player has a dominant strategy, that player will always choose that strategy no matter what the other players do. We

⁷ One can think that the upper limit of production is infinity or that there is a capacity constraint.

⁸ In general, non-finite strategic form games are not analyzed using payoff matrices. The analysis ought to be more abstract.

distinguish between **strictly dominant** and **weakly dominant** actions/strategies. Letting subscript i refer to player i , and $-i$ refer to all other players except i , we formally define them as

Definitions:

- A strategy \hat{s}_i for player i is *strictly dominant* if $u_i(\hat{s}_i, s_{-i}) > u_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}$ and for all $s_i \neq \hat{s}_i \in S_i$
- A strategy \hat{s}_i for player i is *weakly dominant* if $u_i(\hat{s}_i, s_{-i}) \geq u_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}$ with at least one case where $u_i(\hat{s}_i, s'_{-i}) > u_i(s_i, s'_{-i})$ for $s'_{-i} \in S_{-i}$ and for all $s_i \neq \hat{s}_i \in S_i$

The former implies that no matter the choices of all other players, player i will always get strictly more anytime she plays the (strictly) dominant strategy compared to any other strategy available to her. The latter implies that sometimes choosing the (weakly) dominant strategy, player i will be better off and sometimes as well as when choosing some other strategy (but never worse off).

Similarly, in some situations of strategic interdependencies, it is possible that the best interest of a player is never to choose a specific action/strategy. We call such an action/strategy a **strictly dominated strategy**. A strictly dominated strategy is a strategy that provides a player with a strictly lower payoff than some other strategy, regardless of the strategies chosen by the other players. In other words, if a player has a strictly dominated strategy, that player should never choose that strategy. It is worth to be noted that a strategy, say \hat{s}_i , being strictly dominated implies that there is another strategy, say \bar{s}_i , that dominates it⁹! We distinguish between **strictly dominated** and **weakly dominated** actions/strategies. We formally define them as

Definitions:

- A strategy $\hat{s}_i \in S_i$ for player i is a *strictly dominated* strategy if $u_i(\hat{s}_i, s_{-i}) < u_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}$ and for some $s_i \neq \hat{s}_i \in S_i$
- A strategy $\hat{s}_i \in S_i$ for player i is a *weakly dominated* strategy if $u_i(\hat{s}_i, s_{-i}) \leq u_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}$ with at least one case where $u_i(\hat{s}_i, s'_{-i}) < u_i(s_i, s'_{-i})$ for $s'_{-i} \in S_{-i}$ and for some $s_i \neq \hat{s}_i \in S_i$

The latter implies that most of the times choosing the weakly dominated strategy player i will be getting strictly less than under some other specific strategy. However, sometimes by choosing the weakly dominated strategy, player i will be as well as when choosing that other strategy. It is worth to be noted that a strategy, say \hat{s}_i , being strictly (weakly) dominated implies that there is another strategy, say \bar{s}_i , that strictly (weakly) dominates it¹⁰!

⁹ Note that a strategy might be dominated by either a pure or a mixed strategy (we will discuss mixed strategies later).

¹⁰ In case we consider only pure strategies, a strategy is strictly dominated if it is never chosen by a player. The requirement of that strategy being dominated by another strategy is redundant.

Example 3: Consider the 2-player game described by the payoff matrix below.

		Player 2		
		Left	Center	Right
Player 1	Up	0, 4	2, 2	1, 3
	Down	4, 2	5, 1	0, 0

First, let us think of the strategies of player 1. For any possible action of player 2, namely “Left”, “Center”, or “Right”, we will be checking player 1’s payoff for each of the two actions available to her, namely “Up” and “Down”. Observe below that, if

		Player 2	
		Left	Center
Player 1	Up	0, 4	
	Down	4, 2	

player 2 plays “Left”, player 1 will get a payoff of 0 by choosing “Up” or a payoff of 4 by choosing “Down”. To put it more formally,

$$u_1(\text{Down, Left}) = 4 > u_1(\text{Up, Left}) = 0$$

Obviously, this is not enough to characterize player 1’s strategy “Down” as strictly dominant, or “Up” as strictly dominated. “Down” must yield more compared to “Up”

		Player 2	
		Center	Right
Player 1	Up	2, 2	
	Down	5, 1	

for player 1 no matter what player 2 is choosing. Observe then, that if player 2 plays “Center”, player 1 will get a payoff of 2 by choosing “Up” or a payoff of 5 by choosing “Down”. To put it more formally,

$$u_1(\text{Down, Center}) = 5 > u_1(\text{Up, Center}) = 2$$

Finally, note that if player 2 plays “Right”, player 1 will get a payoff of 1 by choosing

		Player 2	
		Right	Center
Player 1	Up	1, 3	
	Down	0, 0	

“Up” or a payoff of 0 by choosing “Down”. To put it more formally,

$$u_1(\text{Down, Right}) = 0 < u_1(\text{Up, Right}) = 1$$

Our analysis shows that sometimes it is better for player 1 to choose “Up” and sometimes it is better to choose “Down” **depending on** what player 2 is assumed to be doing. Hence, there aren’t any strictly (or weakly) dominant or dominated strategies for player 1.

Now, let us think of the strategies of player 2. For any possible action of player 1, namely “Up” and “Down”, we will be checking player 2’s payoff for each of the two

		Player 2		
		Left	Center	Right
Player 1	Up	0, 4	2, 2	1, 3
	Down			

actions available to him, namely “Left”, “Center”, or “Right”. Observe below that, if player 1 plays “Up”, player 2 will get a payoff of 4 by choosing “Left”, a payoff of 2 by choosing “Center” or a payoff of 3 by choosing “Right”. To put it more formally,

$$u_2(\text{Left}, \text{Up}) = 4 > u_2(\text{Right}, \text{Up}) = 3 > u_2(\text{Center}, \text{Up}) = 2$$

Obviously, this is not enough to characterize player 2’s strategy “Left” as strictly dominant, or “Center” as strictly dominated. For “Left” to be strictly dominant it must yield more compared to both “Center” and “Right” for player 2 no matter what player 1 is choosing. For “Center” to be strictly dominated it must yield strictly less compared to either “Left” and/or “Right” (or any combination of the two) for player 2 no matter what player 1 is choosing. Observe then, that if player 1 plays “Up”, player 2 will get a payoff of 2 by choosing “Left”, a payoff of 1 by choosing “Center” or a payoff of 3 by

		Player 2		
		Left	Center	Right
Player 1	Up	4, 2	5, 1	0, 0
	Down			

choosing “Right”. To put it more formally,

$$u_2(\text{Left}, \text{Down}) = 2 > u_2(\text{Right}, \text{Down}) = 1 > u_2(\text{Center}, \text{Down}) = 0$$

Our analysis shows that it is always better for player 2 to choose “Left” no matter what player 1 is assumed to be doing. Hence, player 2’s strategy “Left” is strictly dominant and, consequently, strategies “Center” and “Right” strictly dominated.

In general, the discussion above clearly shows that analyzing strategies in a game entails a complex process that necessitates considering the strategies of all players and their potential responses to various strategies. More importantly though it should be intuitively clear that the outcome of a situation that involves strategic interactions can never include choices (*i.e.*, strategies) of the agents involved that are strictly dominated, hence all strictly dominated strategies can be eliminated. Any prediction about the final outcome of a game should not include any player choosing some strictly dominated action¹¹. Similarly, if a player has a strictly dominant strategy, then any prediction about the final outcome of the game should definitely include strictly dominant action of the respective player¹². In the context of *Example 3*, this implies that our predictions about the final outcome of the game should include the prediction that player 2 will definitely choose “Left” (and, consequently, never choose “Center” or “Right”).

How important is this? Consider again the game of *Example 3*. Knowing that player 2 will definitely choose “Left” simplifies our analysis by reducing the payoff

		Player 2	
		Left	
Player 1	Up	0, 4	
	Down	4, 2	

¹¹ However, a solution may very well include weakly dominated strategies.

¹² However, it is possible a solution not to include a weakly dominant strategy.

matrix from a 2x3 to a 2x1 table! So, instead of having to analyze six possible outcomes we only must consider two. As a matter of fact, observe now that making a prediction about the final outcome in the above reduced form table is fairly simple: It is only player 1 who has to make a choice and doing so requires just simple optimization!

A natural question can be: Can we do the same if a strategy is weakly dominant? The answer is unfortunately no. In the context of Example 3 assume a small change in the payoffs. Specifically, let the payoff of player 2 under the strategy profile (“Down”, “Center”) to be 2 instead of 1. Then it is easy to show that player 2’s strategy “Left” is weakly dominant. But the strategy “Center” is not strictly dominated, as it sometimes yields the same payoff as the strategy “Left”. Therefore, we can only

		Player 2		
		Left	Center	
Player 1	Up	0, 4	2, 2	
	Down	4, 2	5, 1	

eliminate “Right” as strictly dominated and the payoff matrix reduces to the one above with four remaining possible outcomes instead of six.

What if every player in a game has a strictly dominant strategy? Then, as we have claimed, the “solution”¹³ of the game should predict every player choosing her strictly dominant strategy. For example, considering the game of the Prisoner’s Dilemma, it is easy to confirm that each prisoner has a strictly dominant strategy to “Confess”. Denoting with C_i and L_i player i ’s strategies “Confess” and “Lie”, respectively, and using the familiar $-i$ for the strategy of the “other” player, we have

$$u_i(C_i, L_{-i}) = 0 > u_i(L_i, L_{-i}) = -1$$

$$u_i(C_i, C_{-i}) = -6 > u_i(L_i, C_{-i}) = -9$$

In other words, the strategy “Lie” of both players is strictly dominated and can be eliminated. By only keeping the strictly dominant actions of the two players in the

		Player 2	
		Confess	
Player 1	Confess		-6, -6

Prisoner’s Dilemma game, the payoff matrix reduces to the one above and there is only one potential outcome: we have found the equilibrium in the Prisoner’s Dilemma game! In equilibrium both prisoners will choose to “Confess”, and each will be sentenced for six years. Is this enough to solve any game? Unfortunately, no. As we have already discussed, in most of the games there is no strictly dominant strategy for every player. Is there anything more to say about strict dominance? According to what follows there is.

¹³ Henceforth, we will be using the term “equilibrium” of a game.

Iterated Elimination of Strictly Dominated Strategies

We have established that the existence of strictly dominant strategies for some players and/or the existence of strictly dominated strategies can significantly reduce the “size” of the game. Is there something more we can do after eliminating any strictly dominated strategy in a game? The answer is yes! Once we eliminate such strategies, say for player 1, we should go and check again about dominances on player 2’s strategies in the reduced payoff matrix. This process can reveal that, although previously player 2 did not have any strictly dominant or dominated actions, after the elimination of some strategies of player 1 and the derivation of the reduced payoff matrix, player 2 might have strictly dominated actions. If yes, strictly dominated actions in the reduced game must be eliminated and the process continues by checking again player 1, and so on so forth. This process is called **Iterated Elimination of Strictly Dominated Strategies** and simplifies strategic interactions between rational players. The idea is to eliminate from consideration all strategies that are strictly dominated by another strategy, and then iterate the process until only a single strategy or a set of strategies that are not dominated remains.

Example 4:

Consider the game described by the payoff matrix below. Following the analysis used in *Example 3* for identifying dominances, you can easily confirm that player 1 does not

		Player 2		
		Left	Middle	Right
Player 1	Up	1, 0	1, 2	0, 1
	Down	0, 3	0, 1	2, 0

have a strictly dominant and/or strictly dominated actions, since her strategy “Up” yields a higher payoff for her when

player 2 choose “Left” or “Middle” but her strategy “Down” is better for her when player 2 chooses “Right”. Moving on to player 2, we can again easily confirm that his strategy “Middle” yields always a higher payoff to him compared to his strategy “Right”, hence the latter is a strictly dominated strategy and it can be eliminated (1st

round of elimination of strictly dominated strategies). The reduced game can now be represented by the payoff

		Player 2		
		Left	Middle	
Player 1	Up	1, 0	1, 2	
	Down	0, 3	0, 1	

matrix on the right. In this reduced form game note now that by re-examining player 1’s strategies for dominances shows that her strategy “Up” is strictly dominating over

		Player 2		
		Left	Middle	
Player 1	Up	1, 0	1, 2	
	Down			

her strategy “Down”, hence the latter can be eliminated (2nd round of elimination). The new reduced game can now

be represented by the payoff matrix on the left. In this reduced form game note now that by re-examining player 2's strategies for dominances shows that his strategy "Middle" strictly dominates over his strategy "Left", hence the latter can be eliminated (3rd round of elimination). This leaves us with a clear prediction about the final outcome of the game as it is shown in the new reduced payoff matrix below. Specifically, we can predict with a very high probability¹⁴ that player 1 will choose "Up" and player 2 will choose "Middle".

		Player 2	
		Middle	Left
Player 1	Up	1, 2	1, 2
	Down	1, 2	1, 2

It appears, as shown in *Example 4*, that this process of iterated elimination of strictly dominated strategies can lead to identifying the equilibrium of the game. However, we should be warned about a few things.

- First, this process can only be used for strict dominances, implying that we should not ever eliminate weakly dominated strategies. The reason we should avoid the elimination of weakly dominated strategies is because there is no guarantee that a weakly dominated strategy will not be part of the equilibrium strategy profile. This is true only for strictly dominated strategies.
- Second, the iteration of strictly dominated strategies can stop at any round if a player does not have a strictly dominated strategy in a reduced form game. Hence, although the process can help to reduce the complexity of the analysis required to reach a prediction about the final outcome (*i.e.*, an equilibrium) by reducing the number of potential outcomes, it does not need to lead to the equilibrium of the game.
- Third, and most importantly, this process requires additional assumptions about the behavior of the players. We have assumed that the players are rational implying that, as individuals, are trying to maximize their wellbeing. In the process of iterative elimination of strictly dominated strategies moving from one round to the next requires not only that players are rational but that their rivals know that they are rational. In *Example 4*, moving from the first round to the second, this assumption implies that player 2 is rational (hence he eliminates "Right") and player 1 knows that player 2 is rational so that player 1 is now facing a reduced form payoff matrix. This might not sound like anything special but there is more! Player 2 knows that player 1 knows that player 2 is rational, hence in round three player 2 eliminates "Left". If this "chain" of knowledge of rationality breaks at any point before the last (where player 2 knows that player 1 knows that player 2 is rational) the iterative elimination of strictly dominated strategies will stop and there will be no way to identify the equilibrium/solution of the game. In Game

¹⁴ We are almost sure that this will be the final outcome. However, we cannot exclude even a very tiny little probability that one of the players makes a mistake!

Theory we overcome this situation of requiring constant reassurance about the players' rationality by assuming that players possess "**common knowledge rationality**".

Nash Equilibrium

What should be an admissible characteristic of an equilibrium? We can claim that an equilibrium situation is characterized by inertia, that is a tendency to do nothing or to remain unchanged. In other words, a possible outcome resulting from a strategy profile combining one strategy for every player, can be an equilibrium of a game if there is not a single player wishing to change her strategy for the given strategies of the other players. Alternatively, a possible outcome cannot be an equilibrium of a game if there is at least one player wishing to change her strategy given the strategies of the others. This is the idea behind the Nash Equilibrium (NE) concept: given the choices of all other players (no matter if these choices are "best," "rational," etc.) should I change my choice? No, I shouldn't if I have chosen the action that yields what is best for me (i.e., best response) for the given choices of others. Formally, player i 's best response is defined for every possible combination of the strategies of all other players, denoted as $-i$, as

Definition: A best response function (or correspondence) is defined as

$$B_i(a_{-i}) = \{a_i \in A_i \mid (a_i, a_{-i}) \succeq_i (a_i', a_{-i}), \forall a_i' \in A_i\}$$

According to the definition above, a best response is a "rule" (i.e., a function or correspondence) that assigns an action that yields the best result to player i for some given set of actions of all the other players. Consider again the "Prisoner's dilemma" game. What is the best response correspondence of player 1 for example? Since "other players" from the perspective of player 1 is only player 2 we will identify the best response of player 1 for any given choice of player 2. If player 2 plays "Lie" then, according to the payoff matrix below (where "Confess" of player 2 is hidden) player 1

		Player 2	
		L ₂	C ₂
Player 1	L ₁	-1, -1	
	C ₁	0, -9	

can either get -1 (by choosing "Lie") or 0 (by choosing "Confess"). Therefore, when player 2 plays "Lie", player 1's best response is "Confess" since 0 is greater than -1. Formally,

$$U_1(C_1, L_2) = 0 > U_1(L_1, L_2) = -1 \Leftrightarrow (C_1, L_2) \succeq_1 (L_1, L_2)$$

hence

$$B_1(L_2) = C_1$$

Working similarly, if player 2 plays “Confess” then, according to the payoff matrix below (where “Lie” of player 2 is hidden) player 1 can either get -9 (by choosing “Lie”)

		Player 2	
		C ₂	L ₂
Player 1	L ₁		
	C ₁		
		-9, 0	-6, -6

or -6 (by choosing “Confess”). Therefore, when player 2 plays “Confess”, player 1’s best response is “Confess” since -6 is greater than -9. Formally,

$$U_1(C_1, C_2) = -6 > U_1(L_1, C_2) = -9 \Leftrightarrow (C_1, C_2) \succcurlyeq_1 (L_1, C_2)$$

hence

$$B_1(C_2) = C_1$$

Note that, in general, this level of formality is not required when discussing best responses. In strategic form games represented by a payoff matrix there is a fast way of identifying best responses.

Following the discussion above, if every player is choosing a best response to the given strategies of the others, there will be no unilateral incentive to deviate, *i.e.*, no one will want to change her strategy. It is therefore straightforward to define Nash Equilibrium in terms of best responses as

Definition: A Nash Equilibrium of a strategic (normal) form game is a strategy profile $a^* \in A = \times_{i \in \mathcal{N}} A_i$ such that $(a_i^*, a_{-i}^*) \succcurlyeq_i (a_i, a_{-i}^*)$ for all $a_i \in A_i$ and for all $i \in \mathcal{N}$. Alternatively, A Nash Equilibrium is a profile $a^* \in A$ such that $a_i^* \in B_i(a_{-i}^*) \forall i \in \mathcal{N}$

According to this definition, a Nash equilibrium contains one strategy for every player such that every strategy is a Best Response. Continuing the example with the “Prisoner’s Dilemma” game, we showed that player 1’s best response can be summarized as

- if player 2 plays “Lie”, player 1’s best response is “Confess”, *i.e.*, $B_1(L_2) = C_1$
- if player 2 plays “Confess”, player 1’s best response is “Confess”, *i.e.*, $B_1(C_2) = C_1$

Due to the symmetry of the game, it is straightforward to confirm that player 2’s best response function can be summarized as

- if player 1 plays “Lie”, player 2’s best response is “Confess”, *i.e.*, $B_2(L_1) = C_2$
- if player 1 plays “Confess”, player 2’s best response is “Confess”, *i.e.*, $B_2(C_1) = C_2$

Combining the best responses from above, there is only one set of strategies where each strategy is a best response to the other. Specifically, we have

$$\begin{array}{c}
 B_1(C_2) = C_1 \\
 \swarrow \quad \searrow \\
 B_2(C_1) = C_2
 \end{array}$$

that is, both players choose “Confess”. Note, that when player 2 chooses “Confess” is best for player 1 to choose “Confess” and, moreover, when player 1 chooses “Confess” is best for player 2 to choose “Confess”. Therefore, the strategy profile $\{C_1, C_2\}$ is a

Nash Equilibrium of the Prisoner’s Dilemma game. As a counterexample, consider the combination of strategies below:

$$\begin{aligned}
 B_1(L_2) &= C_1 \\
 B_2(C_1) &= C_2
 \end{aligned}$$

where when player 2 chooses “Lie” is best for player 1 to choose “Confess” BUT when player 1 chooses “Confess” is best for player 2 to choose “Confess” and not “Lie”.

Note again that, in general, this level of formality is not required when discussing Nash Equilibria in this class. We ought to do couple of things though when finding a Nash Equilibrium. First, a suggested Nash Equilibrium requires us to report that many strategies as the number of players. In the Prisoner’s dilemma example above two strategies must be reported. Second, we must explain (in words or using math) why the specific strategy profile is indeed a Nash Equilibrium.¹⁵

An immediate implication of the definition of Nash Equilibrium as a collection of best responses is that no player wishes to do something different! If a player has chosen the best strategy for her (given the choices of the others) then she has no incentive to change her decision. This gives rise to an alternative way of identifying if a strategy profile is a Nash Equilibrium. We must check that specific strategy profile if a player wants to deviate from that state for fixed choices others (*i.e.*, unilateral deviation). If, for that given profile, there is at least one player that wants to change her action the given profile CANNOT be a Nash Equilibrium. If, for that given profile, no player wants to change his action then the given profile is a Nash Equilibrium.

Example 5: Consider the game described by the payoff matrix below. What is (are) the Nash Equilibrium (Equilibria) in pure strategies? Explain your answer.

		Player 2		
		Left	Center	Right
Player 1	Up	1, 4	1, 2	3, 3
	Middle	4, 4	0, 2	1, 2
	Down	5, 0	2, 1	0, 0

Here, we are going to find the NE using the best responses approach. We will not be formal in the way we describe the best responses (no math!) but we will be using underlines in the payoff matrix! First, we consider player 1. We fix the strategy of player 2 to “Left.” Then it is best for player 1 to choose “Down”, since 5 is greater than 4 and 1, hence we underline the payoff of player 1 in the respective cell:

		Player 2	
		Left	Center
Player 1	Up	1, 4	
	Middle	4, 4	
	Down	<u>5</u> , 0	

¹⁵ In some cases, we must also explain why other strategy profiles Nash Equilibria are not.

Then we fix the strategy of player 2 to “Center.” Then it is best for player 1 to choose “Down”, since 2 is greater than 0 and 1, hence underline the payoff of player 1 in the respective cell:

		Player 2		
		Center		
Player 1	Up		1, 2	
	Middle		0, 2	
	Down		<u>2</u> , 1	

Finally, we fix the strategy of player 2 to “Right.” Then it is best for player 1 to choose “Up”, since 3 is greater than 0 and 1, hence underline the payoff of player 1 in the respective cell:

		Player 2	
		Right	
Player 1	Up		<u>3</u> , 3
	Middle		1, 2
	Down		0, 0

Now, we consider player 2. We fix the strategy of player 1 to “Up.” Then it is best for player 2 to choose “Left”, since 4 is greater than 3 and 2, hence underline the payoff of player 2 in the respective cell:

		Player 2		
		Left	Center	Right
Player 1	Up	1, <u>4</u>	1, 2	3, 3
	Down			

Then, we fix the strategy of player 1 to “Middle.” Then it is best for player 2 to choose “Left”, since 3 is greater than 0 and 1, hence underline the payoff of player 1 in the respective cell:

		Player 2		
		Left	Center	Right
Player 1	Middle	4, <u>4</u>	0, 2	1, 2
	Down			

Finally, we fix the strategy of player 1 to “Down.” Then it is best for player 2 to choose “Center”, since 3 is greater than 0 and 1, hence underline the payoff of player 1 in the respective cell:

		Player 2		
		Left	Center	Right
Player 1	Down	5, 0	2, <u>1</u>	0, 0
	Up			

Putting everything together, the payoff matrix with “best responses” as underlines appears on the left. Note that only in the cell (Down, Center) both payoffs are underlined. Hence, the Nash Equilibrium is {“Down”, “Center”}.

		Player 2		
		Left	Center	Right
Player 1	Up	1, <u>4</u>	1, 2	<u>3</u> , 3
	Middle	4, <u>4</u>	0, 2	1, 2
	Down	<u>5</u> , 0	2, <u>1</u>	0, 0

Hence, the Nash Equilibrium is {“Down”, “Center”}.

Note that when describing the equilibrium of the game presented in *Example 5* we do not use the payoffs (*i.e.*, we do not write “the equilibrium is (2,1)”) but we specify the actions/strategies chosen by the players. This should be a rule for you: we describe the solution of the game by identifying what each player will do!

In *Example 5* we used the best responses approach to identify the Nash Equilibrium of the game. In the same example we can also use the “no unilateral deviation” rule to find the Nash Equilibrium. Consider, for example, the encircled cell below that corresponds to the strategy profile {Up, Center}. Does player 1 have an

		Player 2		
		Left	Center	Right
Player 1	Up	1, 4	1, 2	3, 3
	Middle	4, 4	0, 2	1, 2
	Down	5, 0	2, 1	0, 0

incentive to unilaterally change her decision given the choice of player 2? Player 1 can only move along the vertical axis on this table (i.e., it can only change rows).

Therefore, for her to move, a higher payoff than 1 is required. But if player 1 changes her choice to “Down” (and given player 2’s choice “Center”) her payoff will increase to 2. Therefore, player 1 wishes to unilaterally deviate and the strategy profile {Up, Center} is not a Nash Equilibrium. One must follow the same “checking” procedure for unilateral deviation to every one of the nine cells of example 5’s payoff matrix. This is left to the students as an exercise. However, we will confirm, using the no unilateral deviation rule that the strategy profile {Down, Center} is a Nash Equilibrium. Checking the options of player 1 we can confirm that moving away from {Down, Center} and going to either {Middle, Center} or {Up, Center} she will receive a lower payoff (0 or 1 instead of 2 as shown on the table to the right). Therefore, player 1 does

		Player 2		
		Left	Center	Right
Player 1	Up	1, 4	1, 2	3, 3
	Middle	4, 4	0, 2	1, 2
	Down	5, 0	2, 1	0, 0

not want to unilaterally deviate. Similarly, checking the options of player 2 we can confirm that moving away from {Down, Center} and going to either {Down, Left} or

		Player 2		
		Left	Center	Right
Player 1	Up	1, 4	1, 2	3, 3
	Middle	4, 4	0, 2	1, 2
	Down	5, 0	2, 1	0, 0

{Down, Right} he 1 will receive a lower payoff (of 0 in either case instead of 1 as shown on the table to the left). Therefore, player 2 doesn’t want to unilaterally

deviate. Since no player has an incentive to unilaterally deviate, the strategy profile {Down, Center} is a NE.

Which approach is better? There is no “rule” on what approach one should use to identify Nash Equilibria with the lowest possible effort. It all depends on the type of the problem and the experience one has by solving many different games. However, we can consider a few key facts:

- For strategic form games described by a payoff matrix the best responses approach seems to be a little less tiring compared to the “no unilateral deviation” rule when the number of cells in the payoff matrix is large.
- When in a game the preferences of players are appropriately described by differentiable objective functions (e.g., differentiable utility and/or profit functions) it is usually more convenient to use the best responses approach.

Be aware though that this approach, if not used with caution, might not lead to the identification of all Nash Equilibria (we will see that in an example that follows).

- In games with a higher level of abstractness it is more common to use the “no unilateral deviation” rule to identify the Nash Equilibria (we will see a few examples that fall into this category).

As we saw before, strategic form games cannot always be represented by a payoff matrix. Payoff matrices are convenient when the number of players is no greater than 3 (preferably only 2) and the players have a limited number of strategies to choose from. What if, as we have seen in Example 2, the players of the game have infinitely many actions/strategies to choose from?

Example 6: Consider again a duopoly where firms produce an identical product (at zero per unit cost) and “simultaneously” choose quantities (*i.e.*, competition à la Cournot). Let the inverse market demand be $p = 1 - q_1 - q_2$.

Under the usual assumption that firms are profit maximizers, the optimization problem facing a firm $i \in \{1, 2\}$ in this duopoly setup is expressed by

$$\max_{q_i} \{\Pi_i = (1 - q_i - q_{-i})q_i\}$$

with first order condition of

$$\frac{\partial \Pi_i}{\partial q_i} = 0 \Rightarrow 1 - 2q_i - q_{-i} = 0 \Rightarrow q_i = 0.5 - 0.5q_{-i}$$

The above expression defines the best response function of firm i . Solving the reaction functions of the two firms as a system yields the strategy profile that is a Nash Equilibrium:

$$q_1^* = q_2^* = 1/3$$

Given the optimal quantities above one can easily confirm that the equilibrium price is $p^* = 1/3$ and equilibrium profits are $\Pi_i^* = 1/9$.

Previously, we have found a distinct Nash Equilibrium in each example that we discussed. Does this hold true for all cases? Absolutely not, because there are certain games that do not have any Nash equilibrium at all, and there are also others that have multiple Nash equilibria. Determining whether a game possesses a Nash equilibrium is of great significance in many cases. There are several theorems about the existence of at least one Nash Equilibrium in different categories of strategic form games¹⁶ and, in general, existence did not seem to be an insurmountable obstacle. On the other hand, multiplicity of Nash Equilibria is a far more serious issue. If a particular strategy profile is a Nash Equilibrium, it means that no player has an incentive to unilaterally change their strategy. What Nash Equilibrium doesn't explain is how or why players

¹⁶ Perhaps the most prominent existence theorem establishes that there is at least one Nash equilibrium in any strategic game with a finite number of players, each having a finite set of actions.

select their strategies. Of course, if there is only one Nash Equilibrium, we can rely on logical analysis and assume that common knowledge rationality suffices to explain how players choose their strategies. In cases, however, where there are several Nash Equilibria, relying solely on common knowledge rationality may not be enough to explain why one Nash Equilibrium prevails over another. Consequently, predictions about the outcome of the game may be entirely incorrect. Interestingly, in some situations with multiple Nash Equilibria, the outcome of the game could even be one that isn't a Nash Equilibrium (for example, the "Battle of the Sexes" game, as explained later). In what immediately follows, *Example 7* is a game without a Nash Equilibrium and *Example 8* shows a game with two Nash Equilibria.

Example 7 (Matching the Pennies) Consider the game described by the payoff matrix below. It is straightforward to confirm, as shown on the payoff matrix with the

		Player 2	
		Heads	Tails
Player 1	Heads	<u>1</u> , -1	-1, <u>1</u>
	Tails	-1, <u>1</u>	<u>1</u> , -1

underlying identifying best responses, that this game has no Nash Equilibrium in pure strategies. This seems to contradict a fundamental theorem in Game Theory (see footnote 13). However, as we will see later, this game has a Nash Equilibrium in mixed strategies.

Example 8 (Battle of the Sexes) Consider the game described by the payoff matrix below. Just as an underlying story, accept that the players are Mary (player 1) and Peter (player 2), a couple that is crazy in love. In a specific evening, and while they have no

		Boxing Match	Opera
		Mary	Boxing Match
Opera	0, 0		<u>1</u> , <u>2</u>

way of communicating with each other, two major events are taking place in their city. Each knows that the other will be in one of the two events (but does not know which). They are not the same characters (opposites attract!) so Mary, a dynamic and athletic woman, prefers to attend a boxing match while Peter, a sensitive artist, prefer to listen to the Opera. However, they both prefer being together at any event to being alone. It is straightforward to confirm, as shown on the payoff matrix with the underlying identifying best responses, that this game has two Nash Equilibria in pure strategies, namely one where they both go the Boxing Match (*i.e.*, $NE = \{\text{Boxing, Boxing}\}$) and the other where both go to the Opera (*i.e.*, $NE = \{\text{Opera, Opera}\}$). The uncertainty in this scenario arises from the fact that there's no basis for choosing one Nash Equilibrium over the other. In other words, both Nash Equilibria are equally probable. For example, if Peter is more adaptable (and Mary is aware of this, and Peter is aware that Mary is aware of this, etc.), then it's more likely that Peter will give in and go to

the boxing match, where Mary is eagerly waiting for him. Conversely, if Peter is unresponsive (and Mary is aware of this, and Peter is aware that Mary is aware of this, etc.), then it's more likely that Mary will give in and go to the boxing match, where Peter is waiting for her! This situation may not seem concerning since, ultimately, the outcome is a Nash Equilibrium. The critical factor is determining who is more likely to compromise. One objection to this situation is that this information (regarding who is more likely to compromise) should be included in the payoffs. However, even if we accept the payoffs as they are, the issue persists. What if both individuals make incorrect assumptions about their partner's emotional state? What if they both compromise because they want to make the other happier? In such a scenario, Peter will attend the boxing match, and Mary will attend the opera (and both will feel very unhappy). This is a typical example of coordination failure (that we will discuss later).

There are two more factors to consider when it comes to Nash Equilibrium. Firstly, Nash Equilibria may be challenging to identify (sometimes they hide in ... the corners; see *Example 9* below). Secondly, Nash Equilibrium is a theoretical concept, and its ability to predict outcomes must be tested when possible. Researchers have conducted experiments demonstrating that in some cases, the theoretical predictions of Nash Equilibrium do not align with the actual outcomes (see *Example 10* below).

Example 9 Consider a Cournot triopoly where firms have zero marginal and zero fixed costs, and the inverse demand is described by

$$p = \begin{cases} A - q_1 - q_2 - q_3, & \text{if } \sum_{i=1}^3 q_i < A \\ 0, & \text{otherwise} \end{cases}$$

Under the usual assumption that firms are profit maximizers, the optimization problem facing a firm $i \in \{1, 2, 3\}$ in this triopoly setup is expressed by

$$\max_{q_i} \left\{ \Pi_i = (1 - q_i - \sum_{j \neq i}^3 q_j) q_i \right\}$$

with first order condition of

$$\frac{\partial \Pi_i}{\partial q_i} = 0 \Rightarrow 1 - 2q_i - \sum_{j \neq i}^3 q_j = 0 \Rightarrow q_i = 0.5 - 0.5 \sum_{j \neq i}^3 q_j$$

The above expression defines the best response function of firm i . Solving the reaction functions of the three firms as a system yields the strategy profile that is a Nash Equilibrium $\{q_1^*, q_2^*, q_3^*\} = \{1/4, 1/4, 1/4\}$. Given the optimal quantities above one can easily confirm that the equilibrium price is $p^* = 1/4$ and equilibrium profits are $\Pi_i^* = 1/16$. However, note that this is not a unique Nash Equilibrium. In fact, there are infinitely many Nash Equilibria in this game. A more careful reader will notice that, since the price zeros when total quantity in the market is no less than A , the reaction function of firm i is actually given by

$$\begin{cases} q_i = 0.5 - 0.5 \sum_{j \neq i}^3 q_j, & \text{if } \sum_{k=1}^3 q_k < A \\ q_i \in [0, \infty), & \text{otherwise} \end{cases}$$

It is then straightforward to verify that any strategy combination $\{q_1^{**}, q_2^{**}, q_3^{**}\}$ where the sum of any two strategies is at least A is also a Nash Equilibrium.¹⁷

Example 10 Consider the following scenario describing the “Traveler’s Dilemma” game, as it was formulated by K. Basu (1994): An airline loses two identical suitcases belonging to two different travelers. Both suitcases contain identical antiques. The airline manager responsible for settling the claims of both travelers explains that the airline is only liable for a maximum of \$100 per suitcase and is unable to determine the exact value of the antiques. To determine a fair appraised value of the antiques, the manager separates the travelers so that they cannot communicate with each other. He then instructs both travelers to write down the value of their antiques, which must be between \$2 and \$100. If both travelers write down the same value, the manager will consider that value as the true dollar value of both suitcases and reimburse both travelers that amount. However, if one traveler writes down a lower value than the other, the lower value will be considered as the true dollar value, and both travelers will receive that amount along with a bonus/malus: the traveler who wrote down the lower value will receive an additional \$2, while the traveler who wrote down the higher amount will be charged a \$2 deduction. What strategy should both travelers adopt to decide the value they should write down?

It is easy to confirm that the game described above can be represented by the payoff matrix below and that this game has a unique Nash Equilibrium, namely one

where both players claim just \$2 each (highlighted in blue the best response of player 1 and highlighted in yellow the best responses of player 2)! However, in experiments

		#2						
		100	99	98	97	...	3	2
#1	100	100, 100	97, 101	96, 100	95, 99	...	1, 5	0, 4
	99	101, 97	99, 99	96, 100	95, 99	...	1, 5	0, 4
	98	100, 96	100, 96	98, 98	95, 99	...	1, 5	0, 4
	97	99, 95	99, 95	99, 95	97, 97	...	1, 5	0, 4
	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
	3	5, 1	5, 1	5, 1	5, 1	...	3, 3	0, 4
	2	4, 0	4, 0	4, 0	4, 0	...	4, 0	2, 2

conducted by various researchers, individuals were asked to play the “Traveler’s Dilemma” game. When they were facing a large bonus/malus parameter (*i.e.*, the

¹⁷ A reader who pays close attention will verify that in all these situations, the profit for each company is equal to zero. On the other hand, in the interior Nash Equilibrium where $q_1^* = q_2^* = q_3^* = 1/4$ the profits of each equal $1/16$. It is possible to assert that the players’ rational behavior will lead them to select the only viable solution with a positive profit.

“penalty” if you claim higher amount than the other player was large, say \$20) Nash Equilibrium was a relatively good predictor of the individuals’ behavior. When, however, they were facing a small bonus/malus parameter (*i.e.*, the “penalty” if you claim higher amount than the other player was small, say \$2 like in our example) most players tend to choose a value that is higher than the Nash equilibrium and closer to \$100, thus making Nash Equilibrium a bad predictor of the individuals’ behavior.

Coordination Failure

Consider a scenario reminiscent of two drivers navigating toward each other on a narrow roadway. Each driver faces the critical decision of swerving either left or right to avoid a potential collision. If both drivers opt to swerve in the same direction (*e.g.*, each driver swerves to her right) a collision can be averted, resulting in a preferable outcome for both parties. However, should they swerve in opposing directions, a collision becomes inevitable, leading to adverse consequences for both drivers. In this scenario, there exist two potential Nash equilibria: one where both drivers swerve left, and another where both swerve right. However, in the absence of effective communication or coordination between the drivers, the risk of swerving in conflicting directions looms, ultimately culminating in a collision and an unfavorable outcome for all involved parties.

The scenario described above exemplifies a case of *coordination failure*, where players in a game encounter challenges in aligning their actions effectively, thus resulting in a suboptimal outcome for all participants. To discuss the concept of coordination failure we first describe Pareto efficiency (or Pareto optimality) as a situation (*e.g.*, an allocation of resources) where the circumstances of an individual cannot be improved by moving to a different situation without making at least another individual worse off.

So, what is Coordination Failure then? We define coordination failure in terms of Pareto efficiency. Specifically, a coordination failure is a situation where the outcome of the interaction is not Pareto efficient. This happens when the objectives of the players are not aligned. In ascending order of objectives alignment, we have the following strategic situations:

1. “Prisoner’s dilemma” type of games

The unique Nash equilibrium does not coincide with the Pareto Efficient outcome. As a result, non-cooperative players will never reach Pareto Efficiency.

Example 11: Coordination failure in Prisoner’s Dilemma

Consider again the typical “prisoner’s dilemma” presented in the table below for which we know that the unique Nash Equilibrium strategy profile is {Confess, Confess}. What can one tell about the outcome of the strategy profile {Lie, Lie} and

		Player 2	
		Lie	Confess
Player 1	Lie	-1, -1	-9, 0
	Confess	0, -9	-6, -6

the possibility of moving to a different strategy profile? Specifically, if we move from the strategy profile $\{\text{Lie}, \text{Lie}\}$, where both prisoners receive a payoff of -1, to the strategy profile $\{\text{Confess}, \text{Lie}\}$, we improve the wellbeing of prisoner 1 (whose payoff becomes 0 instead of -1) but we hurt prisoner 2 (whose payoff becomes -9 instead of -1). Similarly, if we move from the strategy profile $\{\text{Lie}, \text{Lie}\}$ to the strategy profile $\{\text{Lie}, \text{Confess}\}$, we improve the wellbeing of prisoner 2 (whose payoff becomes 0 instead of -1) but we hurt prisoner 1 (whose payoff becomes -9 instead of -1). Finally, if we move from the strategy profile $\{\text{Lie}, \text{Lie}\}$ to the strategy profile $\{\text{Confess}, \text{Confess}\}$ we hurt both prisoners (whose payoffs become -6 instead of -1). In other words, examining whether there exists an alternative to the $\{\text{Lie}, \text{Lie}\}$ strategy profile whose outcome clearly improves one prisoner's wellbeing while avoiding harm to the other, reveals that no such alternative exists. This implies that the strategy profile $\{\text{Lie}, \text{Lie}\}$ is Pareto efficient. Working similarly and comparing any specific strategy profile against all alternative strategy profiles, we can confirm that the strategy profiles $\{\text{Confess}, \text{Lie}\}$ and $\{\text{Lie}, \text{Confess}\}$ also are Pareto efficient.

On the other hand, the strategy profile $\{\text{Confess}, \text{Confess}\}$, where each prisoner receives a payoff of -6, is not Pareto efficient since if we move to the strategy profile $\{\text{Lie}, \text{Lie}\}$, where each prisoner receives a payoff of -2, we improve the wellbeing of both prisoners. In other words, examining whether there exists an alternative to the $\{\text{Confess}, \text{Confess}\}$ strategy profile whose outcome clearly improves one prisoner's wellbeing while avoiding harm to the other, reveals that such alternative exists: it is the strategy profile $\{\text{Lie}, \text{Lie}\}$ is a Pareto improvement to the strategy profile $\{\text{Confess}, \text{Confess}\}$, hence the latter is not Pareto efficient.

What is the outcome of this game? As we know the Nash Equilibrium of the prisoner's dilemma game is unique. This allows us to assume with confidence that the outcome of the game will be the NE strategy profile $\{\text{Confess}, \text{Confess}\}$. However, this outcome is not Pareto efficient, hence there is the problem of coordination failure in the prisoner's dilemma game.

Note that the prisoner's dilemma game has three strategy profiles that are Pareto efficient but, in general, they are not interpersonally comparable. Prisoner 1's most and least preferred outcomes are the ones resulting from the strategy profiles $\{\text{Confess}, \text{Lie}\}$ and $\{\text{Lie}, \text{Confess}\}$, correspondingly. On the other hand, the order of preferences for prisoner 2 is the reverse: most and least preferred outcomes are the ones resulting from the strategy profiles $\{\text{Lie}, \text{Confess}\}$ and $\{\text{Confess}, \text{Lie}\}$, correspondingly. Can we find a means to evaluate and contrast the three Pareto efficient outcomes? One common approach is to examine the payoffs or well-being levels associated with each outcome. By comparing the numerical values of these payoffs, we can determine which outcome provides a higher level of overall wellbeing or utility. Additionally, we can analyze the distribution of payoffs between the individuals involved to assess the fairness or equity of each outcome. Other factors, such as the preferences or priorities of the individuals, could also be taken into

consideration when comparing the three outcomes. Ultimately, the specific method of comparison may depend on the context and the criteria used to evaluate wellbeing or utility. We are not going to examine these approaches here.

Most importantly, in prisoner's dilemma type of games, coordination failure is inevitable as individual interests clash with the Pareto Efficient allocation, resulting in a collectively unfavorable outcome despite rational behavior on an individual level. In such scenarios, coordination can be attained by implementing mechanisms that enforce cooperation and discourage players from straying from the Pareto Efficient outcome. This concept echoes the strategies employed by Cartels. Such mechanisms may involve repetitive gameplay or adjustments to payoffs to transform the Pareto Efficient outcome into a Nash Equilibrium, albeit altering the game's fundamental nature away from a prisoner's dilemma.

2. "Battle of sexes" type of games

There are multiple Nash Equilibria, and they are all Pareto efficient. However, none of them Pareto dominates the others: different players prefer different Nash equilibria. As a result, non-cooperative players might choose actions that constitute a strategy profile that it is not a Nash equilibrium (so, not Pareto efficient either).

Example 12: Coordination failure in Battle of sexes

Consider again the typical "battle of sexes" game presented in the table below for which we know that it has two Nash Equilibria in pure strategies, namely one where they both go the Boxing Match (*i.e.*, $NE = \{\text{Boxing, Boxing}\}$) and the other where both go to the Opera (*i.e.*, $NE = \{\text{Opera, Opera}\}$). What can one tell about the outcome of

		Player 2	
		Boxing Match	Opera
Player 1	Boxing Match	2, 1	0, 0
	Opera	0, 0	1, 2

the strategy profile $\{\text{Opera, Opera}\}$ and the possibility of moving to a different strategy profile? If we move from the strategy profile $\{\text{Opera, Opera}\}$, where players 1 and 2 receive a payoff of 1 and 2, respectively, to the strategy profile $\{\text{Boxing, Boxing}\}$, we improve the wellbeing of player 1 (whose payoff becomes 2 instead of 1) but we hurt player 2 (whose payoff becomes 1 instead of 2). Furthermore, if we move from the strategy profile $\{\text{Opera, Opera}\}$ to the strategy profile $\{\text{Opera, Boxing}\}$ or the $\{\text{Boxing, Opera}\}$, both players become worse off by receiving a payoff of zero. In other words, examining whether there exists an alternative to the $\{\text{Opera, Opera}\}$ strategy profile whose outcome clearly improves one prisoner's wellbeing while avoiding harm to the other, reveals that no such alternative exists. This implies that the strategy profile $\{\text{Opera, Opera}\}$ is Pareto efficient. Working similarly and comparing any specific strategy profile against all alternative strategy profiles, we can confirm that the strategy profile $\{\text{Boxing, Boxing}\}$ also is Pareto efficient. On the other hand, the

strategy profiles {Opera, Boxing} and {Boxing, Opera} are not Pareto efficient, as one can easily confirm that any Nash equilibrium in this game is Pareto superior (i.e., improves the wellbeing of both players) to these strategy profiles.

Coordination failure is a common challenge in battle-of-the-sexes types of games due to strategic uncertainty, where players face the risk of selecting conflicting strategies—like one choosing A while the other opts for B. In these games, players often have preferences for different Nash equilibria as there's no clear Pareto Dominance, leading to uncertainty about the final outcome. Each pure Nash Equilibrium can be perceived as somewhat "unfair" in layman's terms, as it will be the most preferred by some but not by all players!

To address these inefficiencies, one potential solution is to explore *correlated equilibrium*. In a correlated equilibrium, players receive signals or guidance that influence their strategy choices, fostering better coordination and potentially yielding more favorable outcomes for all players involved. However, we won't be discussing this equilibrium concept further.

3. Pure coordination games

There are multiple Nash Equilibria, but only one is Pareto efficient. Despite the multiplicity of Nash equilibria, players' objectives are perfectly aligned. As a result, all players will, almost certainly, end up choosing their strategy that corresponds to the Nash equilibrium that Pareto dominates¹⁸ all other Nash equilibria.

It may seem that coordination failure isn't a significant concern in such games, but is that really the case? While payoff dominance plays a significant role in such games, it's not the sole determinant. Consider this scenario: you're a participant in a game where one specific Nash Equilibrium clearly outshines all others in terms of Pareto dominance. Consequently, you opt for a strategy aligned with the outcome associated with that particular Nash Equilibrium, expecting rational players to do the same. However, even if you're fully confident that your opponents are rational, there's always a slight chance that someone might make an error and select a different strategy than the one consistent with the specific Nash Equilibrium. This risk should be factored in when comparing Nash Equilibria.

Thus, alongside the notion of Pareto dominance, the concept of risk dominance emerges as a means of evaluating different Nash Equilibria. Risk dominance evaluates equilibria based on their risk levels, considering the uncertainty surrounding opponents' actions. An equilibrium is deemed risk-dominant if it entails less risk compared to all other Nash Equilibria in the game. To gauge and compare risks across different Nash Equilibria, one must calculate the product of deviation losses. The higher the product of deviation losses is, the more likely that the player will be

¹⁸ A Nash Equilibrium where each player obtains a payoff higher than what they receive in another Nash Equilibrium is referred to as Pareto dominant over that other Nash Equilibrium.

“careful” when choosing and the less likely to make a mistake, hence the risk is reduced. Examples 13 and 14 provided subsequently will elucidate this concept further.

Example 13

Two firms produce complementary products. However, for the products to be used together there must be technologically compatible. Before launching the products in the market, the two firms simultaneously choose one of two available technologies, A or B. If the technologies are incompatible, both products will fail and the firms will make no profits. Conversely, if the technologies are compatible, profits will be generated. Nevertheless, both firms receive higher profits when technology A is chosen. The payoff matrix is presented below where the best responses are indicated

		Firm 2	
		A ₂	B ₂
Firm 1	A ₁	<u>10</u> , <u>8</u>	0, 0
	B ₁	0, 0	<u>6</u> , <u>5</u>

by underlines. It is straightforward to confirm that this game possesses two Nash Equilibria, namely, $N.E._1 = \{A_1; A_2\}$ and $N.E._2 = \{B_1; B_2\}$. That is, in equilibrium we are expecting the two firms to choose the same technology. Why is there a possibility

for coordination failure? As there are multiple (two) Nash Equilibria there is always a chance that firm 1, for example, will choose technology A (expecting that firm 2 will do the same) while firm 2 will choose technology B (expecting, of course, that firm 1 will choose technology B). The possibility of this happening can be evaluated by checking for (a) *Pareto Dominance*, and (b) *Risk Dominance*.

- Note that both firms are getting higher profits should $N.E._1$ becomes the final outcome of the game. Specifically, firm 1 earns a profit of €10 under $N.E._1$ instead of a profit of €6 under $N.E._2$. Similarly, firm 2 earns a profit of €8 under $N.E._1$ instead of a profit of €5 under $N.E._2$. Therefore, $N.E._1$ Pareto dominates $N.E._2$.
- To analyze risk dominance, we need to consider the potential losses a firm might incur when erroneously deviates from a specific Nash Equilibrium while the other firm selects an action aligned with that equilibrium. For instance, let's take Nash Equilibrium 1 where both firms are expected to choose technology A, and suppose firm 2 correctly selects A_2 . If firm 1 mistakenly chooses B_1 instead of A_1 , the game's outcome will be $\{B_1; A_2\}$, resulting in firm 1 receiving €0 instead of the anticipated €10. This translates to a loss of $(10 - 0) = 10$ for firm 1 due to this error. Similarly, assuming firm 1 correctly selects A_1 , if firm 2 erroneously chooses B_2 instead of A_2 , the game's outcome will be $\{A_1; B_2\}$, causing firm 2 to receive €0 instead of the expected €8. Consequently, firm 2 experiences a loss of $(8 - 0) = 8$ due to this mistake. Therefore, the combined loss resulting from deviations for Nash Equilibrium 1 (*i.e.*, the product of deviation losses) is $10 \times 8 = 80$.

Now, let's turn our attention to Nash Equilibrium 2 where both firms are expected to choose technology B. If firm 2 correctly chooses B_2 , but firm 1 mistakenly chooses A_1 instead of B_1 , the outcome will be $\{A_1; B_2\}$, resulting in firm 1 receiving €0 instead of the anticipated €6. This leads to a loss of $(6 - 0) = 6$ for firm 1 due to this error. Similarly, assuming firm 1 correctly selects B_1 , if firm 2 erroneously chooses A_2 instead of B_2 , the outcome will be $\{B_1; A_2\}$, causing firm 2 to receive €0 instead of the expected €5. Consequently, firm 2 experiences a loss of $(5 - 0) = 5$ due to this mistake. Therefore, the combined loss resulting from deviations for Nash Equilibrium 2 is $6 \times 5 = 30$.

Clearly, the combined loss resulting from deviations for Nash Equilibrium 1 exceeds that of Nash Equilibrium 2 (*i.e.*, 80 versus 30), indicating that the two firms stand to lose more (and thus will be more cautious) if something goes wrong with Nash Equilibrium 1. In other words, Nash Equilibrium 1 exhibits risk dominance over Nash Equilibrium 2.

Example 13 is a typical example of pure coordination games where the possibility of coordination failure is slim. This is since one and the same Nash Equilibrium both Pareto and Risk dominates the other Nash Equilibrium. *Example 14* that follows presents a slightly different picture with one Nash Equilibrium being Pareto Dominant and the other Nash Equilibrium being Risk Dominant. Consequently, in that case the likelihood of coordination failure becomes more pronounced.

Example 14: (assurance or stag-hunt game)

Now consider a slight modification in the previous game. Two firms produce complementary products. However, for the products to be used together they must be technologically compatible. Firms choose (simultaneously) to introduce a new technology A or to stick with the traditional technology B prior to introducing the product in the market. If technologies do not match the product that adopts the new tech fails while the other earns some profit. If technologies match, profits are earned.

		Firm 2	
		A ₂	B ₂
Firm 1	A ₁	<u>10</u> , <u>8</u>	0, 4
	B ₁	5, 0	<u>6</u> , <u>5</u>

However, both firms receive higher profits when new technology A is chosen. The payoff matrix is presented to the left where the best responses are indicated by underlines. It is again straightforward to confirm that this game possesses two Nash

Equilibria, namely, $N.E._1 = \{A_1; A_2\}$ and $N.E._2 = \{B_1; B_2\}$. We will be checking again for (a) *Pareto Dominance*, and (b) *Risk Dominance*.

- (a) Note that both firms are getting higher profits should $N.E._1$ becomes the final outcome of the game. Specifically, firm 1 earns a profit of €10 under $N.E._1$ instead of a profit of €6 under $N.E._2$. Similarly, firm 2 earns a profit of €8

under $N.E._1$ instead of a profit of €5 under $N.E._2$. Therefore, $N.E._1$ Pareto dominates $N.E._2$.

- (b) First, consider Nash Equilibrium 1 where both firms are expected to choose technology A, and suppose firm 2 correctly selects A_2 . If firm 1 mistakenly chooses B_1 instead of A_1 , the game's outcome will be $\{B_1; A_2\}$, resulting in firm 1 receiving €5 instead of the anticipated €10. This translates to a loss of $(10 - 5) = 5$ for firm 1 due to this error. Similarly, assuming firm 1 correctly selects A_1 , if firm 2 erroneously chooses B_2 instead of A_2 , the game's outcome will be $\{A_1; B_2\}$, causing firm 2 to receive €4 instead of the expected €8. Consequently, firm 2 experiences a loss of $(8 - 4) = 4$ due to this mistake. Therefore, the combined loss resulting from deviations for Nash Equilibrium 1 (*i.e.*, the product of deviation losses) is $5 \times 4 = 20$.

Now, let's turn our attention to Nash Equilibrium 2 where both firms are expected to choose technology B. If firm 2 correctly chooses B_2 , but firm 1 mistakenly chooses A_1 instead of B_1 , the outcome will be $\{A_1; B_2\}$, resulting in firm 1 receiving €0 instead of the anticipated €6. This leads to a loss of $(6 - 0) = 6$ for firm 1 due to this error. Similarly, assuming firm 1 correctly selects B_1 , if firm 2 erroneously chooses A_2 instead of B_2 , the outcome will be $\{B_1; A_2\}$, causing firm 2 to receive €0 instead of the expected €5. Consequently, firm 2 experiences a loss of $(5 - 0) = 5$ due to this mistake. Therefore, the combined loss resulting from deviations for Nash Equilibrium 2 is $6 \times 5 = 30$.

Clearly, the combined loss resulting from deviations for Nash Equilibrium 2 exceeds that of Nash Equilibrium 1 (*i.e.*, 30 versus 20), indicating that the two firms stand to lose more (and thus will be more cautious) if something goes wrong with Nash Equilibrium 2. In other words, Nash Equilibrium 2 exhibits risk dominance over Nash Equilibrium 1.

Mixed Strategies

In numerous strategic scenarios, players typically opt for singular actions from a predefined set of options, a concept known as *pure strategies*. However, there arise situations where a player might find it advantageous to introduce an element of randomness into their decision-making process, particularly in scenarios where players face uncertainty or lack perfect information about their opponents' actions. When a player decides to distribute their choices probabilistically across available actions, they engage in what is termed a *mixed strategy*, contrasting with the certainty of a pure strategy.

In essence, while pure strategies entail players committing to specific actions with certainty, mixed strategies introduce a degree of unpredictability by allowing players to assign probabilities to their potential actions. This introduces a layer of

complexity and strategic flexibility, as players strategically allocate probabilities to various actions, considering both their own preferences and their anticipation of opponents' moves.

By embracing mixed strategies, players can navigate strategic landscapes with greater versatility, strategically blending certainty and uncertainty to optimize their outcomes. This paradigm shift from pure to mixed strategies not only enriches the strategic depth of games but also mirrors real-world decision-making scenarios where uncertainty and adaptability play crucial roles. Through an exploration of mixed strategies, we unravel the intricate dynamics of strategic decision-making under uncertainty, offering insights into how players can strategically leverage probabilistic approaches to achieve their objectives.

Definition

- A mixed strategy for player i is a probability distribution over her set of m available actions $A_i = (a_{i,1}, a_{i,2}, \dots, a_{i,m})$. In other words, a mixed strategy is an m -dimensional vector $\sigma_i = (p_{i,1} a_{i,1}, p_{i,2} a_{i,2}, \dots, p_{i,m} a_{i,m})$ such that for all $i \in N$ and for all $k = 1, 2, \dots, m$, we have $p_{i,k} \geq 0$ and $\sum_{k=1}^m p_{i,k} = 1$.

To better understand the concept of mixed strategies, consider a player in a game having to choose between three actions, namely a_1 , a_2 , and a_3 . She can, of course, choose to play a_1 , which is a pure strategy¹⁹. However, the player can also choose a strategy according to a simple rule that involves the roll of a dice: if 1, 2, or 3 appears after rolling the dice the player will play a_1 , if 4 or 5 appears after rolling the dice she will play a_2 , and if 6 appears after rolling the dice she will play a_3 . Therefore, before rolling the dice the probabilities of choosing a specific action are 1/2 for a_1 , 1/3 for a_2 , and 1/6 for a_3 . Of course, there are infinitely many ways to assign probabilities over her three available actions. The only constraints are that probabilities must be non-negative and that they must add up to unit.

But why should a person involved in a situation of strategic interdependencies opt for a mixed instead of a pure strategy? One of the reasons is that the use of mixed strategies adds an element of unpredictability and strategic complexity, which can be advantageous in competitive environments where opponents are trying to anticipate and counter each other's moves. Consider the following examples:

- Bluffing in Poker: In poker, a player might choose a mixed strategy of betting aggressively with strong hands and bluffing with weaker hands. By mixing these strategies, the player makes it difficult for opponents to predict their actions, potentially leading to higher gains.
- Sports Strategies: In sports like football, players might adopt mixed strategies when taking penalty shots. For instance, a football player might vary the

¹⁹ Any pure strategy is also a (degenerated) mixed strategy where a probability of zero is assigned to all but one actions and a probability of one is assigned on the remaining action.

direction of his penalty shots randomly between left and right, making it harder for the goalkeeper to anticipate the shot's trajectory.

- **Economic Competition:** In competitive markets, companies might choose mixed pricing strategies, where they randomly adjust prices within a range to prevent competitors from easily predicting their pricing moves. This can help maintain market share and prevent competitors from undercutting prices.
- **Military Tactics:** In warfare, commanders may use mixed strategies to confuse and outmaneuver opponents. For example, a military unit might alternate between attacking head-on and using guerrilla tactics to keep the enemy off balance.
- **Political Negotiations:** In diplomatic negotiations, politicians might employ mixed strategies to negotiate more effectively. They may alternate between taking hardline stances and making concessions to gain leverage and achieve their objectives.

Given that we can now include in our discussion mixed strategies we redefine the concept of Nash Equilibrium, accordingly.

Definition


- A strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*, \dots, \sigma_n^*)$ is a **Mixed Strategy Nash Equilibrium**²⁰ (**MSNE**) if for every player $i \in N$ and all strategies $\sigma'_i \neq \sigma_i^*$ we have

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma'_i, \sigma_{-i}^*)$$

Essentially, as probabilities are involved in the way the game is played, players select strategies aiming to maximize their expected payoffs. Therefore, our objective is to set up the expected payoffs for both players and let each to select a mixed strategy that maximizes their respective expected payoff. As demonstrated in *Example 15*, this process demands some effort. Nonetheless, as illustrated in the same example, there exists a shortcut for identifying the Mixed Strategy Nash Equilibrium (MSNE) in such games.

Example 15: (the penalty-kick game)

A goalkeeper prepares to defend against a penalty shot from the striker of the rival team, who has the option to shoot left or right. The goalkeeper has the choice to dive to the left or right. To prevent any ambiguity regarding the directions above referenced as 'right' and 'left,' please refer to the accompanying image on the right side. The scenario assumes that if the goalkeeper dives in the opposite direction of the striker's shot, the ball will always score. However, if the goalkeeper correctly anticipates the striker's shot



²⁰ Note that, according to the definition of mixed strategies and the concept of MSNE, any Nash Equilibrium in pure strategies is also a MSNE but not vice versa.

direction, he will save the shot. According to the success/failure rates mentioned, the

		Goalie	
		Right	Left
Striker	Right	0, 1	1, 0
	Left	1, 0	0, 1

payoff matrix is presented to the left. By now, you should be able to confirm that this game has no Nash Equilibrium in which each player chooses a pure strategy. But what if

we allow the players to choose mixed strategies? Assume that the striker chooses right with probability p (and thus, left with probability $(1 - p)$) and the goalie chooses right with probability q (and thus, left with probability $(1 - q)$). To better incorporate the impact of randomization on determining a MSNE, the payoff matrix of the game is modified, as seen to the right, to reflect the probabilistic choices made by each player.

		Goalie	
		q Right	$(1 - q)$ Left
Striker	p Right	0, 1	1, 0
	$(1 - p)$ Left	1, 0	0, 1

Calculating the players' expected payoffs requires identifying the likelihood of each potential outcome. If the probabilities p and q are independent, the likelihood of any of the four potential outcomes happening can be determined by multiplying the probabilities assigned by each player to a specific action. For instance, the likelihood that the striker shoots to the right and the goalie dives in the same direction is pq . Likewise, the chance that the striker shoots to the right while the goalie dives in the opposite direction is $p(1 - q)$. Similarly, the probability that the striker shoots to the left while the goalie dives to the right is $(1 - p)q$. Lastly, the probability that the striker shoots to the left and the goalie dives in the same direction is $(1 - p)(1 - q)$. Given these, the expected payoff of the striker is given by²¹

$$EU_S = pq \times 0 + p(1 - q) \times 1 + (1 - p)q \times 1 + (1 - p)(1 - q) \times 0 \Rightarrow$$

$$EU_S = p(1 - q) + (1 - p)q \Rightarrow \boxed{EU_S = p + q - 2pq}$$

Therefore, the striker must choose his probability distribution (*i.e.*, p and $(1 - p)$) over his actions "left" and "right" so that the above expression is maximized. Formally, this process can be expressed as

$$\max_{0 \leq p \leq 1} \{EU_S = p + q - 2pq\}$$

Note, however, that the expected payoff of the striker depends linearly on probability p . This suggests that depending on whether the first derivative of the expected payoff with respect to p is positive or negative, the optimal probability p could either be 1 or 0, respectively (or, in an extreme scenario, any value between 0 and 1 inclusive). It is easy to confirm that the first derivative of the EU_S is

²¹ Working on a similar manner you can confirm that the expected payoff of the goalie is $EU_G = 1 - p - q + 2pq$.

$$\frac{\partial EU_S}{\partial p} = 1 - 2q$$

so that this derivative is positive if $q < 1/2$, negative if $q > 1/2$, and zero if $q = 1/2$. Put simply, if increasing the probability p results in a higher expected payoff for the striker (*i.e.*, $q < 1/2$), then the striker should opt for the maximum possible value of p , *i.e.*, $p = 1$. On the other hand, if increasing the probability p results in a lower expected payoff for the striker (*i.e.*, $q > 1/2$), then the striker should opt for the minimum possible value of p , *i.e.*, $p = 0$. In the event the probability is such that the first derivative of the expected payoff with respect to p is zero (*i.e.*, $q = 1/2$), the striker will be indifferent between any value of p , *i.e.*, $p \in [0, 1]$. Did you notice what have we just described? The above is a full description of the striker's **best response**! Formally, the best response of the striker is described by

$$p^* = \begin{cases} 1, & \text{if } q < 1/2 \\ p \in [0, 1], & \text{if } q = 1/2 \\ 0, & \text{if } q > 1/2 \end{cases}$$

Working similarly, one can confirm that the best response of the goalie is given by

$$q^* = \begin{cases} 1, & \text{if } p > 1/2 \\ q \in [0, 1], & \text{if } p = 1/2 \\ 0, & \text{if } p < 1/2 \end{cases}$$

In equilibrium (in mixed strategies), it must be that both best responses are satisfied. Can it be that in an equilibrium $p^* = 1$? No, because since $p^* = 1 > 1/2$, the best response of the goalie will be to choose $q^* = 1$. But given that $q^* = 1 > 1/2$, the striker's best response is $p^* = 0$ and not $p^* = 1$. Working similarly, we can exclude the case where being in an $p^* = 0$ equilibrium. It is left to see that in a MSNE the Striker will be indifferent between any of his actions and so does the goalie. But for this to be realized it must be that the goalie chooses $q^* = 1/2$ and the striker chooses $p^* = 1/2$. Hence, the MSNE of this game is

$$\text{MSNE} = \{p^* R_S(1 - p^*) L_S; q^* R_G(1 - q^*) L_G\} = \left\{ \frac{1}{2} R_S \frac{1}{2} L_S; \frac{1}{2} R_G \frac{1}{2} L_G \right\}$$

All this looks like an awful long process to identify a MSNE. Fortunately, as shown in *Example 16* that follows, there is a much easier way to do it. It turns out this game's MSNE occurs when (1) the striker's optimal probability p is such that goalkeeper is indifferent between choosing Left or Right (*i.e.*, the goalie's expected payoff from choosing Left equals her expected payoff from choosing Right), and (2) the goalie's optimal probability q is such that the striker is indifferent between choosing Left or

Right (i.e., the striker's expected payoff from choosing Left equals his expected payoff from choosing Right).²²

Example 16: (the penalty-kick game revisited)

Taking requirements (1) and (2) stated above into account we can find the MSNE of the game presented in *Example 15* using a different approach. First, we calculate the expected payoff of the striker should he choose to shoot Right, and his expected payoff should he choose to shoot Left. In case he chooses Right he receives either 0 (if the goalie chooses Right – with probability q) or 1 (if the goalie chooses Left – with probability $(1 - q)$). Therefore, his expected payoff from choosing Right depends on the probability the goalie dives to the right:

		Goalie	
		q Right	$(1-q)$ Left
Striker	p Right	0, 1	1, 0
	$(1-p)$ Left	1, 0	0, 1

$$EU_S(q|R_S) = q \times 0 + (1 - q) \times 1 = 1 - q$$

Similarly, in case the striker chooses Left he receives either 1 (if the goalie chooses Right – with probability q) or 0 (if the goalie chooses Left – with probability $(1 - q)$). Therefore, his expected payoff from choosing Right depends on the probability the goalie dives to the right:

$$EU_S(q|L_S) = q \times 1 + (1 - q) \times 0 = q$$

The striker will be indifferent between Right or Left if the two expected payoffs we calculated above are equal, that is

$$EU_S(q|R_S) = EU_S(q|L_S) \Rightarrow 1 - q = q \Rightarrow q^* = 1/2$$

Note that this is indeed the optimal probability distribution for the goalie we found using the long way! Clearly, due to symmetry of the game we can find that the probability distribution chosen by the striker that makes the goalie indifferent is characterized by $p^* = 1/2$. Hence, clearly,

$$MSNE = \{p^* R_S(1 - p^*) L_S; q^* R_G(1 - q^*) L_G\} = \left\{ \frac{1}{2} R_S \frac{1}{2} L_S; \frac{1}{2} R_G \frac{1}{2} L_G \right\}$$

What are the expected payoffs of the two players out of this game? Having found the optimal probability distributions of the two players we can now identify the likelihood of each potential outcome. Since each player chooses either of his actions with probability $(1/2)$, each outcome will arrive with probability $(1/2) \times (1/2) = (1/4)$. Therefore, the striker's expected payoff will be²³

$$EU_S(p^*, q^*) = 0,25 \times 0 + 0,25 \times 1 + 0,25 \times 0 + 0,25 \times 1 = 0,5^{24}$$

²² Once again, make clear that neither player wants to make the other player indifferent. Each player is maximizing his expected payoff. However, the optimal choice of a player (i.e., the optimal probability distribution over his actions) happens to make the other player indifferent.

²³ By the symmetry of the game, the goalie's expected payoff will be the same.

²⁴ Alternatively, one can get the same result by plugging $q^* = 0,5$ in $EU_S(q|L_S)$.

The approach we have used to identify a MSNE in *Example 16* can be proven to be valid for all cases. *Example 17* that follows (the assurance game of *Example 14*) uses the exact same method for finding the MSNE.

Example 17: (The stag-hunt game and mixed strategies)

Two firms produce complementary products. However, for the products to be used together they must be technologically compatible. Firms choose (simultaneously) to introduce a new technology A or to stick with the traditional technology B prior to introducing the product in the market. If technologies do not match the product that

		Firm 2	
		(q) A ₂	(1-q) B ₂
Firm 1	(p) A ₁	10, 8	0, 4
	(1-p) B ₁	5, 0	6, 5

adopts the new tech fails while the other earns some profit. If technologies match, profits are earned. However, both firms receive higher profits when new technology A is chosen. Allowing the two

firms to randomize over their actions, the payoff matrix, along with the probability distributions, is presented to the left.

First, we consider Firm 1. In case Firm 1 chooses A_1 its profits will either be 10 (if the Firm 2 chooses A_2 – with probability q) or 0 (if the Firm 2 chooses B_2 – with probability $(1 - q)$). Therefore, his expected payoff from choosing A_1 depends on the probability Firm 2 assigns to developing technology A_2 :

$$EU_1(q|A_1) = q \times 10 + (1 - q) \times 0 = 10q$$

Similarly, Firm 1 chooses B_1 its profits will either be 5 (if the Firm 2 chooses A_2 – with probability q) or 6 (if the Firm 2 chooses B_2 – with probability $(1 - q)$). Therefore, his expected payoff from choosing B_1 depends on the probability Firm 2 assigns to developing technology A_2 :

$$EU_1(q|B_1) = q \times 5 + (1 - q) \times 6 = 6 - q$$

Firm 1 will be indifferent between the two technologies if the two expected payoffs we calculated above are equal, that is

$$EU_1(q|A_1) = EU_1(q|B_1) \Rightarrow 10q = 6 - q \Rightarrow q^* = 6/11$$

Note that by substituting $q^* = 6/11$ in either $EU_1(q|A_1)$ or $EU_1(q|B_1)$ we get the expected payoff of Firm 1, i.e., $EU_1 = 10(6/11) = 60/11$. Working similarly, you should be able to confirm that $p^* = 5/9$ and $EU_2 = 40/9$. Therefore, we get that

$$MSNE = \left\{ \frac{5}{9}A_1 \frac{4}{9}B_1; \frac{6}{11}A_2 \frac{5}{11}B_2 \right\}$$

Rationalizing mixed strategies

The idea of mixed strategies can initially seem perplexing. One might question why and how players would introduce randomness into their decision-making process. Randomizing between potential actions, a core aspect of mixed strategies, doesn't

often align with typical human behavior. Rarely do individuals base their choices on a lottery-like approach.

This behavioral challenge is further complicated by the cognitive barrier; people struggle to produce genuinely random outcomes without the assistance of a random or pseudo-random generator.

So, why bother with mixed strategies then? Despite these hurdles, mixed strategies remain prevalent due to their ability to yield Nash equilibria in games where equilibrium in pure strategies is unattainable. Does this suffice? Of course not. We must be able to rationalize the use of mixed strategies. Below are some examples of rationalization of mixed strategies.

- Rationalizing mixed strategies involves recognizing them as deliberate choices. That is, players actively introduce randomness into their actions, keeping this randomness undisclosed beforehand. Consider the dynamic between tax authorities and taxpayers, where random audits serve as a parallel. However, authorities might prefer to disclose the probability of an audit before gameplay to make it transparent to taxpayers.
- Mixed strategies emerge from players' beliefs about their opponents' behaviors. For instance, if Plaisio observes that MediaMarkt has offered exclusive deals on tablets in the past, it might base its strategy on the belief that MediaMarkt's actions were strategic responses rather than random occurrences.
- Furthermore, mixed strategies can stem from the varied dispositions or "moods" of a player. For example, a Bertrand competitor might alternate between aggressive and passive tactics based on their mood. This introduces a layer of unpredictability, yet it also complicates strategic analysis.

In summary, while mixed strategies can offer solutions to strategic dilemmas, their implementation raises challenges regarding the interpretation of opponents' behaviors, the rationale behind past actions, and the influence of personal dispositions on decision-making processes.

PART II: EXTENSIVE FORM GAMES OF COMPLETE AND PERFECT INFORMATION

From strategic form to extensive form

In the strategic-form games examined previously, the focus was on players choosing strategies and on whether the resulting strategy profile constituted a Nash equilibrium. This approach is appropriate for situations in which players choose simultaneously, that is, when no player observes the choices of the others before making her own decision. However, many strategic interactions do not unfold in this way. In many strategic situations, players move in sequence. One player acts first, another player observes that action and then responds, and the game continues in this manner until an outcome is reached. In such cases, the order of play is not a secondary detail but a fundamental part of the strategic situation. To study such interactions, we use the **extensive form** of a game.

The importance of timing can be understood easily through simple examples. In a market-entry problem, for instance, a potential entrant may first decide whether to enter an industry, and only if entry occurs does the incumbent firm decide how to react. Similarly, in a bargaining problem one player may first make an offer, and the other player may then decide whether to accept or reject it. In such situations, the player who moves later can condition her action on what has already happened. Therefore, the strategic problem cannot be understood properly unless the sequence of moves is made explicit. A representation that ignores timing may hide essential aspects of rational behavior.

This is precisely the limitation of the strategic form when it is applied to sequential interactions. A strategic-form representation records the players, their available strategies, and the payoffs associated with every strategy profile. For this reason, it is extremely useful and will remain important throughout the analysis of games. Yet, when the game unfolds over time, the strategic form does not show directly **when** decisions are made, **what** players know when they act, and **how** the game develops from one choice to the next. These features are instead represented naturally in the extensive form, where the game is described as a sequence of decision nodes connected by branches corresponding to the actions available to the players.

A second important point concerns the distinction between **action** and **strategy**. In strategic-form games, these two terms coincide. If a player chooses once and for all among a set of possible moves, then her action is also her strategy. In extensive-form games, however, the two concepts are no longer identical. An **action** is a move taken at a particular point of the game. A **strategy**, by contrast, is a complete plan of behavior. It specifies what a player would do at every point in the game where she may be called upon to act, including points that may never actually be reached in

equilibrium. This distinction is essential and will play a central role in the analysis that follows.

To see why this distinction matters, consider a player who may have to make a decision after one of several possible earlier moves by her opponent. In such a case, the player's strategy must indicate what she would do after each possible history of play. Therefore, a strategy in an extensive-form game is a contingent plan: it tells us how the player would respond in every circumstance that could arise within the game. This is one of the main reasons why extensive-form analysis is richer than strategic-form analysis when timing and observation are important.

The move from strategic form to extensive form is therefore not merely a change in notation or in graphical representation. It reflects a deeper change in the way strategic interaction is analyzed. Once timing becomes explicit, we are no longer interested only in whether the players' strategies are mutual best responses in a static sense. We must also ask whether the behavior prescribed by those strategies is reasonable at every stage of the game. A player may announce one course of action at the beginning of the game, but if the relevant decision point is eventually reached, it may no longer be optimal for her to carry it out. This means that making the sequential structure explicit may reveal weaknesses in the Nash equilibrium concept when it is applied to dynamic settings.

For this reason, the transition to extensive-form games will also require a refinement of the equilibrium concept used in strategic-form analysis. In sequential games, Nash equilibrium may sometimes allow predictions based on behavior that is not credible once the moment of action arrives. This is why, after introducing the extensive form of a game, we will proceed to a stronger solution concept, namely the **Subgame Perfect Nash Equilibrium**, which is better suited to the analysis of sequential rationality.

Why Nash equilibrium may be too weak in sequential games

In strategic-form games, the Nash equilibrium concept is based on a simple and powerful idea: each player chooses a best response to the strategies of the others, so no player has an incentive to deviate unilaterally. In simultaneous-move games, this idea often provides a natural and convincing prediction of the outcome. Once we move to sequential interactions, however, an important question arises: does Nash equilibrium remain an appropriate solution concept when players move one after another and later players observe earlier actions?

The answer is not immediate. In some sequential games, Nash equilibrium continues to provide the correct prediction. In other cases, however, it proves too weak, because it may allow equilibria supported by actions that are not credible once the relevant stage of the game is actually reached. To understand why a refinement is needed, it is useful to examine both types of cases.

Example 18: The Prisoner's Dilemma with a twist

Let us begin with a familiar example. Consider again the Prisoner's Dilemma game introduced earlier, but now with a twist. In the first stage, Prisoner 1 chooses whether to Confess or Lie. In the second stage, Prisoner 2 observes Prisoner 1's choice and then chooses whether to Confess or Lie. Payoffs are realized at the end according to the same payoff matrix as in the simultaneous version of the game. This is exactly the kind of introductory example used in your slides to motivate extensive-form analysis by taking a known strategic-form game and changing only the timing of play. The strategic-form payoff matrix is the following, as shown in Part I:

		Player 2	
		Lie (L_2)	Confess (C_2)
Player 1	Lie (L_1)	-1, -1	-9, 0
	Confess (C_1)	0, -9	-6, -6

In the simultaneous version, we already know that the unique Nash equilibrium is {Confess, Confess}. Each prisoner has a strictly dominant strategy to confess, since confessing yields a better payoff no matter what the other prisoner does. Your earlier notes explain this logic in detail.

What changes if the game is played sequentially? Suppose Prisoner 2 observes what Prisoner 1 has done before making his own choice. If Prisoner 1 has lied, then Prisoner 2 compares the payoff from lying, which is -1, with the payoff from confessing, which is 0. Therefore, after observing Lie, Prisoner 2 prefers to Confess. If Prisoner 1 has confessed, then Prisoner 2 compares the payoff from lying, which is -9, with the payoff from confessing, which is -6. Therefore, after observing Confess, Prisoner 2 again prefers to Confess.

Thus, regardless of what Prisoner 1 does in the first stage, Prisoner 2's optimal response in the second stage is to confess. Prisoner 1, anticipating this, understands that if he lies, Prisoner 2 will confess and Prisoner 1 will get -9, whereas if he confesses, Prisoner 2 will still confess and Prisoner 1 will get -6. Since -6 is better than -9, Prisoner 1 will also choose to confess. Hence, the outcome of the sequential version of the game is again {Confess, Confess}.²⁵

This example is important because it shows that making the game sequential does not by itself create a problem for Nash equilibrium. Here, the same outcome arises in both the simultaneous and the sequential version of the game. The reason is that the continuation behavior of the second mover is fully credible. Prisoner 2's choice to confess is optimal no matter what Prisoner 1 has done, so there is no room for threats or promises that would later be abandoned. In this case, sequential

²⁵ However, as we will see later, the equilibrium of the game needs to be expressed more thoroughly than just {Confess, Confess}.

rationality points to the same outcome as the Nash equilibrium of the strategic-form version.

At this stage, one might be tempted to conclude that Nash equilibrium remains sufficient even in sequential settings. However, this conclusion would be premature. There are sequential games in which the strategic-form Nash equilibrium allows outcomes that are sustained only by threats that the players would not actually want to carry out. In such cases, Nash equilibrium is too weak.

Example 19: The entry-deterrence game

To see where the problem arises, consider now a different example. There are two firms. Firm 1 is a potential entrant into a market currently served by Firm 2, the incumbent. Firm 1 can choose either to Enter the market or to Stay Out. Firm 2 can choose either to Fight entry aggressively or to Accommodate it. Suppose the strategic-form payoff matrix is the following:

		Firm 2	
		Fight	Accommodate
Firm 1	Enter	-1, 1	2, 2
	Stay out	0, 5	0, 5

If Firm 1 stays out, Firm 2 remains a monopolist and earns 5, while Firm 1 earns 0. If Firm 1 enters and Firm 2 accommodates, the firms share the market and both earn positive profits. If Firm 1 enters and Firm 2 fights, both firms are worse off than under accommodation, and the entrant is especially harmed.

Let us first analyze this game in strategic form. There are two Nash equilibria. The first is {Enter, Accommodate}. If Firm 2 accommodates, Firm 1 prefers to enter rather than stay out, since 2 is greater than 0. And if Firm 1 enters, Firm 2 prefers to accommodate rather than fight, since 2 is greater than 1. Therefore, neither firm wishes to deviate unilaterally.

The second Nash equilibrium is {Stay Out, Fight}. If Firm 2's strategy is Fight, then Firm 1 prefers to stay out, since 0 is greater than -1. Given that Firm 1 stays out, Firm 2 is indifferent between Fight and Accommodate, because in either case Firm 1 does not enter and Firm 2 gets 5. Therefore, this strategy profile also satisfies the definition of Nash equilibrium.

Now let us reconsider the same interaction as a sequential game. In the first stage, Firm 1 chooses whether to enter or stay out. In the second stage, if entry occurs, Firm 2 observes this and then chooses whether to fight or accommodate. Once the timing is made explicit, the logic of the game changes in an important way.

Suppose Firm 1 enters. At that point, Firm 2 must choose between Fight and Accommodate. By inspecting the payoffs, we see immediately that Firm 2 prefers to accommodate, since 2 is greater than 1. Therefore, if entry occurs, Firm 2 will not fight. Anticipating this, Firm 1 understands that entry will be followed by accommodation

and will therefore yield 2 rather than 0. Hence Firm 1 will choose to enter. The sequentially rational outcome of the game is therefore that Firm 1 enters and Firm 2 accommodates.

The key observation is that the Nash equilibrium {Stay Out, Fight}, although admissible in the strategic-form representation, is not a convincing prediction once the game is understood as sequential. The reason is that Firm 2's threat to fight entry is not credible. If the moment actually comes to choose, Firm 2 will prefer to accommodate rather than fight. The strategic-form Nash equilibrium allows this threat because it asks only whether strategies are mutual best responses for the game as a whole. It does not ask whether each action remains optimal at the stage where it must actually be chosen.

This is the main issue to remember: in extensive-form games of complete and perfect information, credibility is central, and the player who moves later can only be expected to choose actions that are best responses once the earlier player's choice has been observed. A threat to behave otherwise is a non-credible threat and should not be used to support a solution of the game. For this reason, the ordinary Nash equilibrium concept used in static games is not sufficient for extensive-form games.

Why a refinement is needed

The two examples above clarify the issue. The Prisoner's Dilemma with a twist shows that sequential play does not necessarily create a problem for Nash equilibrium. In that game, the continuation behavior of the second mover is credible, and the sequential logic of the game leads to the same outcome as in the simultaneous version. The entry-deterrence game shows, however, that in other sequential interactions Nash equilibrium may be too weak because it does not distinguish between credible and non-credible future actions.

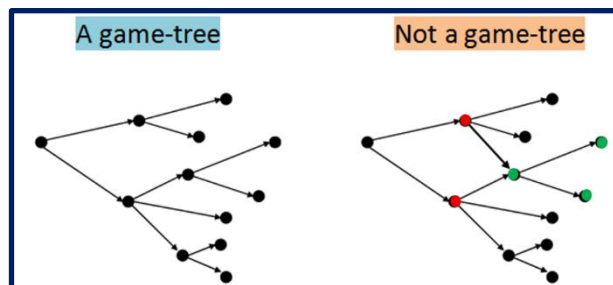
Therefore, once we move from strategic-form to extensive-form analysis, we need a stronger equilibrium concept. We need a concept that requires rational behavior not only at the beginning of the game but also after every possible history of play. In other words, we need a concept that rules out equilibria sustained by non-credible threats or promises. This leads naturally to the concept of **Subgame Perfect Nash Equilibrium**, which is the standard solution concept for extensive-form games of complete and perfect information. In such games, the appropriate method of solution is backward induction, which proceeds from the last stage of the game to the first.

Representation of extensive-form games

The examples discussed above make clear that, once the order of play matters, the strategic-form representation is no longer the most natural way to describe a game. In sequential interactions, we are interested not only in the final outcome and the corresponding payoffs, but also in the way the game unfolds over time. We want to know who moves first, what actions are available at each stage, what later players

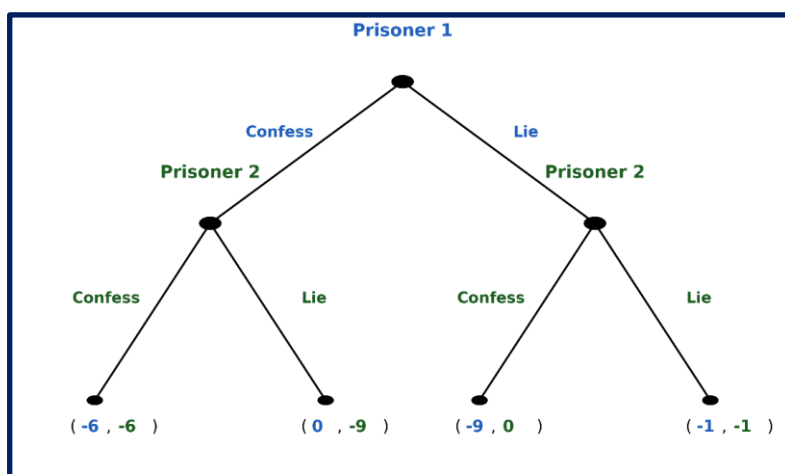
observe before making their decisions, and how the sequence of choices leads to the final outcome. For this reason, extensive-form games are represented by means of game trees.

A game tree is a graphical representation of a sequential interaction. It begins from a single initial node, which represents the starting point of the game. From this node, branches emerge, each corresponding to one of the actions available to the player who moves first. Each branch leads to another node, where either the same player or another player must make a new choice. In this way, the game tree describes the game as a sequence of decision points connected by actions. Nodes that are not followed by any further choices are called terminal nodes. At each terminal node



the game ends and the payoffs of the players are specified. In the graph below, the tree diagram to the left can represent a game since any two different nodes that are not linked to each other do not share any successor nodes. However, this is not true for the tree diagram to the right where the red nodes are “sharing” the green nodes.

To understand the logic of this representation, consider again the Prisoner’s Dilemma with a twist. Prisoner 1 moves first and chooses whether to confess or lie. From the initial node two branches emerge, one corresponding to Confess and the other to Lie. Each of these branches leads to a new decision node belonging to Prisoner 2, who observes Prisoner 1’s choice and then chooses whether to confess or lie. From each of Prisoner 2’s decision nodes, two additional branches emerge, leading finally to terminal nodes where the payoffs are assigned according to the same payoff matrix we used in the strategic-form version. Therefore, the game tree, as shown below, provides a much more transparent description of the sequential structure of the interaction than the payoff matrix alone.



The same applies to the entry-deterrence game for which the game tree is presented below. Firm 1, the entrant, moves first and chooses whether to Enter or Stay

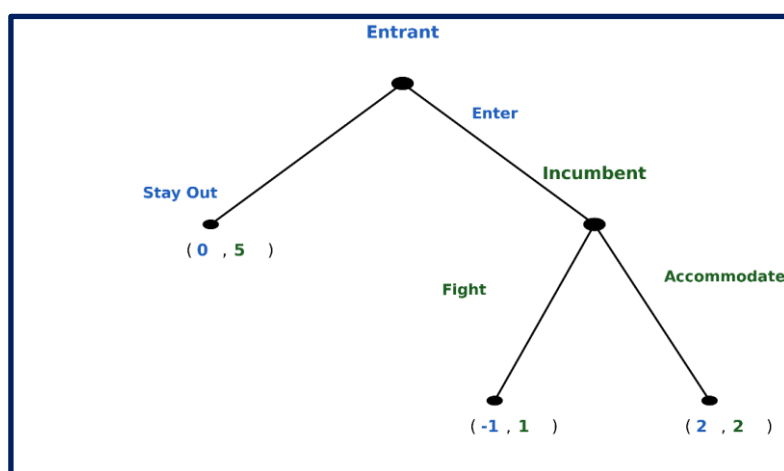
Out. If it stays out, the game ends immediately and the corresponding payoffs are assigned. If it enters, then Firm 2, the incumbent, moves and chooses whether to Fight or Accommodate. Each of these actions leads to a different terminal node with different payoffs. Once this interaction is represented as a game tree, the sequential nature of the problem becomes evident: Firm 2's decision is made only after observing Firm 1's earlier choice. It is precisely this feature that allows us to ask whether Firm 2's threatened response is credible.

The extensive-form representation is therefore valuable not merely because it is visually convenient, but because it makes explicit the elements of the game that are crucial in sequential analysis. In particular, it reveals three aspects that the strategic form tends to conceal.

First, it shows the order of moves. In many games, knowing who moves first and who moves later is essential for understanding the strategic logic of the interaction. A player who moves early must take into account the likely responses of those who follow, while a player who moves later can condition her decision on what has already happened.

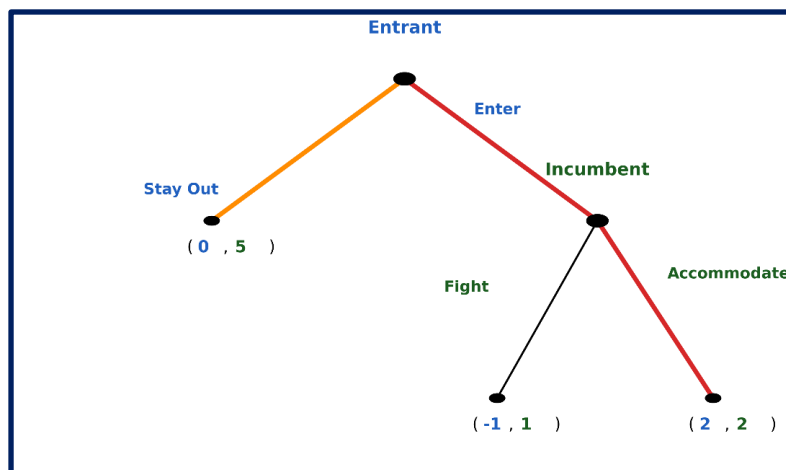
Second, it shows the available actions at each decision point. In a sequential game, a player may face different choices depending on the history of play. The game tree records explicitly what actions are feasible at each node.

Third, it shows the information available to the players when they act. In the games considered in this part, we assume complete and perfect information. Complete information means that all players know the structure of the game, the available actions, and the payoffs associated with every possible outcome. Perfect information means that whenever a player is called upon to move, she observes all previous actions that have been taken. In other words, no player is uncertain about the history of play when making her choice. This distinction was already introduced in your earlier notes, and here it becomes central because the extensive form is designed precisely to capture such information conditions.



A further concept that deserves emphasis at this point is that of a path of play. A path of play is the sequence of actions actually chosen by the players from the initial

node to one of the terminal nodes. For example, in the entry-deterrence game, one possible path of play is the red-colored path in the game-tree below: Firm 1 chooses Enter, then Firm 2 chooses Accommodate. This path leads to the terminal node with payoffs (2,2). Another possible path is the orange-colored path: Firm 1 chooses Stay Out, in which case the game ends immediately with payoffs (0,5). The game tree therefore allows us to distinguish clearly between all possible histories of the game and the final outcome generated by each one of them.



At this stage, one should also note that the extensive-form representation prepares the ground for a more refined discussion of strategies. In the strategic form, it is often natural to think of a strategy as simply the action chosen by a player. In an extensive-form game, however, this is no longer sufficient. A player may have to act at more than one decision node, and her strategy must specify what she would do at each of them. Thus, while an action is a move taken at a particular node, a strategy is a complete contingent plan describing the player's behavior throughout the game. This distinction, already mentioned earlier, becomes much easier to understand once the game is represented as a tree.

The examples and definitions above show that the extensive form is the appropriate framework for studying dynamic strategic interaction. It enables us to describe how the game unfolds, what each player knows when called upon to act, and which future actions are available after every possible history of play. Most importantly, it allows us to analyze whether the behavior prescribed by an equilibrium is credible at every stage of the game. For this reason, the next step is to examine more carefully what a strategy means in an extensive-form game and why it differs from a single action chosen at one point in time.

Actions and strategies in extensive-form games

At this point, we are ready to examine more carefully the distinction between an action and a strategy in an extensive-form game. This distinction is crucial. In the strategic form, it is often natural to identify a player's strategy with the action she chooses, because she typically chooses only once. In an extensive-form game,

however, this is no longer sufficient. A player may be called upon to act at more than one point in the game, and her strategy must specify what she would do at each such point.

An action is a move taken at a particular decision node. For example, in the entry-deterrence game, the entrant has two possible actions at the initial node: Enter or Stay Out. If entry occurs, the incumbent then has two possible actions at her own decision node: Fight or Accommodate. These are actions because they refer to choices made at specific points of the game.

A strategy, by contrast, is a complete contingent plan. It specifies what a player will do at every decision node where she may be called upon to act, even if some of those nodes are never reached along the path of play that actually occurs. In a simple game where each player moves only once, action and strategy coincide. But in games where a player may move after different histories of play, a strategy must describe what that player would do in every such contingency.

To understand this better, consider again the Prisoner's Dilemma with a twist. Prisoner 1 moves first and chooses either Confess or Lie. Prisoner 2 then observes Prisoner 1's action and chooses either Confess or Lie. Prisoner 1 moves only once, so her set of strategies is simply the set of her available actions: Confess or Lie. For Prisoner 2, however, things are different. Prisoner 2 may have to act after Prisoner 1 has confessed or after Prisoner 1 has lied. Therefore, a strategy for Prisoner 2 must specify what he would do in both cases.

This means that Prisoner 2 has four possible strategies:

- confess if Prisoner 1 confessed, and confess if Prisoner 1 lied
- confess if Prisoner 1 confessed, and lie if Prisoner 1 lied
- lie if Prisoner 1 confessed, and confess if Prisoner 1 lied
- lie if Prisoner 1 confessed, and lie if Prisoner 1 lied

Only one of these contingent actions will actually be observed in play, because only one of Prisoner 1's choices will occur. Nevertheless, Prisoner 2's strategy must specify what he would do after each possible action of Prisoner 1. This is exactly what makes the concept of strategy in extensive-form games richer than in the strategic form.

The same logic applies to all extensive-form games. **Whenever a player may face different decision nodes depending on how the game has unfolded, her strategy must specify an action for every such node.** This is important because equilibrium concepts are defined over strategies, not merely over the actions that happen to be taken along the realized path of play. Therefore, when analyzing an extensive-form game, we must always distinguish between the actual sequence of actions observed in equilibrium and the full contingent strategies chosen by the players.

This distinction also helps explain why the Nash equilibrium concept can be too weak in dynamic settings. A strategy profile may be a Nash equilibrium even if some actions prescribed off the equilibrium path would not actually be optimal if the

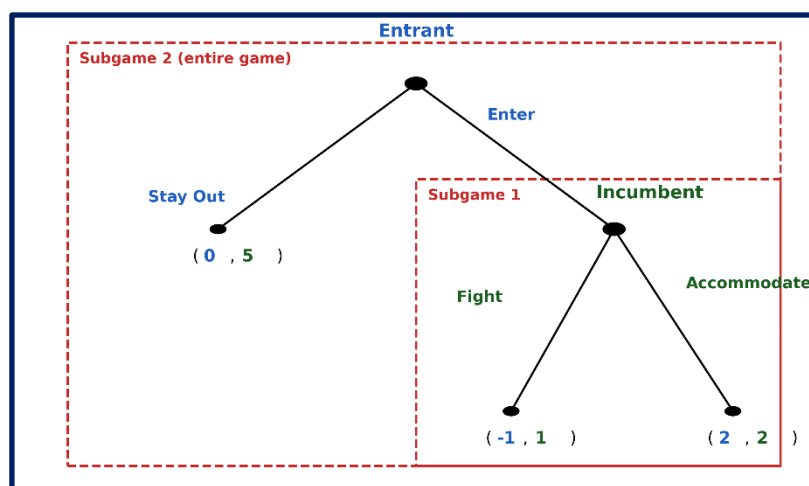
corresponding decision node were reached. To exclude such cases, we need a stronger solution concept. Before introducing it, however, we must first define one more important object: the subgame.

Subgames

The concept of a subgame plays a central role in the analysis of extensive-form games. Intuitively, a subgame is a smaller game contained within the larger game. It begins at a particular decision node and includes that node together with all the nodes, branches, and terminal payoffs that follow from it.

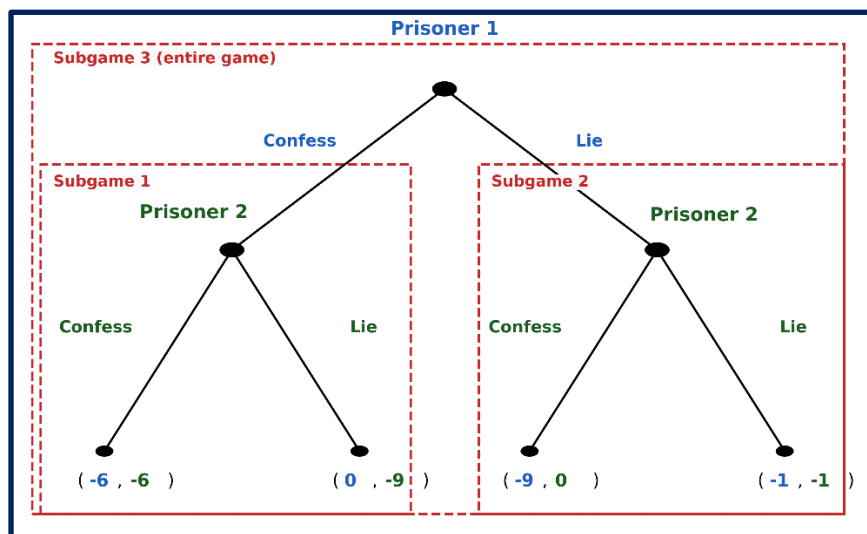
In games of complete and perfect information, the idea is relatively simple. Since every player observes all previous moves before acting, every decision node can potentially mark the beginning of a smaller game. Once a particular history of play has occurred and the corresponding node has been reached, the remainder of the game can be analyzed as a game in its own right. This remaining part of the game is a subgame.

Consider again the entry-deterrence game. The whole game begins when the entrant chooses whether to Enter or Stay Out. If the entrant chooses Enter, the incumbent then chooses whether to Fight or Accommodate. The part of the game that begins after entry has occurred is itself a subgame (i.e., Subgame 1 in the game tree below). At that point, the entrant's earlier decision is already fixed, and only the incumbent remains to move. Thus, the incumbent's decision problem after entry can be viewed as a smaller game embedded in the original one. At the same time, the entire game is itself a subgame (i.e., Subgame 2 in the game tree below).



Similarly, in the Prisoner's Dilemma with a twist, the part of the game beginning after Prisoner 1 has chosen Confess is a subgame (i.e., Subgame 1 in the following game tree), and so is the part beginning after Prisoner 1 has chosen Lie (i.e., Subgame 2 in the following game tree). In each case, once Prisoner 1's action is known, the remainder of the interaction consists of Prisoner 2 choosing between Confess and Lie,

followed by the realization of payoffs. At the same time, the entire game is itself a subgame (i.e., Subgame 3 in the game tree below).



The importance of subgames lies in the following idea: if an equilibrium is to provide a convincing prediction of rational behavior in a sequential game, then the prescribed strategies should be optimal not only for the game as a whole, but also in every subgame. In other words, once any point of the game is reached, the remaining strategy choices should still form a sensible equilibrium of the continuation game. This is the basic intuition behind the refinement of Nash equilibrium that we now introduce.

Subgame Perfect Nash Equilibrium

As we have seen, the ordinary Nash equilibrium concept may be too weak in sequential games because it allows strategy profiles that are supported by non-credible threats or promises. What we need is a stronger equilibrium notion that requires rational behavior at every stage of the game. This leads to the concept of Subgame Perfect Nash Equilibrium.

Definition:

A Subgame Perfect Nash Equilibrium (SPNE) is a strategy profile that induces a Nash equilibrium in every subgame of the game.

This definition captures the idea of **sequential rationality** in a precise way. A strategy profile is not enough merely because no player wishes to deviate at the beginning of the game. It must also be the case that, after every possible history of play, the continuation strategies remain mutually optimal. Thus, SPNE rules out equilibria that rely on non-credible threats.

Every SPNE is a Nash equilibrium, because if a strategy profile is a Nash equilibrium in every subgame, then in particular it must be a Nash equilibrium in the

game as a whole. The converse, however, is not true. A Nash equilibrium need not be subgame perfect, because it may fail to prescribe equilibrium behavior in one or more proper subgames.

The entry-deterrence game illustrates this point very clearly. Recall that when the game is represented in strategic form, there are two Nash equilibria: {Enter, Accommodate} and {Stay Out, Fight}. The second of these is not subgame perfect. To see why, focus on the subgame that begins after the entrant has chosen Enter. In that subgame, the incumbent must choose between Fight and Accommodate. Since accommodation yields a payoff of 2 while fighting yields a payoff of 1, the incumbent prefers Accommodate. Therefore, the strategy Fight does not constitute a Nash equilibrium in that subgame. It follows that the strategy profile {Stay Out, Fight} is not subgame perfect.

By contrast, the strategy profile {Enter, Accommodate} is a Nash equilibrium both in the whole game and in the subgame following entry. It is therefore a Subgame Perfect Nash Equilibrium.

This example shows that SPNE refines the set of Nash equilibria by eliminating those that depend on non-credible behavior. For this reason, SPNE is the standard solution concept for extensive-form games of complete and perfect information.

A more formal definition of SPNE

To fully define an extensive form game the followings must be defined

- $\mathcal{N} = \{1, 2, \dots, n\}$ players
- $X = \{v_0, v_1, \dots, v_m\}$ is the set of nodes and it is partitioned in $n + 1$ non-joint sets
- If $v_l \in X$ but $v_l \notin X_i$ for all $i \in \mathcal{N}$, then v_l is a terminal node (*i.e.*, a node that it does not belong to any player)
- v_0 is the starting node and may or may be not belong to a player
- X_i the set of nodes belonging to player $i \in \mathcal{N}$, where $X_i \subset X$
- $X(v_l)$ is the set of v_l and its successors
- Starting from any node a subgame consists of $X(v_l)$
- For a non-terminal node v_l that belongs to a player i , $A_{v_l} = \{a_{v_l}^1, a_{v_l}^2, \dots, a_{v_l}^k\}$ is the set of actions of player i at the node $v_l \in X_i$
- $S_i = \{s_i^1, s_i^2, \dots, s_i^m\}$ is the set of strategies of player i
- A strategy of player is a function $s_i: X_i \rightarrow A_i$ such that $s_i(v_l) \in A_{v_l}$ (*i.e.*, the strategy assigns an action of player i to every node that belongs to her).
- $u_i: S \rightarrow \mathcal{R}$, where $S = \times_{i \in \mathcal{N}} S_i$ is the payoff function of player i

Definition: Let $s^* = (s_1^*, s_2^*, \dots, s_n^*) \in S$ be a strategy profile of an extensive-form game, where $S = \times_{i \in \mathcal{N}} S_i$. For any node $v_\ell \in X$, let the subgame starting at v_ℓ consist of $X(v_\ell)$. Then s^* is a **Subgame Perfect Nash Equilibrium (SPNE)** if, for every node $v_\ell \in X$, the restriction of s^* to the subgame $X(v_\ell)$ is a Nash equilibrium of that subgame. Equivalently, s^* is a SPNE $\Leftrightarrow \forall v_\ell \in X, s^*|_{X(v_\ell)} \in NE(X(v_\ell))$.

Backward induction

How do we find a Subgame Perfect Nash Equilibrium in a finite extensive-form game of complete and perfect information? The standard method is backward induction.

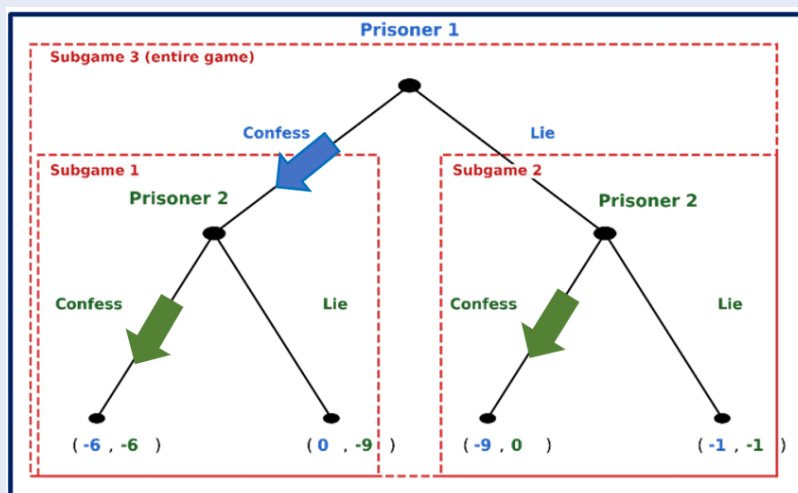
Backward induction is a procedure that starts from the final stages of the game and works backward toward the initial node. The logic is straightforward. At the last stage of the game, the player who moves there chooses the action that maximizes her payoff. Once that optimal continuation is identified, the player who moves one stage earlier anticipates it and chooses the action that is best for her given the future response. This reasoning continues backward until the initial node is reached.

Thus, backward induction solves the game by replacing future decision problems with their optimal outcomes. The resulting path of play is the one generated by sequentially rational behavior at every stage.

It is important to understand the relationship between backward induction and SPNE. Subgame perfection is the equilibrium concept; backward induction is the method used to find it in finite games of complete and perfect information. In such games, the outcome generated by backward induction is a Subgame Perfect Nash Equilibrium outcome.

Example 20:

Let us apply this logic to the Prisoner's Dilemma with a twist with the help of the game tree below. At the second stage, Prisoner 2 chooses after observing Prisoner 1's move. If Prisoner 1 has confessed (i.e., in Subgame 1), Prisoner 2 compares confessing, which yields -6, with lying, which yields -9. Therefore, Prisoner 2 will confess (hence, the green arrow on the edge representing the choice Confess). If Prisoner 1 has lied (i.e., in Subgame 2), Prisoner 2 compares confessing, which yields 0, with lying, which yields -1. Again, Prisoner 2 will confess (hence, the green arrow on the edge representing the choice Confess). Thus, at every terminal decision node, Prisoner 2's optimal action is Confess.



Prisoner 1 anticipates this. Therefore, in Subgame 3, if Prisoner 1 confesses, the final outcome will be {Confess, Confess}, giving Prisoner 1 a payoff of -6. If Prisoner 1 lies, Prisoner 2 will still confess, so the final outcome will be {Lie, Confess}, giving Prisoner 1 a payoff of -9. Since -6 is better than -9, Prisoner 1 chooses Confess (hence, the blue arrow representing the choice Confess). Hence, the backward-induction outcome is {Confess, Confess}, exactly as we found earlier.

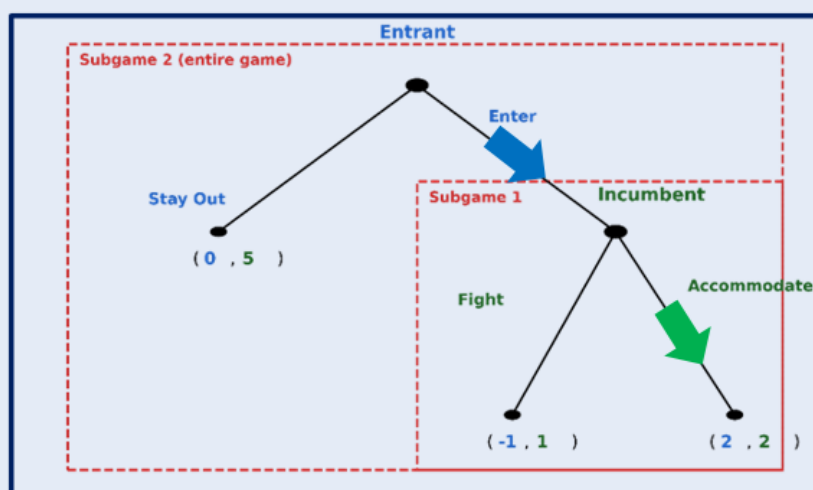
One important point must be emphasized here: the backward-induction outcome {Confess, Confess} is not, strictly speaking, the full solution of the game. When we describe how a sequential strategic interaction will unfold, we must be precise by reporting the equilibrium strategy profile, i.e., a strategy for every player. Remember that, in an extensive-form game, a strategy of a player must specify as many actions as the number of nodes the player must make a choice. Therefore, the solution must specify not only the actions that are actually observed along the equilibrium path, but also the actions each player would choose at every possible contingency. In the present case, it is therefore not enough to say simply that both prisoners confess. We must state that Prisoner 1 chooses Confess in the first stage, and that Prisoner 2 chooses Confess after **either** possible action of Prisoner 1. In other words, Prisoner 2 confesses if Prisoner 1 confessed, and also confesses if Prisoner 1 lied. More formally, the Subgame Perfect Nash Equilibrium of the game is

$$SPNE = \{C_1; C_2 \text{ if } C_1 \text{ and } C_2 \text{ if } L_1\}$$

This example shows why backward induction is such a useful tool. It translates the idea of sequential rationality into a simple step-by-step procedure. By starting from the end and moving backward, we ensure that every action prescribed by the solution is credible at the moment it must actually be chosen.

Example 21:

Now consider the entry-deterrence game. At the second stage, if entry has occurred (i.e., in Subgame 1), the incumbent compares Fight, which yields 1, with Accommodate, which yields 2. Therefore, the incumbent will choose Accommodate



(hence the green arrow on the edge representing the choice Accommodate). Anticipating this, the entrant compares Stay Out, which yields 0, with Enter, which will now be followed by Accommodate, yielding 2. Therefore, the entrant chooses Enter (hence, the blue arrow on the edge representing the choice Enter). The backward-induction outcome is thus {Enter, Accommodate}. In this case the SPNE looks very similar to the backward-induction outcome

$$SPNE = \{\text{Enter ; Accommodate if firm 1 enters}\}$$

Worked example: Battle of the Sexes with a twist

Let us now consider one more example that illustrates how the sequential structure of the game can sharpen the prediction relative to the simultaneous-move version. Recall the Battle of the Sexes game from Part I. Mary and Peter prefer to be together rather than apart, but Mary prefers the Fight while Peter prefers the Opera. In the simultaneous version, the game has two pure-strategy Nash equilibria, one where both go to the Fight and one where both go to the Opera. The Nash equilibrium concept alone does not tell us which one of the two coordinated outcomes will occur.

Now suppose that the game is played sequentially. In the first stage, Mary chooses whether to go to the Fight or to the Opera. In the second stage, Peter observes Mary's choice and then decides where to go.

To solve this game, we apply backward induction. Suppose first that Mary chooses Fight. Peter then compares going to the Fight, where he receives the payoff associated with being together at Mary's preferred event, with going alone to the Opera, where he receives 0. Since being together is better than being apart, Peter will choose Fight. Now suppose instead that Mary chooses Opera. Peter then compares going to the Opera, where he is together with Mary, with going alone to the Fight. Again, he prefers to be together, so he will choose Opera.

Thus, Peter's optimal continuation strategy is simple: after Fight, choose Fight; after Opera, choose Opera. Mary anticipates this. She therefore knows that if she chooses Fight, Peter will follow and the outcome will be the coordinated outcome Mary prefers. If she chooses Opera, Peter will again follow, but this leads to the coordinated outcome Peter prefers. Since Mary prefers the first of these two coordinated outcomes, she chooses Fight.

Therefore, the backward-induction outcome of the sequential Battle of the Sexes is the coordinated outcome in which both players go to the Fight. This example shows that once timing is made explicit, the game may produce a much sharper prediction than in the simultaneous-move version. The extensive form reveals a first-mover advantage for Mary: by moving first and being observed, she can induce the coordinated outcome she prefers.

This example also reinforces the broader lesson of Part II. In sequential games, the timing of moves and the observability of earlier actions can be decisive for

equilibrium selection. The extensive-form representation allows us to capture these features explicitly, and backward induction allows us to solve the game in a way that respects sequential rationality.

Concluding remarks on extensive-form games of complete and perfect information

The analysis in this part has shown why the extensive form is the appropriate framework for studying sequential strategic interaction. Unlike the strategic form, it makes explicit the order of moves, the information available to players when they act, and the continuation structure of the game after every possible history of play.

This additional structure matters because the Nash equilibrium concept, while fundamental, may be too weak in dynamic settings. In some sequential games, such as the Prisoner's Dilemma with a twist, Nash equilibrium continues to provide the correct prediction. In other cases, such as the entry-deterrence game, it may allow equilibria supported by non-credible threats. To rule out such outcomes, we use the concept of Subgame Perfect Nash Equilibrium, which requires equilibrium behavior in every subgame.

In finite extensive-form games of complete and perfect information, Subgame Perfect Nash Equilibria can be found by backward induction. This method captures the basic logic of sequential rationality: players acting earlier anticipate the optimal responses of those acting later and choose accordingly. As a result, the equilibrium outcome is built from credible choices at every stage of the game.