IA168 Algorithmic Game Theory

Štěpán Zapadlo

2024-03-02

Table of contents

Preface

This book follows the course IA168 Algorithmic Game Theory taught by assoc. prof. Brázdil on FI MUNI as it appeared in Fall 2023. All of this course is almost 1-to-1 based on his materials and I only typesetted it into a (in my opinion) more readable format.

In recent years, huge amount of research has been done at the borderline between game theory and computer science, largely motivated by the emergence of the Internet. The aim of the course is to provide students with basic knowledge of fundamental game theoretic notions and results relevant to applications in computer science. The course will cover classical topics, such as general equilibrium theory and mechanism design, together with modern applications to network routing, scheduling, online auctions etc. We will mostly concentrate on computational aspects of game theory such as complexity of computing equilibria and connections with machine learning.

Some notes:

- more mathematically oriented than typically FI course
- lectures should be enough to pass the course

All slides will be available.

Evaluation

Homework is mandatory and there are 3 homework assignments. Also there is a threshold for *pass/not pass* of the homework.

There are public homeworks from last years

Afterwards follows an oral exam:

- there will be at least 3 or 4 exam dates every week
  - there are as many tries as we want
- we need to know **everything** (and it's not a joke :))
- at least 1h for the exam (but very possibly more)
- precise mathematical communication needed for the exam

Overview of the course

This is a *theoretical* course aimed at some fundamental results of game theory, often related to computer science

- We start with strategic form games (such as the Prisoner's dilemma), investigate several solution concepts (dominance, equilibria) and related algorithms.
- Then we consider repeated games which allow players to learn from history and/or to react to deviations of the other players.
- Subsequently, we move on to incomplete information games and auctions.
- Finally, we consider (in)efficiency of equilibria (such as the Price of Anarchy) and its properties on important classes of routing and network formation games.
- Remaining time will be devoted to selected topics from extensive form games, games on graphs etc.

1 Introduction & Outline

1.1 What is game theory?

> One of the possible definitions of Game Theory
>
> *The study of mathematical models of conflict and cooperation between intelligent rational decision-makers*

It is effectively a multi-objective optimization, where multiple loss functions need to be optimized. The algorithmic part means that there will be algorithms for finding concepts discussed.

1.2 Prisoner's dilemma

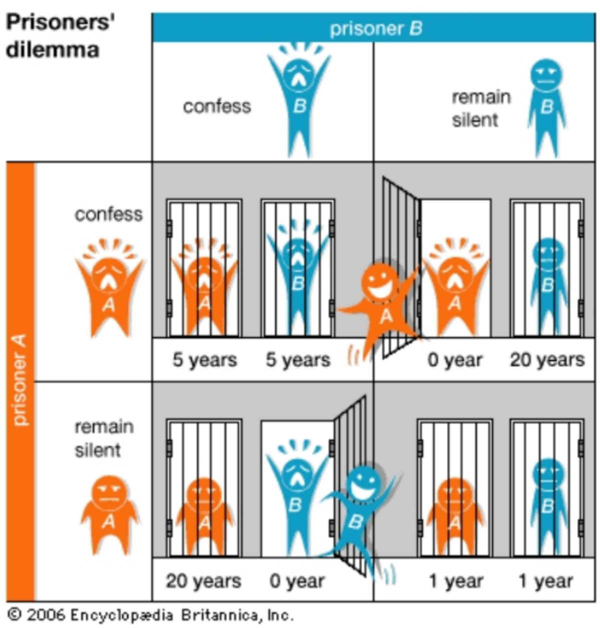

Two suspects of a serious crime are arrested and imprisoned. Police have enough evidence of only petty theft, and to nail the suspects for the serious crime they need testimony from at least one of them. The suspects are interrogated separately without any possibility of communication Each of the suspects is offered a deal:

- If he confesses (C) to the crime, he is free to go.
- The alternative is not to confess, that is remain silent (S)

One prisoner is said to be a *row player* and the other is a *column player* and it following table there are payoffs row the row and column players.

	C	S
C	$(-5, -5)$	$(0, -20)$
S	$(-20, 0)$	$(-1, -1)$

Rational "row" suspect (or his adviser) may reason as follows:

- If my colleague chooses C, then playing C gives me -5 and playing S gives -20.
- If my colleague chooses S, then playing C gives me 0 and playing S gives -1

In both cases, C is clearly better (it *strictly dominates* the other strategy). If the other suspect's reasoning is the same, both choose C and get a 5-year sentence

There is a solution (S, S) which is better for both players but needs some "central" authority to control the players

This suboptimality will be visible throughout the course.

1.3 Nash equilibria

1.3.1 Battle of Sexes

A couple agreed to meet this evening, but cannot recall if they will be attending the opera or a football match. One of them wants to go to the football game. The other one to the opera. Both would prefer to go to the same place rather than different ones.

Battle of Sexes can be modeled as a game of two players (the couple) with the following payoffs.

	O	F
O	$(2, 1)$	$(0, 0)$
F	$(0, 0)$	$(1, 2)$

Apparently, no strategy of any player is dominant. So what could be a "solution"? Note that whenever both players play O, then neither of them wants to unilaterally deviate from his strategy! Here, (O, O) is an example of a *Nash equilibrium* (as is (F, F)).

> **Interpretation**
>
> This can be interpreted as advice to both players that they would have to deviate from.

The thinking process is the same as in physics. First, we make a model of a social situation then make a prediction and validate it.

1.3.2 Rock, Paper, Scissors

Again, the existence of a Nash equilibrium is not guaranteed - as can be seen in *Rock, Paper, Scissors*

	R	P	S
R	$(0,0)$	$(-1,1)$	$(1,-1)$
P	$(1,-1)$	$(0,0)$	$(-1,1)$
S	$(-1,1)$	$(1,-1)$	$(0,0)$

This game is a **zero-sum** game - whatever one player wins, the other one loses.

We can further generalize the players' behavior of the game into *mixed strategies*: Each player plays each pure strategy with probability $\frac{1}{3}$. The expected payoff of each player is 0 (even if one of the players changes his strategy, he still gets 0!).

It is always assumed that each player tries to maximize their payoff (they're playing rationally).

1.4 Dynamic Games

So far we have seen games in *strategic form* that are unable to capture games that unfold over time (such as chess).

For such purpose, we need to use *extensive* form games. So it can be said that they are *iterated* or *repeated* games.

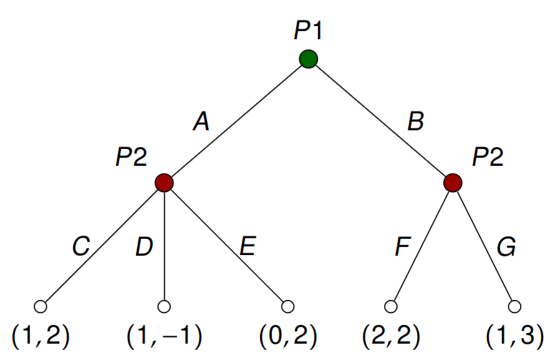

In this tree a player *P1* chooses an action A or B. Then the player *P2* follows by choosing from actions C, D, E, F, G.

1.4.1 Imperfect information

Some decisions in the game tree may happen by chance and are controlled by neither player (e.g. Poker, Backgammon, etc.)

Sometimes a player may not be able to distinguish between several "positions" because he does not know all the information in them (Think a card game with the opponent's cards hidden). For this purpose, we will introduce so-called *information sets* (and these games are *imperfect information* games).

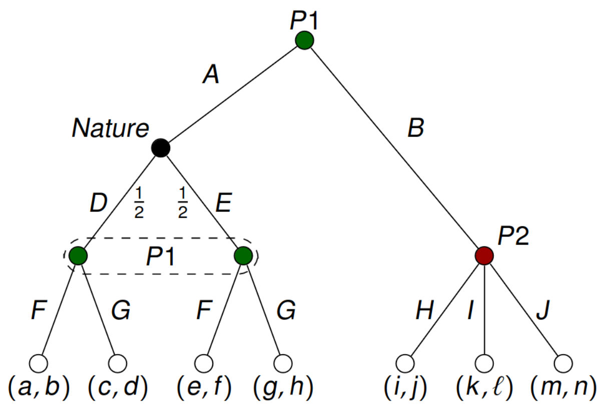

1.4.2 Incomplete information

In all previous games, the players knew all the details of the game they played, and this fact was a "common knowledge". This is not always the case.

1.4.2.1 Sealed Bid Auction

Two bidders are trying to purchase the same item.

- The bidders simultaneously submit bids b_1 and b_2 and the item is sold to the highest bidder at his bid price (first price auction)
- The payoff of the player 1 (and similarly for player 2) is calculated by

$$u_1(b_1, b_2) = \begin{cases} v_1 - b_1 & b_1 > b_2 \\ \frac{1}{2}(v_1 - b_1) & b_1 = b_2 \\ 0 & b_1 < b_2 \end{cases}$$

  Here v_1 is the private value that player 1 assigns to the item and so player 2 **does not know** u_1.

1.5 Inefficiency of Equilibria

In Prisoner's Dilemma, the selfish behavior of suspects (the Nash equilibrium) results in a somewhat worse-than-ideal situation

Defining a *welfare function* W which to every pair of strategies assigns the sum of payoffs, we get $W(C, C) = -10$ but $W(S, S) = -2$.

> *Such choice of a welfare function is generally arbitrary, but results vary depending on this choice.*

The ratio $\frac{W(C,C)}{W(S,S)} = 5$ measures the inefficiency of "selfish-behavior" (C, C) w.r.t. the optimal "centralized" solution. *Price of Anarchy* is the maximum ratio between values of equilibria and the value of an optimal solution.

1.5.1 Selfish routing

Consider a transportation system where many agents are trying to get from some initial location to a destination. Consider the welfare to be the average time for an agent to reach the destination.

There are two versions:

- "Centralized": A central authority tells each agent where to go.
- "Decentralized": Each agent selfishly minimizes his travel time.

Price of Anarchy measures the ratio between average travel time in these two cases.

Problem: Bound the price of anarchy over all possible routing games?

1.6 Games in Computer Science

Game theory is a core foundation of mathematical economics. But what does it have to do with CS?

- Games in AI: modeling of "rational" agents and their interactions.
- Games in machine learning: Generative adversarial networks (GANs), reinforcement learning
- Games in Algorithms: several game theoretic problems have a very interesting algorithmic status and are solved by interesting algorithms
- Games in modeling and analysis of reactive systems: program inputs viewed "adversarially", bisimulation games (the tester is trying to kill the program by weird inputs), etc.
- Games in computational complexity: Many complexity classes are definable in terms of games: PSPACE, polynomial hierarchy, etc.
- Games in Logic: modal and temporal logics, Ehrenfeucht-Fraisse games, etc.

2 Complete-Information Static Games

2.1 Intuition

Proceed in two steps:

1. Players *simultaneously and independently* choose their strategies. This means that players play without observing **strategies** chosen by other players.
2. Conditional on the players' strategies, **payoffs** are distributed to all players.

Complete information means that the following is *common knowledge* among players:

- all possible strategies of all players,
- what payoff is assigned to each combination of strategies.

Definition 2.1 (Common knowledge)**.** A fact E is a *common knowledge* among players $\{1, \dots, n\}$ if for every sequence $i_1, \dots, i_k \in \{1, \dots, n\}$ we have that i_1 knows that i_2 knows that ... i_{k-1} knows that i_k knows E.

The goal of each player is to maximize his payoff (and this fact is a common knowledge).

2.2 Strategic-Form games

To formally represent static games of complete information we define *strategic-form* games

Definition 2.2 (Strategic-form games)**.** A game in *strategic-form* (or normal-form) is an ordered triple $G = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$, in which:

- $N = \{1, \dots, n\}$ is a finite set of players;
- S_i is a set of (pure) strategies of player i, for every $i \in N$.

  A *strategy profile* $\boldsymbol{s}$ is a vector of strategies of all players $\boldsymbol{s} = (s_1, \dots, s_n) \in S_1 \times \cdots \times S_n$.

  We denote the set of all strategy profiles by $S = S_1 \times \cdots \times S_n$;
- $u_i : S \to \mathbb{R}$ is a function associating each strategy profile $\boldsymbol{s} = (s_1, \dots, s_n) \in S$ with the payoff $u_i(\boldsymbol{s})$ to player i, for every player.

Definition 2.3 (Zero-sum games)**.** A *zero-sum* game G is one in which for all $\boldsymbol{s} = (s_1, \ldots, s_n) \in S$ we have $\sum_{i \in N} u_i(\boldsymbol{s}) = 0$

Two examples are provided in the slides Prisoners' dilemma and Cournot Doupoly.

2.3 Solution Concepts

A *solution concept* is a method of analyzing games with the objective of restricting the set of *all possible outcomes* to those that are more reasonable than others. We will use the term *equilibrium* for any one of the strategy profiles that emerge as one of the solution concepts' predictions.

We follow the approach of Steven Tadelis here, even though it is not completely standard.

> Tip
>
> *Nash equilibrium* is a solution concept. That is, we "solve" games by finding Nash equilibria and declaring them to be reasonable outcomes.

2.4 Assumptions

Throughout the lecture, we assume that:

1. Players are **rational**: a *rational* player is one who chooses his strategy to maximize his payoff.
2. Players are **intelligent**: An *intelligent* player knows everything about the game (actions and payoffs) and can make any inferences about the situation that we can make.
3. **Common knowledge**: The fact that players are rational and intelligent is common knowledge among them.
4. **Self-enforcement**: Any prediction (or equilibrium) of a solution concept must be *self-enforcing*.

Here 4th assumption implies non-cooperative game theory: Each player is in control of his actions, and he will stick to an action only if he finds it to be in his best interest

2.5 Evaluating Solution Concepts

In order to evaluate our theory as a methodological tool we use the following criteria:

1. **Existence** (i.e., how often does it apply?): The solution concept should apply to a wide variety of games.
    - E.g. We shall see that mixed Nash equilibria exist in all two-player finite strategic-form games.

2. **Uniqueness** (How much does it restrict behavior?): We demand our solution concept to restrict the behavior as much as possible.
    - E.g. So-called strictly dominant strategy equilibria are always unique as opposed to Nash eq.

2.5.1 Pure Strategies

We will consider the following solution concepts:

- strict dominant strategy equilibrium;
- iterated elimination of strictly dominated strategies (IESDS);
- rationalizability;
- Nash equilibria.

> **Note**
>
> These concepts will be first discussed in terms of pure strategies only at first.

3 Solution concepts

3.1 Domination of strategies

Let $N = \{1, \ldots, n\}$ be a finite set and for each $i \in N$ let X_i be a set. Moreover, let $X := \prod_{i \in N} X_i = \left\{(x_1, \ldots, x_n) | x_j \in X_j,\ j \in N\right\}$. We shall now introduce the following notation:

- for $i \in N$ we define $X_{-i} := \prod_{j \neq i} X_j$, i.e.,

$$X_{-i} = \left\{(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n) | x_j \in X_j, \forall j \neq i\right\};$$

- an element of X_{-i} will be denoted by

$$\boldsymbol{x}_{-i} = (x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n);$$

  as a slight abuse of notation we write $(x_i, \boldsymbol{x}_{-i})$ to denote $(x_1, \ldots, x_i, \ldots, x_n) \in X$.

Definition 3.1 (Strictly dominated strategies)**.** Let $s_i, s_i' \in S_i$ be strategies of player i as per Definition 2.2. Then s_i' is **strictly dominated** by s_i (write $s_i \succ s_i'$) if for any possible combination of the other players' strategies, $\boldsymbol{s}_{-i} \in S_{-i}$, we have

$$u_i(s_i, \boldsymbol{s}_{-i}) > u_i(s_i', \boldsymbol{s}_{-i})$$

for all $\boldsymbol{s}_{-i} \in S_{-i}$.

Conjecture 3.1. *An intelligent and rational player will never play a strictly dominated strategy.*

Proof. Clearly, intelligence implies that the player should recognize dominated strategies, and rationality implies that the player will avoid playing them. □

Definition 3.2 (Strictly dominant strategy)**.** We say $s_i \in S_i$ is *strictly dominant* if every other pure strategy of player i is strictly dominated by s_i. Similarly, we say $s_i \in S_i$ is *strictly dominated in game G* if there exists a pure strategy $s_i' \in S_i$ such that $s_i' \succ s_i$.

Observe that every player has *at most* one strictly dominant strategy and that strictly dominant strategies do **not** have to exist. Furthermore, we can make a similar claim to Conjecture 3.1:

Conjecture 3.2. *Any rational player will play the strictly dominant strategy if it exists.*

Definition 3.3 (Strictly dominant strategy equilibrium)**.** A strategy profile $\boldsymbol{s} \in S$ is a *strictly dominant strategy equilibrium* if s_i is strictly dominant for all $i \in N$.

Corollary 3.1. *If the strictly dominant strategy equilibrium exists, it is unique and rational players will play it.*

3.1.1 Examples

Example 3.1 (Prisoner's dilemma)**.** In the Prisoner's dilemma, (C, C) is the strictly dominant strategy equilibrium.

	C	S
C	$(-5, -5)$	$(0, -20)$
S	$(-20, 0)$	$(-1, -1)$

Example 3.2 (Battle of Sexes)**.** In the Battle of Sexes, no strictly dominant strategies exist.

	O	F
O	$(2, 1)$	$(0, 0)$
F	$(0, 0)$	$(1, 2)$

i Indiana Jones and the Last Crusade

Indiana Jones, his father, and the Nazis have all converged at the site of the Holy Grail. The two Joneses refuse to help the Nazis reach the last step. So the Nazis shoot Indiana's dad. Only the healing power of the Holy Grail can save the senior Dr. Jones from his mortal wound. Suitably motivated, Indiana leads the way to the Holy Grail. But there is one final challenge. He must choose between literally scores of chalices, only one of which is the cup of Christ. While the right cup brings eternal life, the wrong choice is fatal. The Nazi leader impatiently chooses a beautiful gold chalice, drinks the holy water, and dies from the sudden death that follows from the wrong choice. Indiana picks a wooden chalice, the cup of a carpenter. Exclaiming "There's only one way to find out" he dips the chalice into the font and drinks what he hopes is the cup of life. Upon discovering that he has chosen wisely, Indiana brings the cup to his father and the water heals the mortal wound.

In this scene, Indy behaved "suboptimally (irrationally)", because he overlooked his strictly dominant strategy, which would be to give the water to his father without tasting. Here are the possible outcomes:

- If Indiana has chosen the right cup, his father is still saved.
- If Indiana has chosen the wrong cup, then his father dies but Indiana is spared.

Testing the cup before giving it to his father doesn't help, since if Indiana has made the wrong choice, there is no second chance – Indiana dies from the water and his father dies from the wound.

3.2 Iterated Strict Dominance in Pure Strategies

We know that, by Conjecture 3.1, no rational player ever plays strictly dominated strategies. As each player knows that each player is rational, each player knows that his opponents will not play strictly dominated strategies and thus all opponents know that effectively they are facing a "smaller" game. As rationality is common knowledge, everyone knows that everyone knows that the game is effectively smaller. Thus everyone knows, that nobody will play strictly dominated strategies in the smaller game (and such strategies may indeed exist). Because it is common knowledge that all players will perform this kind of reasoning again, the process can continue until no more strictly dominated strategies can be eliminated.

This principle (or reasoning) yields the **Iterated Elimination of Strictly Dominated Strategies**.

Definition 3.4 (Iterated Elimination of Strictly Dominated Strategies)**.** Define a sequence $D_i^0, D_i^1, D_i^2, \dots$ of strategy sets of player i. Also denote by G_{DS}^k the game obtained from G by restricting to D_i^k, $i \in N$. We shall call the following algorithm *"Iterated Elimination of Strictly Dominated Strategies"*:

1. Initialize $k = 0$ and $D_i^0 = S_i$ for each $i \in N$.
2. For all players $i \in N$: Let D_i^{k+1} be the set of all pure strategies of D_i^k that are **not** strictly dominated in G_{DS}^k.
3. If $D_i^{k+1} = D_i^k$ for all players $i \in N$, then stop. Otherwise, let $k := k + 1$ and go to 2.

We say that $s_i \in S_i$ **survives IESDS** if $s_i \in D_i^k$ for all $k = 0, 1, 2, \dots$ (or until stop).

Caution

I modified the algorithms (or definitions) 3.4 and 3.9 to stop when nothing changes across two iterations. In the slides, this condition is **not** present.
This change was motivated to make each of these processes stop at some point when no further iterations are necessary. But just to be sure, maybe use/remember the original/official definitions.

Definition 3.5. A strategy profile $\boldsymbol{s} = (s_1, \dots, s_n) \in S$ is an **IESDS equilibrium** if each s_i survives IESDS. A game is **IESDS solvable** if it has a **unique** IESDS equilibrium.

Remark. If all S_i are finite, then in 2. we may remove only some of the strictly dominated strategies (not necessarily all). The result is not affected by the order of elimination since strictly dominated strategies remain strictly dominated even after removing some other strictly dominated strategies.

3.2.1 Examples

Example 3.3. In the Prisoner's dilemma, the strategy profile (C, C) is the only one surviving the first round of IESDS.

	C	S
C	$(-5, -5)$	$(0, -20)$
S	$(-20, 0)$	$(-1, -1)$

Example 3.4. In the Battle of Sexes, all strategies survive **all rounds** (i.e. IESDS says: "*anything may happen, sorry*")

	O	F
O	$(2,1)$	$(0,0)$
F	$(0,0)$	$(1,2)$

Exercise 3.1. Presume the game is given by the Table 3.5:

Table 3.5: Table defining a game

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

Solution. Surely, from the Definition 3.4 follows

$$D_1^0 = S_1 = \{L, C, R\}, \qquad D_2^0 = S_2 = \{L, C, R\}.$$

Notice that for player 2, playing C is strictly dominated by R. On the other hand, in this iteration, there are no strictly dominated strategies for the first player. Therefore next iteration reads

$$D_1^1 = D_1^0 = \{L, C, R\}, \qquad D_2^1 = \{L, R\}$$

and the table transforms to

$$\begin{bmatrix} (4,3) & (6,2) \\ (2,1) & (3,6) \\ (3,0) & (2,8) \end{bmatrix}.$$

Now the strategy L dominates both C, R for player 1. For player 2, no strategy is strictly dominant. Now the second iteration has the form

$$D_1^2 = \{L\}, \qquad D_2^2 = \{L, R\}$$

with the table

$$\begin{bmatrix} (4,3) & (6,2) \end{bmatrix}.$$

In this case, the R strategy is *strictly dominated* by L for the second player and we get

$$D_1^3 = \{L\}, \qquad D_2^3 = \{L\}$$

Here, the iterations stop and the strategies L are the only one that survives IESDS (for both players). Hence the (L, L) is the IESDS equilibrium of this game.

3.2.2 Political Science Example: Median Voter Theorem

In this example of a two-player game by Hotelling (1929) and Downs (1957), we have

$$N = \{1, 2\}, \qquad S_i = \{1, 2, \dots, 10\},$$

where each player (a candidate) plays his position on a political and ideological spectrum, see Figure 3.1 – 10 voters belong to each position (or in real life, 10% of voters belong to each position). Voters then vote for their closest candidate and if there is a tie, then half the votes go to each candidate. The payoff in this game is the number of voters for the candidate, and each candidate (selfishly) strives to maximize this number.

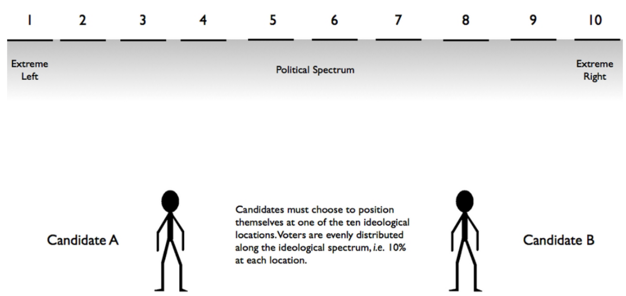

Figure 3.1: Illustration of the spectrum

Here in $G = G^0_{DS}$ surely, 1 and 10 are the (only) strictly dominated strategies $\implies D^1_1 = D^1_2 = \{2, \dots, 9\}$. Then again, in G^1_{DS} the 2 and 9 are the (only dominated strategies) $\implies D^2_1 = D^2_2 = \{3, \dots, 8\}$. If we repeat this, only strategies 5 and 6 emerge not dominated and survive IESDS.

3.3 Belief & Best Response

IESDS eliminated apparently unreasonable behavior (leaving "reasonable" behavior implicitly untouched). What if we rather want to actively preserve reasonable behavior? But a question arises, what is reasonable? The simple answer is what we believe is reasonable :). To build an intuition, consider the following situation

- Imagine that your colleague did something stupid
- What would you ask him? Usually something like "What were you thinking?"
- The colleague may respond with a reasonable description of his *belief* in which his action was (one of) the best he could do

You may, of course, question the reasonableness of the belief

To formalize this kind of reasoning, we start with the next few definitions.

Definition 3.6 (Belief)**.** A *belief* of player i is a pure strategy profile $\mathbf{s}_{-i} \in S_{-i}$ of his opponents.

Definition 3.7 (Best Response)**.** A strategy $s_i \in S_i$ of player i is a *best response* to a belief $\boldsymbol{s}_{-i} \in S_{-i}$ if

$$u_i(s_i, \boldsymbol{s}_{-i}) \geq u_i(s_i', \boldsymbol{s}_{-i})$$

for all $s_i' \in S_i$.

Conjecture 3.3. *A rational player who believes that his opponents will play $\boldsymbol{s}_{-i} \in S_{-i}$ always chooses a best response to $\boldsymbol{s}_{-i} \in S_{-i}$.*

Definition 3.8. A strategy $s_i \in S_i$ is *never best response* if it is not a best response to any belief $\boldsymbol{s}_{-i} \in S_{-i}$.

Clearly, a rational player never plays any strategy that is never best response.

Proposition 3.1. *If s_i is strictly dominated for the player i, then it is never best response.*

The opposite, though, does not have to be true in pure strategies. Consider the game given by Table 3.6.

Table 3.6: Counter-example to the opposite of Proposition 3.1

	X	Y
A	$(1, 1)$	$(1, 1)$
B	$(2, 1)$	$(0, 1)$
C	$(0, 1)$	$(2, 1)$

Here A is never best response but it is not strictly dominated either by B, or by C.

3.3.1 Elimination of Stupid Strategies

Using similar iterated reasoning as for IESDS, see Definition 3.4, strategies that are never best response can be iteratively eliminated.

Definition 3.9 (Rationalizable)**.** Define a sequence $R_i^0, R_i^1, \ldots$ of strategy sets of player i. Also, denote by G_{Rat}^k the game obtained from G by restricting to R_i^k for $i \in N$. Consider the following algorithm

1. Initialize $k = 0$ and $R_i^0 = S_i$ for each $i \in N$.
2. For all players $i \in N$: Let R_i^{k+1} be the set of all strategies of R_i^k that are best responses to some beliefs in G_{Rat}^k.
3. If $R_i^{k+1} = R_i^k$ for all players $i \in N$, then stop. Otherwise, let $k := k + 1$ and go to 2.

We say that $s_i \in S_i$ is *rationalizable* if $s_i \in R_i^k$ for all $k = 0, 1, 2, \ldots$ (or until stop).

Definition 3.10. A strategy profile $\boldsymbol{s} = (s_1, \ldots, s_n) \in S$ is a *rationalizable equilibrium* if each s_i rationalizable.

We say that a game is *solvable by rationalizability* if it has a **unique** rationalizable equilibrium.

> **Warning**
>
> For some reason, rationalizable strategies are almost always defined using mixed strategies!

3.3.2 Examples

Example 3.5. In the Prisoners' dilemma, the strategy profile (C, C) is the only rationalizable equilibrium.

	C	S
C	$(-5, -5)$	$(0, -20)$
S	$(-20, 0)$	$(-1, -1)$

Example 3.6. In the Battle of Sexes, all strategies are rationalizable.

	O	F
O	$(2, 1)$	$(0, 0)$
F	$(0, 0)$	$(1, 2)$

3.3.3 Cournot Doupoly

Consider a game $G = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$ with $N = \{1, 2\}$, $S_i = [0, \infty)$ and payoffs of form

$$u_1(q_1, q_2) = q_1(\kappa - q_1 - q_2) - q_1 c_1 = (\kappa - c_1)q_1 - q_1^2 - q_1 q_2$$
$$u_2(q_1, q_2) = q_2(\kappa - q_1 - q_2) - q_2 c_2 = (\kappa - c_2)q_2 - q_2^2 - q_1 q_2$$

For simplicity, we shall assume $c_1 = c_2 = c$ and denote $\theta = \kappa - c$.

> **Economical background**
>
> Cournot competition is an economic model used to describe an industry structure in which companies compete on the amount of output they will produce, which they decide on independently of each other and at the same time. It is named after Antoine Augustin Cournot (1801–1877) who was inspired by observing competition in a spring water duopoly. It has the following features:
>
> - There is more than one firm and all firms produce a homogeneous product, i.e., there is no product differentiation;
> - Firms do not cooperate, i.e., there is no collusion;
> - Firms have market power, i.e., each firm's output decision affects the good's price;
> - The number of firms is fixed;
> - Firms compete in quantities rather than prices; and
> - The firms are economically rational and act strategically, usually seeking to maximize profit

given their competitors' decisions.

In this context, $q_i \in S_i$ is spent manufacturing time of i-th firm, c_i its fixed costs per one manufactured product and κ scales the variable costs based on the total manufactured quantity.

What is a best response of player 1 to a given q_2?

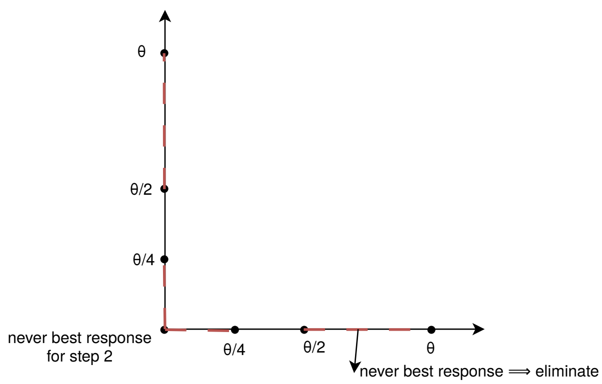

Figure 3.2: Diagram for Cournot douploy

Solve $\frac{\partial u_1}{\partial q_1} = \theta - 2q_1 - q_2 = 0$, which gives that $q_1 = (\theta - q_2)/2$ is the only best response of player 1 to q_2. Similarly, $q_2 = (\theta - q_1)/2$ is the only best response of player 2 to q_1. Since $q_2 \geq 0$ (and therefore $(\theta - q_2)/2 \leq \theta/2$), we obtain that q_1 is never best response iff $q_1 > \theta/2$ and just as well, q_2 is never best response iff $q_2 > \theta/2$. Thus $R_1^1 = R_2^1 = [0, \theta/2]$. Now, in G_{Rat}^1, we still have that $q_1 = (\theta - q_2)/2$ is the best response to q_2, and $q_2 = (\theta - q_1)/2$ the best response to q_1. Since $q_2 \in R_2^1 = [0, \theta/2]$, we obtain that q_1 is never best response iff $q_1 \in [0, \theta/4)$. Similarly q_2 is never best response iff $q_2 \in [0, \theta/4)$. In general, after $2k$ iterations we have $R_1^{2k} = R_2^{2k} = [l_k, r_k]$, where

- $r_k = (\theta - l_{k-1})/2$ for $k \geq 1$
- $l_k = (\theta - r_k)/2$ for $k \geq 1$ and $l_0 = 0$

Solving the recurrence we obtain

- $l_k = \theta/3 - \left(\frac{1}{4}\right)^k \theta/3$
- $r_k = \theta/3 + \left(\frac{1}{4}\right)^{k-1} \theta/6$

Hence, $\lim_{k\to\infty} l_k = \lim_{k\to\infty} r_k = \theta/3$ and thus $(\theta/3, \theta/3)$ is the only rationalizable equilibrium. But does this mean, that $q_i = \theta/3$ provide the best outcomes possible?

$$u_1(\theta/3, \theta/3) = u_2(\theta/3, \theta/3) = \theta^2/9$$

No, when we consider $q_i = \theta/4$, we get

$$u_1(\theta/4, \theta/4) = u_2(\theta/4, \theta/4) = \theta^2/8,$$

but this would require cooperation between players (firms).

3.4 IESDS vs Rationalizability in Pure Strategies

Theorem 3.1. *Assume that S is finite. Then for all k we have that $R_i^k \subseteq D_i^k$ (from Definition 3.4 and Definition 3.9). That is, in particular, all rationalizable strategies survive IESDS.*

The opposite inclusion does not have to hold in pure strategies. As a counter-example consider a game given by a table

Table 3.9: Counter-example to the opposite of Theorem 3.1, see also Table 3.6

	X	Y
A	$(1,1)$	$(1,1)$
B	$(2,1)$	$(0,1)$
C	$(0,1)$	$(2,1)$

Recall, from Table 3.6, that A is never best response but it is not strictly dominated by either B, or C. That is, A survives IESDS but is not rationalizable.

Lemma 3.1. *If s_i is a best response to $\mathbf{s}_{-i}$ in G^k_{Rat}, then s_i is a best response to $\mathbf{s}_{-i}$ in G.*

Proof. Let us prove Lemma 3.1 by induction on k. For $k = 0$ we have $G^k_{Rat} = G^0_{Rat} = G$ and the claim holds trivially.

Now assume that the claim is true for some k and that s_i is a best response to $\mathbf{s}_{-i}$ in G^{k+1}_{Rat}. Let s_i' be a best response to $\mathbf{s}_{-i}$ in G^k_{Rat}. Then $s_i' \in G^{k+1}_{Rat}$, since s_i' is *not* eliminated from G^k_{Rat}. However, since s_i is a best response to $\mathbf{s}_{-i}$ in G^{k+1}_{Rat}, we get $u_i(s_i, \mathbf{s}_{-i}) \geq u_i(s_i', \mathbf{s}_{-i})$. Thus, by Definition 3.7, s_i is a best response to $\mathbf{s}_{-i}$ in G^k_{Rat}.

By the induction hypothesis, s_i is a best response to $\mathbf{s}_{-i}$ in G and the lemma has been proven. □

> **!** Important
>
> The proof of Theorem 3.1 is not necessary for the final exam!

Proof. Let us now focus on Theorem 3.1 and prove $R_i^k \subseteq D_i^k$ for all players i by induction on k.

For $k = 0$ we have $R_i^0 = S_i = D_i^0$ by definition (see Definition 3.4 and Definition 3.9).

Assume that $R_i^k \subseteq D_i^k$ for some $k \geq 0$ and prove that $R_i^{k+1} \subseteq D_i^{k+1}$. Let $s_i \in R_i^{k+1}$, then there must be $\boldsymbol{s}_{-i} \in R_{-i}^k$ such that s_i **is a best response to $\boldsymbol{s}_{-i}$ in** G_{Rat}^k, as s_i must have **not been** eliminated in G_{Rat}^k. By the Lemma 3.1, s_i is a best response to $\boldsymbol{s}_{-i}$ in G as well. Also, by the induction hypothesis, $s_i \in R_i^{k+1} \subseteq R_i^k \subseteq D_i^k$ and $\boldsymbol{s}_{-i} \in R_{-i}^k \subseteq D_{-i}^k$.

However, then s_i is a best response to $\boldsymbol{s}_{-i}$ in G_{DS}^k, as this follows from the fact, that the "best response" relationship of s_i and $\boldsymbol{s}_{-i}$ is preserved when removing arbitrarily many other strategies. Thus s_i **is not strictly dominated in** G_{DS}^k and $s_i \in D_i^{k+1}$.

This concludes the proof. □

3.5 Nash equilibria

We may raise certain criticisms of previous approaches:

- Strictly dominant strategy equilibria often do not exist;
- IESDS and rationalizability may not remove any strategies.

A typical example is the Battle of Sexes:

Table 3.10: Battle of Sexes

	O	F
O	$(2,1)$	$(0,0)$
F	$(0,0)$	$(1,2)$

Here all strategies are equally reasonable according to the above concepts. But are all strategy profiles really equally reasonable? Assume that each player has a belief about the strategies of other players. By Conjecture 3.3, each player plays a best response to his beliefs, so is (O, F) as reasonable as (O, O) in this respect? Note that if player 1 believes that player 2 plays O, then playing O is reasonable, and if player 2 believes that player 1 plays F, then playing F is reasonable. But such beliefs cannot be correct together! Hence (O, O) can be obtained as a profile where each player plays the best response to his belief and the **beliefs are correct**.

Nash equilibrium can be defined as a set of beliefs (one for each player) and a strategy profile in which every player plays a best response to his belief and each strategy of each player is consistent with the beliefs of his opponents.

Definition 3.11 (Nash Equilibrium)**.** A pure-strategy profile $\mathbf{s}^* = (s_1^*, \dots, s_n^*) \in S$ is a (pure) **Nash equilibrium** if s_i^* is a best response to $\mathbf{s}_{-i}^*$ for each $i \in N$, that is

$$u_i(s_i^*, \mathbf{s}_{-i}^*) \geq u_i(s_i, \mathbf{s}_{-i}^*)$$

for all $s_i \in S_i$ and all $i \in N$.

> **Note**
>
> Note that this definition is equivalent to the previous one in the sense that $\mathbf{s}_{-i}^*$ may be considered as the (consistent) belief of player i to which he plays a best response s_i^*.

Example 3.7. In the Prisoner's dilemma, (C, C) is the only Nash equilibrium.

	C	S
C	$(-5, -5)$	$(0, -20)$
S	$(-20, 0)$	$(-1, -1)$

Example 3.8. In the Battle of Sexes, only (O, O) and (F, F) are Nash equilibria.

	O	F
O	$(2, 1)$	$(0, 0)$
F	$(0, 0)$	$(1, 2)$

In Cournot duopoly, see Section 3.3.3, $(\theta/3, \theta/3)$ is the only Nash equilibrium, as best response relations $q_1 = (\theta - q_2)/2$ and $q_2 = (\theta - q_1)/2$ are both satisfied only by $q_1 = q_2 = \theta/3$.

3.5.1 Stag Hunt

Two (in some versions more than two) hunters, players 1 and 2, can each choose to hunt

- stag (S) = a large tasty meal,
- hare (H) = also tasty but small,

but hunting stag is much more demanding and the forces of both players need to be joined (hare can be hunted individually). In a strategy-form game model, we have $N = \{1, 2\}, S_1 = S_2 = \{S, H\}$ and payoff of form

Table 3.13: Stag Hunt payoffs

	S	H
S	$(5,5)$	$(0,3)$
H	$(3,0)$	$(3,3)$

Clearly, there are two Nash equilibria – (S,S) and (H,H), where the former is strictly better for each player than the latter! Which one is more reasonable?

If each player believes that the other one will go for a hare, then (H,H) is a reasonable outcome $\implies$ a society of individualists who do not cooperate at all. On the other hand, if each player believes that the other will cooperate, then this anticipation is self-fulfilling and results in what can be called a cooperative society.

> **Note**
>
> This is supposed to explain that in the real world, there are societies that have similar endowments, access to technology and physical environment but have very different achievements, all because of self-fulfilling beliefs (or norms of behavior).

Another point of view might be to notice that (H,H) is less risky. The minimum secured by playing S is 0 as opposed to 3 by playing H (We will get to this minimax principle later). So it seems to be rational to expect (H,H)?

Theorem 3.2.

1. *If $\mathbf{s}^*$ is a strictly dominant strategy equilibrium, then it is the unique Nash equilibrium.*
2. *Each Nash equilibrium is rationalizable and survives IESDS.*
3. *If S is finite, neither rationalizability nor IESDS creates new Nash equilibria.*

Corollary 3.2. *Assume that S is finite. If rationalizability or IESDS results in a unique strategy profile, then this profile is a Nash equilibrium.*

3.5.2 Interpretations of Nash Equilibria

Although we have seen the rigorous definition of a Nash equilibrium, it can be more intuitively interpreted in the following two ways.

When the goal is to give advice to all of the players in a game (i.e., to advise each player on what strategy to choose), any advice that was not an equilibrium would have the unsettling property that there would always be some player for whom the advice was bad, in the sense that, if all other players followed the parts of the advice directed to them, it would be better for some player to do differently than he was advised. If the advice is an equilibrium, however, this will not be the case, because the advice to each player is the best response to the advice given to the other players.

On the other hand, when the goal is prediction rather than prescription, a Nash equilibrium can also be interpreted as a potential stable point of a dynamic adjustment process in which individuals adjust their behavior to that of the other players in the game, searching for strategy choices that will give them better results.

4 Mixed Strategies in Strategic-Form Games

4.1 Mixed Strategies

As pointed out before, neither of the solution concepts discussed so far has to exist in pure strategies, e.g. Example 4.1.

Example 4.1 (Rock-Paper-Scissors)**.** Consider a game

	R	P	S
R	$(0, 0)$	$(-1, 1)$	$(1, -1)$
P	$(1, -1)$	$(0, 0)$	$(-1, 1)$
S	$(-1, 1)$	$(1, -1)$	$(0, 0)$

There are no strictly dominant pure strategies and no strategy is strictly dominated (IESDS removes nothing). Also, each strategy is a best response to some strategy of the opponent (rationalizability removes nothing). Finally, there are no pure Nash equilibria, as no pure strategy profile allows each player to play a best response to the strategy of the other player.

How to solve this? Let the players randomize their choice of pure strategies.

Definition 4.1. Let A be a finite set. A **probability distribution over** A is a function $\sigma : A \to [0, 1]$ such that $\sum_{a \in A} \sigma(a) = 1$

We denote by $\Delta(A)$ the set of *all probability distributions over* A.

> **Note**
>
> Consider two strategies A_1, A_2, then their respective probabilities are p, $1 - p$. For three strategies A_1, A_2, A_3 they span a subspace of $\mathbb{R}^3$ such that $1 = \sigma(A_1) + \sigma(A_2) + \sigma(A_3)$ – this gives a part of the plane (simplex).

Example 4.2. Consider $A = \{a, b, c\}$ and a function $\sigma : A \to [0, 1]$ such that $\sigma(a) = 1/4, \sigma(b) = 3/4$ and $\sigma(c) = 0$. Then σ in $\Delta(A)$.

Let us fix a strategic-form game $G = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$.

! Important

From now on, **assume two players and both S_i finite!**

Definition 4.2 (Mixed Strategy)**.** A **mixed strategy** of player i is a probability distribution $\sigma \in \Delta(S_i)$ over S_i. We denote by $\Sigma_i = \Delta(S_i)$ the set of all mixed strategies of player i. We define $\Sigma = \Sigma_1 \times \Sigma_2$, the set of all **mixed strategy profiles**.

We identify each $s_i \in S_i$ with a mixed strategy σ that assigns probability one to s_i (and zero to other pure strategies). Let $\boldsymbol{\sigma} = (\sigma_1, \sigma_2)$ be a mixed strategy profile. Intuitively, we assume that each player i **randomly** selects his pure strategy according to σ_i and **independently** of his opponents. Thus for $\boldsymbol{s} = (s_1, s_2) \in S = S_1 \times S_2$ we have that

$$\boldsymbol{\sigma}(\boldsymbol{s}) := \sigma_1(s_1) \cdot \sigma_2(s_2)$$

is the probability that the players randomly select the pure strategy profile $\boldsymbol{s}$ according to the mixed strategy profile $\boldsymbol{\sigma}$.

i Note

We abuse notation a bit here: $\boldsymbol{\sigma}$ denotes two things, a vector of mixed strategies as well as a probability distribution on S.

Example 4.3 (Rock-Paper-Scissors)**.** Consider a game of Rock-Paper-Scissors given by

	R	P	S
R	$(0,0)$	$(-1,1)$	$(1,-1)$
P	$(1,-1)$	$(0,0)$	$(-1,1)$
S	$(-1,1)$	$(1,-1)$	$(0,0)$

As an example of mixed strategy σ_1 consider

$$\sigma_1(R) = \frac{1}{2}, \quad \sigma_1(P) = \frac{1}{3}, \quad \sigma_1(S) = \frac{1}{6},$$

which we sometimes write as $(\frac{1}{2}(R), \frac{1}{3}(P), \frac{1}{6}(S))$, or only $(\frac{1}{2}, \frac{1}{3}, \frac{1}{6})$ if the order of pure strategies is fixed. Now assume a mixed strategy profile $\boldsymbol{\sigma} = (\sigma_1, \sigma_2)$ where

$$\sigma_1 = \left(\frac{1}{2}(R), \frac{1}{3}(P), \frac{1}{6}(S)\right), \quad \sigma_2 = \left(\frac{1}{3}(R), \frac{2}{3}(P), 0(S)\right).$$

Then the probability $\boldsymbol{\sigma}(R, P)$ that the pure strategy profile (R, P) will be played by players playing the mixed profile (σ_1, σ_2) is

$$\sigma_1(R) \cdot \sigma_2(P) = \frac{1}{2} \cdot \frac{2}{3} = \frac{1}{3}$$

But now what is the suitable notion of payoff?

Definition 4.3 (Expected Payoff)**.** The **expected payoff** of player i under a mixed strategy profile $\boldsymbol{\sigma} \in \Sigma$ is

$$u_i(\boldsymbol{\sigma}) := \sum_{\boldsymbol{s} \in S} \boldsymbol{\sigma}(\boldsymbol{s}) u_i(\boldsymbol{s}) = \sum_{s_1 \in S_1} \sum_{s_2 \in S_2} \sigma_1(s_1)\sigma_2(s_2) u_i(s_1, s_2),$$

i.e. it is the "weighted average" of what player i wins under each pure strategy profile $\boldsymbol{s}$, weighted by the probability of that profile.

Proposition 4.1. *Every rational player strives to maximize his own expected payoff.*

> **Note**
>
> This assumption is not always completely convincing… Not everyone must be fond of lottery (and still be rational).

Example 4.4 (Matching Pennies)**.** Consider a game Matching Pennies given by

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

Each player secretly turns a penny into heads or tails, and then they reveal their choices simultaneously. If the pennies match, player 1 (row) wins, if they do not match, player 2 (column) wins. Now, consider $\sigma_1 = (1/3(H), 2/3(T))$ and $\sigma_2 = (1/4(H), 3/4(T))$, then

$$u_1(\sigma) = \sum_{(X,Y) \in \{H,T\}^2} \sigma_1(X)\sigma_2(Y)u_1(X,Y) = \cdots = \frac{1}{6}$$

$$u_2(\sigma) = \sum_{(X,Y) \in \{H,T\}^2} \sigma_1(X)\sigma_2(Y)u_2(X,Y) = \cdots = -\frac{1}{6}$$

4.2 Solution Concepts

We revisit the following solution concepts in mixed strategies:

- strictly dominant strategy equilibrium,
- IESDS equilibria,
- rationalizable equilibria,
- Nash equilibria;

! Important

From now on, when we say a *strategy* we implicitly mean a **mixed strategy.**

In order to deal with efficiency issues we assume that the size of the game G is defined by $|G| := |N| + \sum_{i \in N} |S_i| + \sum_{i \in N} |u_i|$ where $|u_i| = \sum_{s \in S} |u_i(\boldsymbol{s})|$ and $|u_i(\boldsymbol{s})|$ is the length of a binary encoding of $u_i(\boldsymbol{s})$ (we assume that rational numbers are encoded as quotients of two binary integers). Note that, in particular, $|G| > |S|$. This will be later needed for the complexity of certain algorithms in relation to game size.

Definition 4.4. Let $\sigma_1, \sigma_1' \in \Sigma_1$ be (mixed) strategies of player 1. Then σ_1' is **strictly dominated** by σ_1 (write $\sigma_1' \prec \sigma_1$) if

$$u_1(\sigma_1, s_2) > u_1(\sigma_1', s_2)$$

for all $s_2 \in S_2$. Symmetrically for player 2.

i Note

The above condition is equivalent to

$$u_1(\sigma_1, \sigma_2) > u_1(\sigma_1', \sigma_2)$$

for all $\sigma_2 \in \Sigma_2$.

Example 4.5. Consider a game

Table 4.4: Payoffs for one player

	X	Y
A	3	0
B	0	3
C	1	1

Is there a strictly dominated strategy? Here $\boldsymbol{\sigma} = (1/2, 1/2, 0)$ dominates $(0, 0, 1)$, aka the pure C, strategy. Then expected payoffs for strategies of this game can be plotted as seen in Figure 4.1.

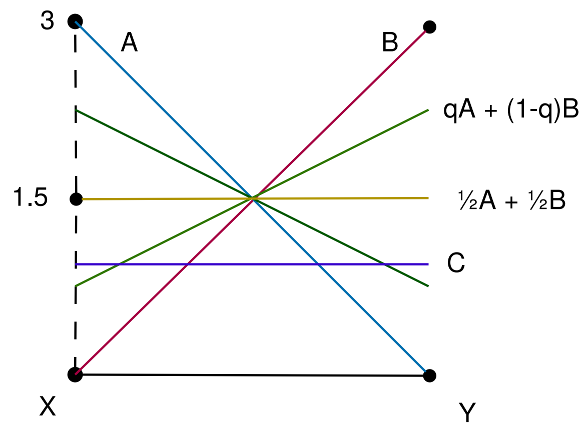

Figure 4.1: Example of a game with a pure strategy strictly dominated by a mixed strategy

Question: Is there a game with at least one strictly dominated strategy but without strictly dominated **pure** strategies? Yes, consider the following game, where C strictly dominates the mixed strategy $(1/2, 1/2, 0)$:

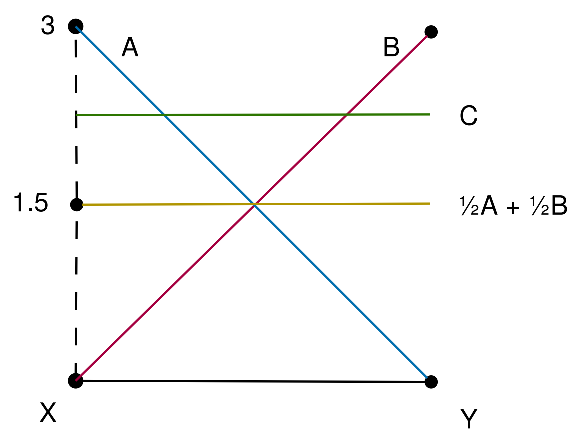

Definition 4.5. A mixed strategy $\sigma_i \in \Sigma_i$ is strictly dominant if every other mixed strategy of player i is strictly dominated by σ_i.

Definition 4.6. A strategy profile $\boldsymbol{\sigma} \in \Sigma$ is a strictly dominant strategy equilibrium if $\sigma_i \in \Sigma_i$ is strictly dominant for all $i \in N$.

Proposition 4.2. *If the strictly dominant strategy equilibrium exists, it is unique, all its strategies are pure, and rational players will play it.*

Proof. Let $\boldsymbol{\sigma}^* = (\sigma_1^*, \sigma_2^*) \in \Sigma$ be a strictly dominant strategy equilibrium. First, without loss of generality, let us assume that for player 1 there exists a non-empty non-singular subset of his pure strategies $S_1' \subseteq S_1$ such that $\sigma_1^*(s_{1,j}) > 0$ for all $j \in \{1, \dots, |S_1'|\}$ with $s_{1,j} \in S_1'$, i.e. $|\text{supp}\left(\sigma_1^*\right)| = |S_1'| > 1$. Then per Definition 4.3, we get

$$\begin{aligned} u_1(\sigma_1^*, \sigma_2^*) &= \sum_{s_{1,j} \in S_1'} \sum_{s_2 \in S_2} \sigma_1^*(s_{1,j}) \sigma_2^*(s_2) u_1(s_{1,j}, s_2) \\ &= \sum_{s_{1,j} \in S_1'} \sigma_1^*(s_{1,j}) \underbrace{\sum_{s_2 \in S_2} \sigma_2^*(s_2) u_1(s_{1,j}, s_2)}_{U_1(s_{1,j}; \sigma_2^*)} \\ &= \sum_{s_{1,j} \in S_1'} \sigma_1^*(s_{1,j}) U_1(s_{1,j}; \sigma_2^*) \end{aligned}$$

Because $\boldsymbol{\sigma}^*$ is a *strictly dominant strategy equilibrium*, then per Definition 4.6, it must hold that $U_1(s_{1,j}; \sigma_2^*) \neq U_1(s_{1,k}; \sigma_2^*)$ for $j \neq k$, otherwise one could arbitrarily distribute probabilities of σ_1^* between $s_{1,j}$ and $s_{1,k}$ without changing the payoff for player 1, which is a contradiction. Hence there exists $J \in \{1, \dots, |S_1'|\}$ such $U_1(s_{1,J}; \sigma_2^*) > U_1(s_{1,j}; \sigma_2^*)$ for every $j \in \{1, \dots, |S_1'|\}$ with $j \neq J$. But then by choosing $\sigma_1'(s_{1,J}) = 1$ and $\sigma_1'(s_1) = 0$ for every $s_1 \in S_1$ such that $s_1 \neq s_{1,J}$ (i.e. playing the pure strategy $s_{1,J}$), we get

$$u_1(\sigma_1', \sigma_2^*) = U_1(s_{1,J}; \sigma_2^*) > \underbrace{\sum_{s_{1,j} \in S_1'} \sigma_1^*(s_{1,j})}_{=1} \overbrace{U_1(s_{1,j}; \sigma_2^*)}^{< U_1(s_{1,J}; \sigma_2^*)} = u_1(\sigma_1^*, \sigma_2^*),$$

which is a contradiction with the fact that $\boldsymbol{\sigma}^*$ is a *strictly dominant strategy equilibrium*. Thus any strictly dominant strategy equilibrium is comprised of only pure strategies. Then by Corollary 3.1, we obtain that $\boldsymbol{\sigma}^*$ is unique and that rational players will play it, which concludes the proof. □

Thus, to compute the strictly dominant strategy equilibrium, it is sufficient to consider only pure strategies.

Definition 4.7 (IESDS in Mixed Strategies)**.** Define a sequence $D_i^0, D_i^1, D_i^2, \dots$ of strategy sets of player i. Also denote by G_{DS}^k the game obtained from G by restricting to D_i^k, $i \in N$. We shall call the following algorithm *"Iterated Elimination of Strictly Dominated Strategies"*:

1. Initialize $k = 0$ and $D_i^0 = S_i$ for each $i \in N$.
2. For all players $i \in N$: Let D_i^{k+1} be the set of all pure strategies of D_i^k that are **not** strictly dominated in G_{DS}^k by **mixed strategies**.
3. If $D_i^{k+1} = D_i^k$ for all players $i \in N$, then stop. Otherwise, let $k := k + 1$ and go to 2.

We say that $s_i \in S_i$ **survives IESDS** if $s_i \in D_i^k$ for all $k = 0, 1, 2, \dots$ (or until stop).

Definition 4.8. A strategy profile $\boldsymbol{s} = (s_1, s_2) \in S$ is an **IESDS equilibrium** if both s_1 and s_2 survive IESDS.

Each D_i^{k+1} can be computed in polynomial time using *linear programming*.

Example 4.6. Consider a game

	X	Y
A	3	0
B	0	3
C	1	1

Let us have a look at the first iteration of IESDS. Observe that A, B are not strictly dominated by any mixed strategy. Let us construct a set of constraints on mixed strategies (possibly) strictly dominating C:

$$\begin{aligned} 3x_A + 0x_B + x_C &> 1 \\ 0x_A + 3x_B + x_C &> 1 \\ x_A, x_B, x_C &\geq 0 \\ x_A + x_B + x_C &= 1 \end{aligned}$$

4.3 Linear Programming

Linear programming is a technique for the optimization of a linear objective function, subject to linear (non-strict) inequality constraints. Formally, a linear program in so-called canonical form looks like this:

$$\begin{aligned} &\max \sum_{j=1}^{m} c_j x_j \\ &\text{s.t.} \sum_{j=1}^{m} a_{ij} x_j \leq b_i && 1 \leq i \leq n \\ &\quad x_j \geq 0 && 1 \leq j \leq m \end{aligned}$$

Here a_{ij}, b_k and c_j are real numbers and x_j's are real variables. A **feasible solution** is an assignment of real numbers to the variables $x_j, 1 \leq j \leq m$, so that the constraints are satisfied. An **optimal solution** is a feasible solution that maximizes the objective function $\sum_{j=1}^{m} c_j x_j$. We assume that coefficients a_{ij}, b_k and c_j are encoded in binary (more precisely, as fractions of two integers encoded in binary).

Theorem 4.1 (Khachiyan, Doklady Akademii Nauk SSSR, 1979)**.** *There is an algorithm that for any linear program computes an optimal solution in polynomial time.*

The algorithm uses the so-called ellipsoid method. In practice, the Khachiyan's is not used. Usually simplex algorithm is used even though its theoretical complexity is exponential. There is also a polynomial time algorithm (by Karmarkar) which has better complexity upper bounds than the Khachiyan's and sometimes works even better than the simplex. There exist several advanced linear programming solvers (usually parts of larger optimization packages) implementing various heuristics for solving large-scale problems, sensitivity analysis, etc.

Example 4.7 (continued Example 4.6)**.** The linear program for deciding whether C is strictly dominated: The program maximizes y under the following constraints:

$$\begin{aligned} 3x_A + 0x_B + x_C &\geq 1 + y \\ 0x_A + 3x_B + x_C &\geq 1 + y \\ x_A, x_B, x_C &\geq 0 \\ x_A + x_B + x_C &= 1 \\ y &\geq 0 \end{aligned}$$

Here y just implements the strict inequality using $\geq$, we look for a solution with $y > 0$. The maximum $y = \frac{1}{2}$ is attained at $x_A = \frac{1}{2}$ and $x_B = \frac{1}{2}$. Note that in step 2 of Definition 4.7, it is not sufficient to consider domination by only pure strategies. Consider the following zero-sum game

Table 4.6: Zero-sum game

	X	Y
A	3	0
B	0	3
C	1	1

Here C is strictly dominated by $(\sigma_1(A), \sigma_1(B), \sigma_1(C)) = \left(\frac{1}{2}, \frac{1}{2}, 0\right)$, but no strategy is strictly dominated in pure strategies.

4.4 Best Response in Mixed Strategies

Definition 4.9. A **(mixed) belief** of player 1 is a mixed strategy σ_2 of player 2 (and vice versa).

Definition 4.10. A mixed strategy $\sigma_1 \in \Sigma_1$ is a **best response** to a belief $\sigma_2 \in \Sigma_2$ if

$$u_1(\sigma_1, \sigma_2) \geq u_1(s_1, \sigma_2)$$

for all $s_1 \in S_1$. Denote by $\mathrm{BR}_1(\sigma_2)$ the set of all best responses of player 1. Symmetrically for player 2.

> ℹ Note
>
> The above condition is equivalent to
>
> $$u_1(\sigma_1, \sigma_2) \geq u_1(\sigma_1', \sigma_2)$$
>
> for all $\sigma_1' \in \Sigma_1$.

Example 4.8. Consider a game with the following payoffs for player 1:

Table 4.7: Payoffs for player 1

	X	Y
A	2	0
B	0	2
C	1	1

- Player 1 (row) plays $\sigma_1 = (a(A), b(B), c(C))$.
- Player 2 (column) plays $(q(X), (1-q)(Y))$ (we write just q).

Because

$$u_1(A, q) = 2q$$
$$u_1(B, q) = 2(1-q)$$
$$u_1(C, q) = 1$$

then

$$\mathrm{BR}_1(q) = \begin{cases} A, & q > \frac{1}{2}, \\ B, & q < \frac{1}{2}, \\ \langle (A, B, C), \boldsymbol{c} \rangle, \ \sum \boldsymbol{c} = 1, & q = \frac{1}{2}. \end{cases}$$

For σ_1 such that $\sigma_1(A), \sigma_1(B) > 0$ and

$$u_1(A, q) < u_1(B, q),$$

then $\bar{\sigma}_1$ such that

$$\bar{\sigma}_1(A) = 0,$$
$$\bar{\sigma}_1(B) = \sigma_1(A) + \sigma_1(B)$$

is a better response.

4.5 Rationalizability in Mixed Strategies (Two Players)

We will begin with an assumption: A rational player 1 with a belief σ_2 always plays a best response to σ_2 (the same for player 2).

Definition 4.11. A pure strategy $s_1 \in S_1$ of player 1 is never best response if it is not a best response to any belief σ_2 (similarly for player 2).

No rational player plays a strategy that is never best response.

Definition 4.12 (Rationalizable)**.** Define a sequence $R_i^0, R_i^1, \dots$ of strategy sets of player i. Also, denote by G_{Rat}^k the game obtained from G by restricting to R_i^k for $i \in N$. Consider the following algorithm

1. Initialize $k = 0$ and $R_i^0 = S_i$ for each $i \in N$.
2. For all players $i \in N$: Let R_i^{k+1} be the set of all strategies of R_i^k that are best responses to some **(mixed) beliefs** in G_{Rat}^k.
3. Let $k := k + 1$ and go to 2.

We say that $s_i \in S_i$ is *rationalizable* if $s_i \in R_i^k$ for all $k = 0, 1, 2, \dots$ (or until stop).

Definition 4.13. A strategy profile $\boldsymbol{s} = (s_1, s_2) \in S$ is a **rationalizable equilibrium** if both s_1 and s_2 are rationalizable.

Example 4.9. Consider again a game G given by

	X	Y
A	3	0
B	0	3
C	1	1

What pure strategies of player 1 are strictly dominated? In pure strategies, C is not strictly dominated, but it is never best response. In mixed strategies, the strategy $\left(\frac{1}{2}(A), \frac{1}{2}(B), 0(C)\right)$ strictly dominates C. What pure strategies of player 1 are never best responses? We can make the following observation: The set of strictly dominated pure strategies coincides with the set of pure never best responses! ... and this holds in general for two-player games:

Theorem 4.2. *A pure strategy s_1 of player 1 is never best response to any belief σ_2 if and only if s_1 is strictly dominated by a strategy $\sigma_1 \in \Sigma_1$ (similarly for player 2).*

Corollary 4.1. *It follows that a strategy of S_i survives IESDS $\iff$ it is rationalizable.*

As opposed to pure strategies, the IESDS and rationalizability coincide in mixed strategies.

4.6 Mixed Nash Equilibria

Definition 4.14. A mixed-strategy profile $\boldsymbol{\sigma}^* = (\sigma_1^*, \sigma_2^*) \in \Sigma$ is a **(mixed) Nash equilibrium** if σ_1^* is a best response to σ_2^* and σ_2^* is a best response to σ_1^*. That is

$$u_1(\sigma_1^*, \sigma_2^*) \geq u_1(s_1, \sigma_2^*) \quad \forall s_1 \in S_1,$$
$$u_2(\sigma_1^*, \sigma_2^*) \geq u_2(\sigma_1^*, s_2) \quad \forall s_2 \in S_2.$$

Theorem 4.3 (Nash)**.** *Every finite game in strategic form has a Nash equilibrium.*

Example 4.10. Consider a game given by

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

Player 1 (row) plays $(p(H), (1-p)(T))$ (we write just p) and player 2 (column) plays $(q(H), (1-q)(T))$ (we write q). Compute all Nash equilibria. What are the expected payoffs of playing pure strategies for player 1? From

$$u_1(H, q) = 2q - 1 \qquad \wedge \qquad u_1(T, q) = 1 - 2q,$$

we get

$$u_1(p, q) = pu_1(H, q) + (1-p)u_1(T, q) = p(2q-1) + (1-p)(1-2q).$$

We can obtain the best response correspondence $\mathrm{BR}_1(q)$

$$\mathrm{BR}_1(q) = \begin{cases} T, & q < \frac{1}{2}, \\ p \in [0,1], & q = \frac{1}{2}, \\ H, & q > \frac{1}{2} \end{cases}$$

and we can repeat the same process for player 2. The only situation where they both play their best response to each other – the intersection of $\mathrm{BR}_1(q)$ and $\mathrm{BR}_2(p)$ – is the only Nash equilibrium $\boldsymbol{\sigma} = (\sigma_1, \sigma_2) = (\frac{1}{2}, \frac{1}{2})$.

4.6.1 Properties of Mixed Nash Equilibria

Lemma 4.1. *Every Nash equilibrium* $\boldsymbol{\sigma}^* = (\sigma_1^*, \sigma_2^*) \in \Sigma$ *satisfies*

- $u_1(s_1, \sigma_2^*) = u_1(\boldsymbol{\sigma}^*)$ *for* $s_1 \in \mathrm{supp}\left(\sigma_1^*\right)$*;*
- $u_2(\sigma_1^*, s_2) = u_1(\boldsymbol{\sigma}^*)$ *for* $s_2 \in \mathrm{supp}\left(\sigma_2^*\right)$*.*

Proof. Let's now prove Lemma 4.1. Without loss of generality, consider only player 1 and assume that $\boldsymbol{\sigma}^*$ is a Nash equilibrium. The latter assumption implies $u_1(s_1, \sigma_2^*) \le u_1(\boldsymbol{\sigma}^*)$ for all $s_1 \in S_1$. Now, if there exists $s_1' \in \operatorname{supp}(\sigma_1^*) \subseteq S_1$ satisfying $u_1(s_1', \sigma_2^*) < u_1(\boldsymbol{\sigma}^*)$, then because $\sigma^*(s_1') > 0$ we have

$$\begin{aligned} u_1(\boldsymbol{\sigma}^*) &= \sum_{s_1 \in S_1} \sigma_1^*(s_1) u_1(s_1, \sigma_2^*) \\ &= \sigma_1^*(s_1') \underbrace{u_1(s_1', \sigma_2^*)}_{< u_1(\boldsymbol{\sigma}^*)} + \sum_{s_1 \in S_1 \setminus \{s_1'\}} \sigma_1^*(s_1) \underbrace{u_1(s_1, \sigma_1^*)}_{\le u_1(\boldsymbol{\sigma}^*)} \\ &< \sigma_1^*(s_1') u_1(\boldsymbol{\sigma}^*) + \sum_{s_1 \in S_1 \setminus \{s_1'\}} \sigma_1^*(s_1) u_1(\boldsymbol{\sigma}^*) \\ &= \sum_{s_1 \in S_1} \sigma_1^*(s_1) u_1(\boldsymbol{\sigma}^*) = u_1(\boldsymbol{\sigma}^*), \end{aligned}$$

which is a contradiction. Note that we could only sum over the $\operatorname{supp}(\sigma_1^*) \subseteq S_1$. Thus $u_1(s_1, \sigma_2^*) = u_1(\boldsymbol{\sigma}^*)$ for all $s_1 \in S_1$. ☐

Note

Intuitively, not playing simply a pure strategy can give you a better utility. But if it were to give you less, it would drag the whole averaged utility of the equilibrium down, which would be a contradiction with the initial utility value of the equilibrium.

Example 4.11. Consider again a game given by

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

Player 1 (row) plays $(p(H), (1-p)(T))$ (we write just p) and player 2 (column) plays $(q(H), (1-q)(T))$ (we write q). Compute all Nash equilibria.

Firstly, there are no pure strategy equilibria. Also, there are no equilibria where only player 1 randomizes (in other words, e.g. $|\operatorname{supp}(\sigma_1^*)| = 1$ and $|\operatorname{supp}(\sigma_2^*)| = 2$): Indeed, assume that (p, H) is such an equilibrium. Then by Lemma 4.1

$$1 = u_1(H, H) = u_1(T, H) = -1$$

is a contradiction. Also, (p, T) cannot be an equilibrium. Assume now both players randomize, so $p, q \in (0, 1)$. The expected payoffs of playing pure strategies for player 1:

$$u_1(H, q) = 2q - 1 \quad \wedge \quad u_1(T, q) = 1 - 2q$$

and similarly

$$u_2(p, H) = 1 - 2p \quad \wedge \quad u_2(p, T) = 2p - 1.$$

Again, by Lemma 4.1, such Nash equilibria must satisfy

$$2q - 1 = 1 - 2q \quad \wedge \quad 2p - 1 = 1 - 2p,$$

which are linear equations and from them we get $p = q = \frac{1}{2}$ – the only Nash equilibrium.

> **Note**
>
> In other words, we try to remove the element of randomness as much as we can by considering Lemma 4.1 and pure strategies.

Example 4.12. Consider the Battle of Sexes game given by the following table:

	O	F
O	$(2, 1)$	$(0, 0)$
F	$(0, 0)$	$(1, 2)$

Clearly, there are two pure strategy equilibria (O, O) and (F, F), and no Nash equilibrium where only one player randomizes. Now assume that

- player 1 (row) plays $(p(O), (1 - p)F)$ (we write just p) and
- player 2 (column) plays $(q(O), (1 - q)F)$ (we write just q),

where $p, q \in (0, 1)$. By Lemma 4.1, such Nash equilibria must satisfy

$$2q = 1 - q \quad \wedge \quad p = 2(1 - p),$$

which holds only for $q = \frac{1}{3}$ and $p = \frac{2}{3}$.

What did we do in these examples (see Example 4.11 and Example 4.12)? We went through all support combinations for both players.

(pure, one player mixing, both mixing)

For each pair of support sets we tried to find equilibria in strategies with these supports.

(in Battle of Sexes: two pure, no equilibrium with just one player mixing, one equilibrium when both mixing)

Whenever one of the supports was non-singleton, we reduced the computation of Nash equilibria to linear equations.

Lemma 4.2. *Let $\boldsymbol{\sigma}^* = (\sigma_1^*, \sigma_2^*) \in \Sigma$ be a mixed profile. Assume there exists $w_1, w_2 \in \mathbb{R}$ such that*

- $u_1(s_1, \sigma_2^*) = w_1$ *for* $s_1 \in \operatorname{supp}(\sigma_1^*)$,
- $u_1(s_1, \sigma_2^*) \leq w_1$ *for* $s_1 \notin \operatorname{supp}(\sigma_1^*)$,

- $u_2(\sigma_2^*, s_2) = w_2$ *for* $s_2 \in \operatorname{supp}(\sigma_2^*)$,
- $u_2(\sigma_2^*, s_2) \le w_2$ *for* $s_2 \notin \operatorname{supp}(\sigma_2^*)$.

Then $u_1(\boldsymbol{\sigma}^*) = w_1$ *and* $u_2(\boldsymbol{\sigma}^*) = w_2$, *and* $\boldsymbol{\sigma}^*$ *is a Nash equilibrium.*

Proof. Consider just the player 1 (for player 2 similarly):

$$\begin{aligned} u_1(\boldsymbol{\sigma}^*) &= \sum_{s_1 \in S_1} \sigma_1^*(s_1) u_1(s_1, \sigma_2^*) \\ &= \sum_{s_1 \in \operatorname{supp}(\sigma_1^*)} \sigma_1^*(s_1) u_1(s_1, \sigma_2^*) \\ &= \sum_{s_1 \in \operatorname{supp}(\sigma_1^*)} \sigma_1^*(s_1) w_1 \\ &= w_1 \sum_{s_1 \in \operatorname{supp}(\sigma_1^*)} \sigma_1^*(s_1) = w_1. \end{aligned}$$

Now the fact that $\boldsymbol{\sigma}^*$ is a Nash equilibrium follows from the definition. □

4.7 Computing Nash Equilibria

Every Nash equilibrium $\boldsymbol{\sigma}^*$ can be computed by finding appropriate w_1, w_2 so that

- $u_1(s_1, \sigma_2^*) = w_1$ for $s_1 \in \operatorname{supp}(\sigma_1^*)$,
- $u_1(s_1, \sigma_2^*) \le w_1$ for $s_1 \notin \operatorname{supp}(\sigma_1^*)$,
- $u_2(\sigma_2^*, s_2) = w_2$ for $s_2 \in \operatorname{supp}(\sigma_2^*)$,
- $u_2(\sigma_2^*, s_2) \le w_2$ for $s_2 \notin \operatorname{supp}(\sigma_2^*)$.

Indeed,

- by Lemma 4.2, all $\boldsymbol{\sigma}^*$ and w_1, w_2 satisfying the above inequalities give a Nash equilibrium $\boldsymbol{\sigma}^*$ with $u_1(\boldsymbol{\sigma}^*) = w_1$ and $u_2(\boldsymbol{\sigma}^*) = w_2$,
- by Lemma 4.1, for every Nash equilibrium $\boldsymbol{\sigma}^*$ choosing $w_1 = u_1(\boldsymbol{\sigma}^*)$ and $w_2 = u_2(\boldsymbol{\sigma}^*)$ satisfies the above inequalities.

Suppose that we somehow know the supports $\operatorname{supp}(\sigma_1^*), \operatorname{supp}(\sigma_2^*)$ for some Nash equilibrium $\boldsymbol{\sigma}^* = (\sigma_1^*, \sigma_2^*)$ (which itself is unknown to us). We may consider all $\sigma_i^*(s_i)$'s and both w_1, w_2's as variables and use the above conditions to design a system of inequalities capturing Nash equilibria with the given support sets $\operatorname{supp}(\sigma_1^*), \operatorname{supp}(\sigma_2^*)$.

4.7.1 Support Enumeration

Definition 4.15 (Nash equilibrium given the support sets)**.** To simplify notation, assume that for i we have $S_i = \{1, \dots, m_i\}$. Then $\sigma_i(j)$ is the probability of the pure strategy j in the mixed strategy σ_i. Now fix supports $\text{supp}_i \subseteq S_i$ for every $i \in \{1, 2\}$ and consider the following system of constraints with variables $\sigma_1(1), \dots, \sigma_1(m_1), \sigma_2(1), \dots, \sigma_2(m_2), w_1, w_2$:

1. For all $k \in \text{supp}_1$ and $l \in \text{supp}_2$:
$$u_1(k, \sigma_2) = \sum_{l' \in S_2} \sigma_2(l') u_1(k, l') = w_1, \quad \sum_{k' \in S_1} \sigma_1(k') u_2(k', l) = w_2.$$
2. For all $k \notin \text{supp}_1$ and $l \notin \text{supp}_2$:
$$\sum_{l' \in S_2} \sigma_2(l') u_1(k, l') \leq w_1, \quad \sum_{k' \in S_1} \sigma_1(k') u_2(k', l) \leq w_2.$$
3. For all $i \in \{1, 2\}$: $\sigma_i(1) + \dots + \sigma_i(m_i) = 1$.
4. For all $i \in \{1, 2\}$ and all $k \in \text{supp}_i$: $\sigma_i(k) \geq 0$.

> **i Strict inequality**
>
> Technically, there should be a strict inequality, but then it wouldn't be a linear program anymore. What's more, it can be shown that this relaxation gives us the same solutions.

5. For all $i \in \{1, 2\}$ and all $k \notin \text{supp}_i$: $\sigma_i(k) = 0$.

Therefore we can see that the constraints are **linear** for the two players. So the question now is how to find the supp_1 and supp_2... And we can simply *guess*!

Definition 4.16 (Algorithm for finding a Nash equilibrium)**.** Consider a two-player strategic game G with strategy sets $S_1 = \{1, \dots, m_1\}$ and $S_2 = \{1, \dots, m_2\}$ and rational payoffs u_1, u_2 as an input. The output of this algorithm is *a Nash equilibrium* $\boldsymbol{\sigma}^*$.

For all possible $\text{supp}_1 \subseteq S_1$ and $\text{supp}_2 \subseteq S_2$:

- Check if the corresponding system of linear constraints from Definition 4.15 has a feasible solution $\boldsymbol{\sigma}^*, w_1^*, w_2^*$
- If so, **STOP**: the feasible solution $\boldsymbol{\sigma}^*$ is a Nash equilibrium satisfying $u_i(\boldsymbol{\sigma}^*) = w_i^*$

Unfortunately, there are $2^{(m_1+m_2)}$ possible subsets $\text{supp}_1, \text{supp}_2$ and as such, this algorithm requires worst-case exponential time. What's more, we can formulate the following remarks:

- The algorithm in Definition 4.16 combined with Theorem 4.3 and properties of linear programming imply that every finite two-player game has a rational Nash equilibrium (furthermore, the rational numbers have polynomial representation in binary).

- The algorithm can be used to compute *all* Nash equilibria.

  There are algorithms for computing (a finite representation of) a set of all feasible solutions of a given linear constraint system.

- The algorithm can be used to compute "good" equilibria.

  For example, to find a Nash equilibrium maximizing the sum of all expected payoffs (the "social welfare") it suffices to solve the system of constraints while maximizing $w_1 + w_2$. More precisely, the algorithm can be modified as follows:

  – Initialize $W := -\infty$ (W stores the current maximum welfare)
  – For all possible $\text{supp}_1 \subseteq S_1$ and $\text{supp}_2 \subseteq S_2$:
    * Find the maximum value $\max(w_1 + w_2)$ of $w_1 + w_2$ so that constraints are satisfiable (using linear programming)
    * Put $W := \max\{W, \max(w_1 + w_2)\}$
  – Return W.

4.8 Complexity results

Theorem 4.4. *Given a two-player game in strategic form, a mixed Nash equilibrium can be computed in exponential time.*

Theorem 4.5. *All the following problems are NP-complete: Given a two-player game in strategic form, does it have:*

1. *a Nash equilibrium in which player 1 has utility at least a given amount v?*
2. *a Nash equilibrium in which the sum of expected payoffs of the two players is at least a given amount v?*
3. *a Nash equilibrium with a support of size greater than a given number?*
4. *a Nash equilibrium whose support contains a given strategy s?*
5. *a Nash equilibrium whose support does not contain a given strategy s?*
6. *...*

What's more, NP-hardness can be proved using reduction from SAT, see Figure 4.2.

Definition 4 *Let ϕ be a Boolean formula in conjunctive normal form (representing a SAT instance). Let V be its set of variables (with $|V| = n$), L the set of corresponding literals (a positive and a negative one for each variable[6]), and C its set of clauses. The function $v : L \to V$ gives the variable corresponding to a literal,* e.g., *$v(x_1) = v(-x_1) = x_1$. We define $G_\epsilon(\phi)$ to be the following finite symmetric 2-player game in normal form. Let $\Sigma = \Sigma_1 = \Sigma_2 = L \cup V \cup C \cup \{f\}$. Let the utility functions be*

- $u_1(l^1, l^2) = u_2(l^2, l^1) = n - 1$ *for all* $l^1, l^2 \in L$ *with* $l^1 \neq -l^2$;
- $u_1(l, -l) = u_2(-l, l) = n - 4$ *for all* $l \in L$;
- $u_1(l, x) = u_2(x, l) = n - 4$ *for all* $l \in L$, $x \in \Sigma - L - \{f\}$;
- $u_1(v, l) = u_2(l, v) = n$ *for all* $v \in V$, $l \in L$ *with* $v(l) \neq v$;
- $u_1(v, l) = u_2(l, v) = 0$ *for all* $v \in V$, $l \in L$ *with* $v(l) = v$;
- $u_1(v, x) = u_2(x, v) = n - 4$ *for all* $v \in V$, $x \in \Sigma - L - \{f\}$;
- $u_1(c, l) = u_2(l, c) = n$ *for all* $c \in C$, $l \in L$ *with* $l \notin c$;
- $u_1(c, l) = u_2(l, c) = 0$ *for all* $c \in C$, $l \in L$ *with* $l \in c$;
- $u_1(c, x) = u_2(x, c) = n - 4$ *for all* $c \in C$, $x \in \Sigma - L - \{f\}$;
- $u_1(x, f) = u_2(f, x) = 0$ *for all* $x \in \Sigma - \{f\}$;
- $u_1(f, f) = u_2(f, f) = \epsilon$;
- $u_1(f, x) = u_2(x, f) = n - 1$ *for all* $x \in \Sigma - \{f\}$.

Theorem 1 *If $(l_1, l_2, \ldots, l_n)$ (where $v(l_i) = x_i$) satisfies ϕ, then there is a Nash equilibrium of $G_\epsilon(\phi)$ where both players play l_i with probability $\frac{1}{n}$, with expected utility $n - 1$ for each player. The only other Nash equilibrium is the one where both players play f, and receive expected utility ϵ each.*

Figure 4.2: The reduction

But the real question is what is the **exact complexity** of computing Nash equilibria in two-player games?

Let us concentrate on the problem of computing one Nash equilibrium (sometimes called the sample equilibrium problem). As the class NP consists of decision problems, it cannot be directly used to characterize the complexity of the sample equilibrium problem. We use complexity classes of *function problems* such as FP, FNP, etc. The sample equilibrium problem belongs to the complexity class PPAD (which is a subclass of TFNP) for two-player games.

> **ℹ Note**
>
> A binary relation $P(x, y)$ is in TFNP if and only if there is a deterministic polynomial time algorithm that can determine whether $P(x, y)$ holds given both x and y, and for every x, there exists a y which is at most polynomially longer than x such that $P(x, y)$ holds.

> **ℹ Note**
>
> NP complexity class can be embedded into FNP by considering characteristic functions of the languages.

Can we do better than FNP (i.e. exponential time)? In what follows we show that the sample equilibrium problem can be solved in polynomial time for zero-sum two-player games.

4.8.1 Zero-sum Games

Definition 4.17. A mixed strategy $\sigma_1^* \in \Sigma_1$ is a **maxmin** strategy of player 1 if

$$\sigma_1^* \in \operatorname*{argmax}_{\sigma_1 \in \Sigma_1} \min_{s_2 \in S_2} u_1(\sigma_1, s_2), \tag{4.1}$$

where $s_2 \in S_2$ can be swapped for $\sigma_2 \in \Sigma_2$ (and there can be more than one maxmin strategy).

> **Note**
>
> Intuitively, a **maxmin** strategy σ_1^* maximizes player 1's worst-case payoff in the situation where player 2 strives to cause the greatest harm to player 1 and **knows** what strategy player 1 will play.

Similarly, $\sigma_2^* \in \Sigma_2$ is a **maxmin** strategy of player 2 if

$$\sigma_2^* \in \operatorname*{argmax}_{\sigma_2 \in \Sigma_2} \min_{s_1 \in S_1} u_2(s_1, \sigma_2).$$

Assuming a zero-sum game, i.e. $u_1 = -u_2$, this becomes

$$\sigma_2^* \in \operatorname*{argmin}_{\sigma_2 \in \Sigma_2} \max_{s_1 \in S_1} u_1(s_1, \sigma_2), \tag{4.2}$$

where again $s_1 \in S_1$ can be swapped for $\sigma_1 \in \Sigma_1$. Note the same payoff function for both players in (4.1) and (4.2)

Theorem 4.6 (von Neumann)**.** *Assume a two-player **zero-sum** game. Then*

$$\max_{\sigma_1 \in \Sigma_1} \min_{s_2 \in S_2} u_1(\sigma_1, s_2) = \min_{\sigma_2 \in \Sigma_2} \max_{s_1 \in S_1} u_1(s_1, \sigma_2).$$

Moreover, $\boldsymbol{\sigma}^* = (\sigma_1^*, \sigma_2^*) \in \Sigma$ *is a Nash equilibrium iff both* σ_1^* *and* σ_2^* *are* maxmin.

> **Tip**
>
> The *maxmin* equality in Theorem 4.6 can be equivalently expressed as
>
> $$\max_{\sigma_1 \in \Sigma_1} \min_{s_2 \in S_2} u_1(\sigma_1, s_2) = -\max_{\sigma_2 \in \Sigma_2} \min_{s_1 \in S_1} u_2(s_1, \sigma_2).$$

So to compute a Nash equilibrium it suffices to compute (arbitrary) *maxmin* strategies for both players.

4.8.2 Computing Nash Equilibria in Two-player Zero-sum Games

Assume $S_1 = \{1, \dots, m_1\}, S_2 = \{1, \dots, m_2\}$. We want to compute

$$\sigma_1^* \in \operatorname*{argmax}_{\sigma_1 \in \Sigma_1} \min_{l \in S_2} u_1(\sigma_1, l).$$

Consider a linear program with variables $\sigma_1(1), \dots, \sigma_1(m_1), v$:

$$\begin{aligned} & \max v \\ \text{s.t.} \quad & \sum_{k=1}^{m_1} \sigma_1(k) u_1(k,l) \geq v, \quad l = 1, \dots, m_2, \\ & \sum_{k=1}^{m_1} \sigma_1(k) = 1, \\ & \sigma_1(k) \geq 0, \qquad\qquad k = 1 \dots, m_1. \end{aligned} \tag{4.3}$$

Lemma 4.3. *For a mixed strategy σ_1^* it holds that $\sigma_1^* \in \operatorname{argmax}_{\sigma_1 \in \Sigma_1} \min_{l \in S_2} u_1(\sigma_1, l)$ if and only if assigning $\sigma_1(k) := \sigma_1^*(k)$ and $v := \min_{l \in S_2} u_1(\sigma_1^*, l)$ gives an optimal solution of the linear program (4.3).*

4.9 Summary and Results

As a summary:

- We have reduced the computation of Nash equilibria to the computation of maxmin strategies for both players;
- Maxmin strategies can be computed using linear programming in polynomial time;
- That is, Nash equilibria in zero-sum two-player games can be computed in polynomial time.

We have considered static games of complete information, i.e., "one-shot" games where the players know exactly what game they are playing. We modeled such games using strategic-form games. We have considered both pure strategy setting and mixed strategy setting. In both cases, we considered four solution concepts:

- Strictly dominant strategies;
- Iterative elimination of strictly dominated strategies;
- Rationalizability (i.e., iterative elimination of strategies that are never best responses);
- Nash equilibria.

In pure strategy setting:

1. Strictly dominant strategy equilibrium survives IESDS, rationalizability and is the unique Nash equilibrium (if it exists);
2. In finite games, rationalizable equilibria survive IESDS, and IESDS preserves the set of Nash equilibria;

3. In finite games, rationalizability preserves Nash equilibria.

In mixed setting on the other hand:

1. In finite two-player games, IESDS and rationalizability coincide;
2. Strictly dominant strategy equilibrium survives IESDS (rationalizability) and is the unique Nash equilibrium (if it exists);
3. In finite games, IESDS (rationalizability) preserves Nash equilibria.

> **Tip**
>
> The proofs for 2. and 3. in the mixed setting are similar to the corresponding proofs in the pure setting.

Again, using the expected value as we did gives us weird results like IESDS and rationalizability being the same. Strictly dominant strategy equilibria coincide in pure and mixed settings and can be computed in polynomial time. IESDS and rationalizability can be implemented in polynomial time in the pure setting as well as in the mixed setting. In the mixed setting, linear programming is needed to implement one step of IESDS (rationalizability). Nash equilibria can be computed for two-player games

- in polynomial time for zero-sum games (using von Neumann's Theorem 4.6 and linear programming);
- in exponential time using support enumeration;
- in PPAD using Lemke-Howson (omitted).

4.10 Modes of domination

To simplify, let us consider only **pure strategies**. Recall that for $s_i, s_i' \in S_i$ a strategy s_i' is **strictly dominated** if $u_i(s_i, \mathbf{s}_{-i}) > u_i(s_i', \mathbf{s}_{-i})$ for all $\mathbf{s}_{-i} \in S_{-i}$.

Let $s_i, s_i' \in S_i$. Then s_i' is **weakly dominated** by s_i if $u_i(s_i, \mathbf{s}_{-i}) \geq u_i(s_i', \mathbf{s}_{-i})$ for all $\mathbf{s}_{-i} \in S_{-i}$ and there is $\mathbf{s}'_{-i} \in S_{-i}$ such that $u_i(s_i, \mathbf{s}'_{-i}) > u_i(s_i', \mathbf{s}'_{-i})$.

Let $s_i, s_i' \in S_i$. Then s_i' is **very weakly dominated** by s_i if $u_i(s_i, \mathbf{s}_{-i}) \geq u_i(s_i', \mathbf{s}_{-i})$ for all $\mathbf{s}_{-i} \in S_{-i}$.

Then we say that a strategy is (strictly, weakly, very weakly) dominant if it (strictly, weakly, very weakly) dominates any other strategy.

Theorem 4.7. *Any pure strategy profile $\mathbf{s} \in S$ such that s_i is very weakly dominant is a Nash equilibrium.*

The same claim can be also proven in the mixed strategy setting.

5 Extensive Form Games

Static games (modeled using strategic-form games) cannot capture games that unfold over time. In particular, as all players move simultaneously, there is no way to model situations in which the order of moves is important. Imagine e.g. chess where players take turns, in every round, a player knows all turns of the opponent before making his own turn. There are many examples of dynamic games: markets that change over time, political negotiations, models of computer systems, etc.

We model dynamic games using **extensive-form games**, a tree-like model that allows us to express the sequential nature of games. We start with perfect information games, where each player always knows the results of all previous moves. Then generalize to imperfect information, where players may have only partial knowledge of these results (e.g. most card games).

Example 5.1. Consider a game:

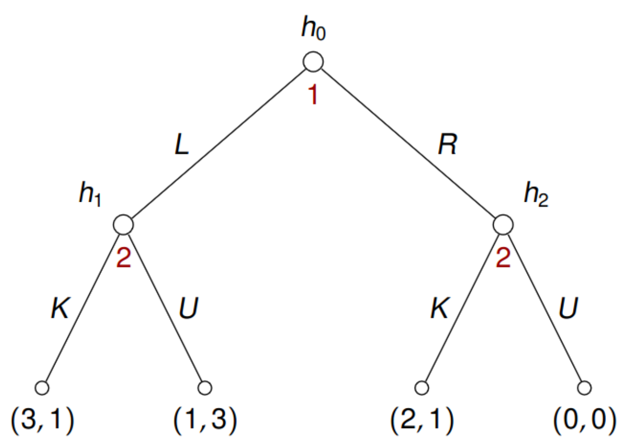

Figure 5.1: Extensive form games

Here h_0, h_1, h_2 are non-terminal nodes, leaves are terminal nodes. Each non-terminal node is owned by a player who chooses an action, e.g. h_1 is owned by player 2 who chooses either K or U. Every action results in a transition to a new node, so choosing L in h_0 results in a move to h_1. When a play reaches a terminal node, players collect payoffs, e.g. the leftmost terminal node gives 3 to player 1 and 1 to player 2.

5.1 Basic Concepts

Definition 5.1. A *perfect-information extensive-form game* is a tuple $G = (N, A, H, Z, \chi, \rho, \pi, h_0, \boldsymbol{u})$ where

- $N = \{1, \dots, n\}$ is a set of n **players**, A is a (single) set of **actions**;
- H is a set of **non-terminal** (choices) nodes, Z is a set of **terminal** nodes (assume $Z \cap H = \emptyset$), denote $\mathscr{H} = H \cup Z$;
- $\chi : H \to (2^A \setminus \{\emptyset\})$ is the **action function**, which assigns to each non-terminal choice node a non-empty set of **enabled** actions;
- $\rho : H \to N$ is the **player function**, which assigns to each non-terminal node a player i who chooses an action there, we define $H_i := \{h \in H | \rho(h) = i\}$;
- $\pi : H \times A \to \mathscr{H}$ is the **successor function**, which maps a non-terminal node and an action to a new node, such that
  - h_0 is the only node that is not in the image of π (i.e. nothing maps to the root);
  - for all $h_1, h_2 \in H$ and for all $a_1 \in \chi(h_1)$ and all $a_2 \in \chi(h_2)$, if $\pi(h_1, a_1) = \pi(h_2, a_2)$, then $h_1 = h_2$ and $a_1 = a_2$;
- $\boldsymbol{u} = (u_1, \dots, u_n)$, where each $u_i : Z \to \mathbb{R}$ is a **payoff function** for player i in the terminal nodes of Z.

We say that h' is a **child** of h, and h is a **parent** of h' if there is $a \in \chi(h)$ such that $h' = \pi(h, a)$. A **path** from $h \in H$ to $h' \in H$ is a sequence $h_1 a_2 h_2 a_3 h_3 \dots h_{k-1} a_k h_k$ where $h_1 = h, h_k = h'$ and $\pi(h_{j-1}, a_j) = h_j$ for every $1 < j \leq k$.

> **Note**
>
> Note that, in particular, h is a path from h to h.

Also, $h' \in \mathscr{H}$ is **reachable** from $h \in \mathscr{H}$ if there is a path from h to h'.

> **Note**
>
> If h' is reachable from h, we say that h' is a descendant of h and h is an ancestor of h'.

Every perfect-information extensive-form game can be seen as a game on a *rooted tree* $(\mathscr{H}, E, h_0)$ where

- $\mathscr{H} = H \cup Z$ is a set of nodes,
- $E \subseteq \mathscr{H} \times \mathscr{H}$ is a set of edges defined by $(h, h') \in E$ iff $h \in H$ and there is a $a \in \chi(h)$ such that $\pi(h, a) = h'$,
- h_0 is the root.

Example 5.2. Consider a game

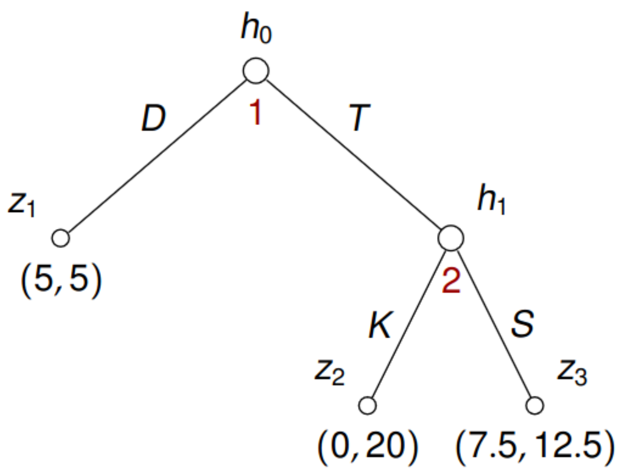

Figure 5.2: Trust game

Two players, both start with 5$ and player 2 is a broker who wants to "help" player 1. Player 1 either distrusts (D) player 2 and keeps the money (payoffs (5, 5)), or trusts (T) player 2 and passes 5$ to player 2. If player 1 chooses to trust player 2, the total money (10) is doubled by the experimenter in the hands of player 2. Player 2 may either keep (K) the additional 15$ (resulting in (0, 20)), or share (S) it with player 1 (resulting in (7.5, 12.5)).

Let us put $N = \{1, 2\}$, $A = \{D, T, K, S\}$, $H = \{h_0, h_1\}$ and $Z = \{z_1, z_2, z_3\}$. Moreover

$$\chi(h_0) = \{D, T\}, \quad \chi(h_1) = \{S, K\}$$

and $\rho(h_0) = 1$, $\rho(h_1) = 2$ with

$$\pi(h_0, D) = z_1, \quad \pi(h_0, T) = h_1, \quad \pi(h_1, K) = z_2, \quad \pi(h_1, S) = z_3$$

and lastly

$$\boldsymbol{u}(z_1) = (5, 5), \quad \boldsymbol{u}(z_2) = (0, 20), \quad \boldsymbol{u}(z_3) = (7.5, 12.5).$$

Example 5.3 (Stackelberg Competition)**.** Very similar to Cournot duopoly, this time we have two identical firms, players 1 and 2, produce some good. Denote by q_1 and q_2 quantities produced by firms 1 and 2, resp. The total quantity of products in the market is $q_1 + q_2$. The price of each item is $\kappa - q_1 - q_2$ where $\kappa > 0$ is fixed. Firms have a common per-item production cost c.

As opposed to Cournot duopoly, firm 1 moves first, and chooses the quantity $q_1 \in [0, \infty)$. Afterward, the firm 2 chooses $q_2 \in [0, \infty)$ (knowing q_1) and then the firms get their payoffs.

As an extensive-form game, we get

$$N = \{1, 2\}, \quad A = [0, \infty),$$
$$H = \left\{h_0, h_1^{q_1} | q_1 \in [0, \infty)\right\}, \quad Z = \{z^{q_1, q_2} | q_1, q_2 \in [0, \infty)\},$$
$$\chi(h_0) = [0, \infty), \quad \chi(h_1^{q_1}) = [0, \infty), \quad \rho(h_0) = 1, \quad \rho(h_1^{q_1}) = 2,$$
$$\pi(h_0, q_1) = h_1^{q_1}, \quad \pi(h_1^{q_1}, q_2) = z_{q_1, q_2},$$
$$\boldsymbol{u}(z^{q_1, q_2}) = \begin{pmatrix} q_1(\kappa - q_1 - q_2) - q_1 c \\ q_2(\kappa - q_1 - q_2) - q_2 c \end{pmatrix}.$$

So this game is **huge** – it is shallow but very wide.

Example 5.4 (Chess)**.** Surely, $N = \{1, 2\}$. Denoting Boards the set of all (appropriately encoded) board positions, we define $\mathscr{H} = B \times \{1, 2\}$ where

$$B = \left\{w \in \text{Boards}^+ | \text{ no board repeats } \geq 3 \text{ times in } w\right\}.$$

Tip

Here Boards^+ is the set of all non-empty sequences of boards

Surely, Z consists of all nodes (wb, i), here $b \in$ Boards, where either b is checkmate for player i, or i does not have a move in b, or every move of i in b leads to a board with three occurrences in w. Also, $\chi(wb, i)$ is the set of all possible moves of player i in wb and $\rho(wb, i) = i$. Then π is defined by $\pi((wb, i), a) = (wbb', 3 - i)$ where b' is obtained from b according to the move a. The initial board is $h_0 = (b_0, 1)$ and $u_j(wb, i) \in \{1, 0, -1\}$, where 1 means "win", 0 means "draw", and -1 means "loss" for player j.

5.2 Pure strategies

Let $G = (N, A, H, Z, \chi, \rho, \pi, h_0, \boldsymbol{u})$ be a perfect-information extensive-form game.

Definition 5.2. A **pure strategy** of player i in G is a function $s_i : H_i \to A$ such that for every $h \in H_i$ we have $s_i(h) \in \chi(h)$.

We denote by S_i the set of all pure strategies of player i in G. Denote by $S = S_1 \times \cdots \times S_n$ the set of all pure strategy profiles. Note that each pure strategy profile $\boldsymbol{s} \in S$ determines a unique path $w_{\boldsymbol{s}} = h_0 a_1 h_1 \ldots h_{k-1} a_k h_k$ form h_0 to a terminal node h_k by

$$a_j = s_{\rho(h_{j-1})}(h_{j-1}) \quad \forall 0 < j \leq k,$$

so we can denote by $O(\boldsymbol{s})$ the terminal node reach by $w_{\boldsymbol{s}}$.

Abusing notation a bit, we denote by $u_i(\boldsymbol{s})$ the value $u_i(O(\boldsymbol{s}))$ of the payoff for player i when the terminal node $O(\boldsymbol{s})$ is reached using strategies of $\boldsymbol{s}$.

> **Note**
>
> The tree is assumed to have **finite depth** here.

Example 5.5. Recall the game

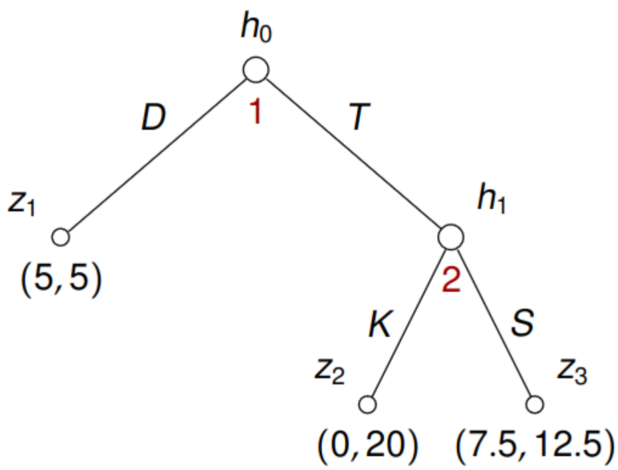

Figure 5.3: Trust game

A pure strategy profile (s_1, s_2) where

$$s_1(h_0) = T \quad \text{and} \quad s_2(h_1) = K.$$

5.3 Extensive-Form vs Strategic-Form

The extensive-form game G determines the corresponding strategic-form game $\hat{G} = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$.

> **Note**
>
> Here note that the set of players N and the sets of pure strategies S_i are the same in G and in the corresponding game.
> The payoff functions u_i in $\hat{G}$ are understood as functions on the pure strategy profiles of $S = S_1 \times \cdots \times S_n$.

With this definition, we may apply all solution concepts and algorithms developed for strategic-form games to the extensive form games.

> **Note**
>
> We often consider the extensive-form to be only a different way of representing the corresponding strategic-form game and do not distinguish between them.

There are some issues, namely whether all notions from the strategic-form area make sense in the extensive-form. Also, naive application of algorithms may result in unnecessarily high complexity.

> **!** Important
>
> For now, let us consider pure strategies only!

Example 5.6. Recall the game

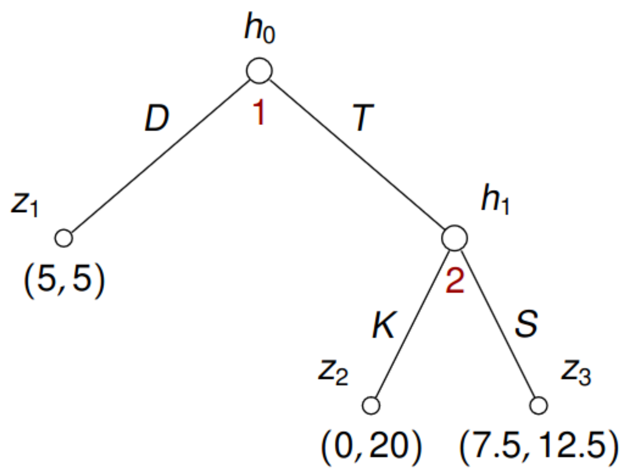

Figure 5.4: Trust game

Is any strategy strictly (weakly, very weakly) dominant? If player 1 distrusts (D), player 2 might just as well play any of the two strategies (so there are **no strictly dominant** strategies). But playing K to h_1 **weakly dominates** playing S.

5.3.1 Criticism of Nash Equilibria

Example 5.7. Consider a game

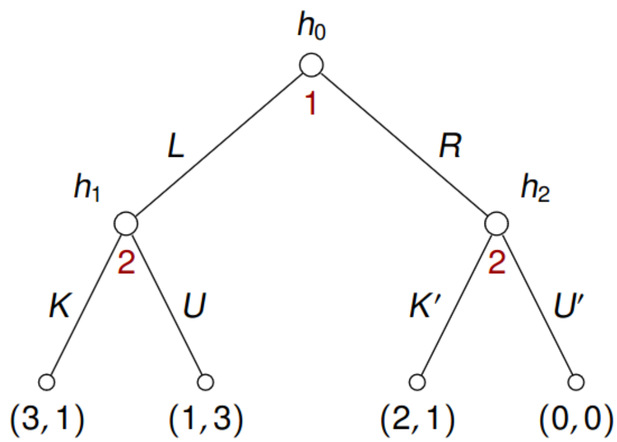

Find all pure strategies of both players – for player 2, there are **4 strategies** (KK', KU', UK', UU' with the usual notation $h_1 \mapsto K, h_2 \mapsto K'$ for KK'). Is any strategy (strictly, weakly, very weakly) dominant? For player 2, UK' is weakly dominant. Is any strategy (strictly, weakly, very weakly) dominated? Is any strategy never best response? Are there Nash equilibria in pure strategies?

Table 5.1: Table for extensive-form game

	KK'	KU'	UK'	UU'
L	(3, 1)	(3, 1)	(1, 3)	(1, 3)
R	(2, 1)	(0, 0)	(2, 1)	(0, 0)

Writing it in a table converts it into a strategic-form game. As such, it is much easier to see Nash equilibria R, UK' and L, UU'. When we examine (L, UU'), we obtain:

- player 2 threats to play U' in h_2,
- as a result, player 1 plays L,
- player 2 reacts to L by playing the best response, i.e. U.

However, the threat is not credible, once a play reaches h_2, a rational player 2 chooses K'. Now examine (R, UK'), which is sensible in the following sense

- player 2 plays the best response in both h_1 and h_2,
- player 1 plays the "best response" in h_0 assuming that player 2 will play his best responses in the future.

This equilibrium is called *subgame perfect*.

Given $h \in \mathscr{H}$, we denote $\mathscr{H}^h$ the set of all nodes reachable from h.

Definition 5.3 (Subgame)**.** A subgame G^h of G rooted at $h \in \mathscr{H}$ is the restriction of G to nodes reachable from h in the game tree. More precisely,

$$G^h = (N, A, H^h, \chi^h, \rho^h, \pi^h, h, \boldsymbol{u}^h)$$

where $H^h = H \cap \mathscr{H}^h$, $Z^h = Z \cap \mathscr{H}^h$, χ^h and ρ^h are restrictions of χ and ρ to H^h, respectively. Moreover

- π^h is defined for $h' \in H^h$ and $a \in \chi^h(h')$ by $\pi^h(h', a) = \pi(h', a)$;
- each u_i^h is a restriction of u_i to Z^h.

> **Tip**
>
> Given a function $f : A \to B$ and $C \subsetneq A$, a restriction of f to C is a function $g : C \to B$ such that $g(x) = f(x)$ for all $x \in C$

Definition 5.4 (Subgame Perfect Equilibrium)**.** A **subgame perfect equilibrium (SPE)** in pure strategies is a pure strategy profile $\boldsymbol{s} \in S$ such that for any subgame G^h of G, the restriction of $\boldsymbol{s}$ to H^h is a Nash equilibrium in pure strategies in G^h.

> **Note**
>
> A restriction of $\boldsymbol{s} = (s_1, \ldots, s_n) \in S$ to H^h is a strategy profile $\boldsymbol{s}^h = (s_1^h, \ldots, s_n^h)$ where $s_i^h(h') = s_i(h')$ for all $i \in N$ and all $h' \in H_i \cap H^h$.

Example 5.8. Recall the Stackelberg competition Example 5.3. Player 1 chooses q_1, and we know that the best response of player 2 is $q_2 = (\theta - q_1)/2$ where $\theta = \kappa - c$. Then

$$u_1(z^{q_1,q_2}) = q_1(\theta - q_1 - \theta/2 - q_1/2) = (\theta/2)q_1 - q_1^2/2$$

which is maximized by $q_1 = \theta/2$, giving $q_2 = \theta/4$. So

$$u_1(z^{q_1,q_2}) = \theta^2/8, \qquad u_2(z^{q_1,q_2}) = \theta^2/16$$

Note that firm 1 has an advantage as a leader.

5.4 Backward Induction

An algorithm for computing SPE for finite perfect-information extensive-form games.

Definition 5.5 (Backward Induction)**.** We inductively "attach" to every node h a pure strategy profile $\boldsymbol{s}^h = (s_1^h, \ldots, s_n^h)$ in G^h, together with a vector of expected payoffs $\boldsymbol{u}(h) = (u_1(h), \ldots, u_n(h))$.

- **Initially** - Attach to each terminal node $z \in Z$ the empty profile $\boldsymbol{s}^z = (\emptyset, \dots, \emptyset)$ and the payoff vector $\boldsymbol{u}(z) = (u_1(z), \dots, u_n(z))$.
- **While** (there is an unattached node h with all children attached):
  1. Let K be the set of all children of h;
  2. Let
  $$h_{\max} \in \operatorname*{argmax}_{h' \in K} u_{\rho(h)}(h');$$
  3. Attach to h a strategy profile $\boldsymbol{s}^h$ where
     - $\pi(h, s^h_{\rho(h)}(h)) = h_{\max}$;
     - for all $i \in N$ and all $h' \in H_i^h \setminus \{h\}$ (at the same time $h' \in H^{\bar{h}} \cap H_i$ for an appropriate $\bar{h} \in K$) define $s_i^h(h') = s_i^{\bar{h}}(h')$ where $\bar{h} \in K$;
  4. Attach to h the vector of expected payoffs $\boldsymbol{u}(h) := \boldsymbol{u}(h_{\max})$.

Example 5.9. Recall the game

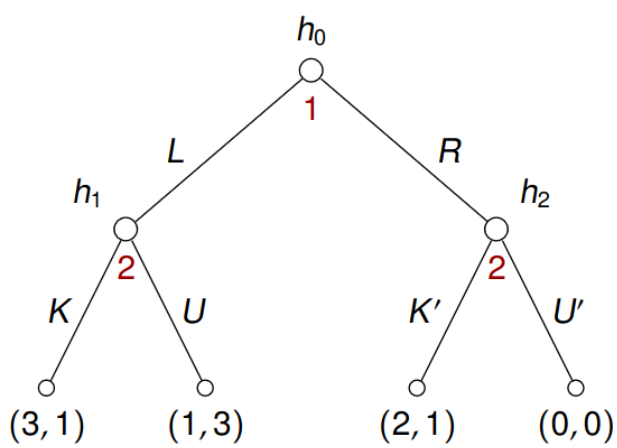

Then clearly
$$\boldsymbol{s}^{z_1} = \boldsymbol{s}^{z_2} = \boldsymbol{s}^{z_3} = \boldsymbol{s}^{z_4} = (\emptyset, \emptyset).$$

Now for h_1, where the second player makes a choice (and his better choice is U getting him $(1,3)$), we get
$$\boldsymbol{s}^{h_1} = (s_1^{h_1}, s_2^{h_2}) = (\emptyset, \{(h_1, U)\}) \implies \boldsymbol{u}(h_1) = (1,3)$$

and similarly
$$\boldsymbol{s}^{h_2} = (\overbrace{\emptyset}^{s_1^{h_2}}, \overbrace{\{(h_2, K')\}}^{s_2^{h_2}}) \implies \boldsymbol{u}(h_2) = (2,1).$$

Finally considering h_0, where $\rho(h_0) = 1$, we see that
$$\boldsymbol{s}^{h_0} = (\{(h_0, R)\}, \{(h_1, U), (h_2, K')\}) \implies \boldsymbol{u}(h_0) = (2,1).$$

Theorem 5.1. *For every finite perfect-information extensive-form game and for each node h the attached $\boldsymbol{s}^h$ is a SPE and the attached vector $\boldsymbol{u}(h)$ satisfies $\boldsymbol{u}(h) = \boldsymbol{u}(\boldsymbol{s}^h) = (u_1(\boldsymbol{s}^h), \dots, u_n(\boldsymbol{s}^h))$.*

Proof. We will lead the proof by induction. In any terminal node z no player has any choice, thus empty strategies make a SPE with payoffs $\boldsymbol{u}(z)$.

Now assume that h is being processed in the while loop. Denote by $\bar{\boldsymbol{s}}^h$ a profile obtained from $\boldsymbol{s}^h$ by changing the strategy of player i. We split the situation into two different cases. First, we assume that player i does **not** control h, i.e. $\rho(h) \neq i$. Let $\bar{\boldsymbol{s}}^{h_{\max}}$ be the restriction of $\bar{\boldsymbol{s}}^h$ to the subgame rooted in $h_{\max}$, see Figure 5.5.

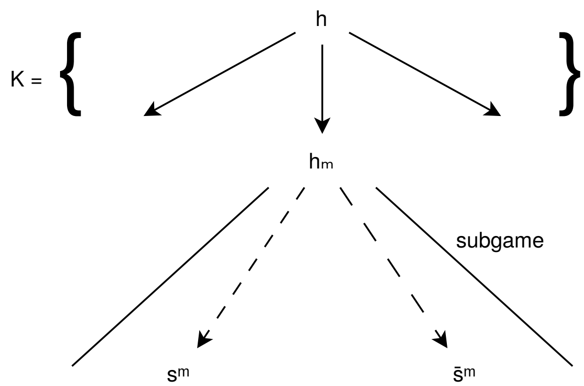

Figure 5.5: Case when player i does not control h

Here $\mathrm{h_m}$ denotes $h_{\max}$ and $\mathrm{s^m}$, and $\bar{\mathrm{s}}^{\mathrm{m}}$, the profiles $\boldsymbol{s}^{h_{\max}}$, and $\bar{\boldsymbol{s}}^{h_{\max}}$ respectively. By induction we then get

$$u_i(\boldsymbol{s}^h) = u_i(\boldsymbol{s}^{h_{\max}}) \geq u_i(\bar{\boldsymbol{s}}^{h_{\max}}) = u_i(\bar{\boldsymbol{s}}^h).$$

Second, we assume $i = \rho(h)$ and denote by $\bar{h} = \bar{s}^h_{\rho(h)}$. Let $\bar{\boldsymbol{s}}^{\bar{h}}$ be the restriction of $\bar{\boldsymbol{s}}^h$ to the subgame rooted in $\bar{h}$. Then also, see Figure 5.6,

$$u_i(\bar{\boldsymbol{s}}^h) = u_i(\bar{\boldsymbol{s}}^{\bar{h}}) \leq u_i(\boldsymbol{s}^{\bar{h}}) \leq u_i(\boldsymbol{s}^{h_{\max}}) = u_i(\boldsymbol{s}^h). \tag{5.1}$$

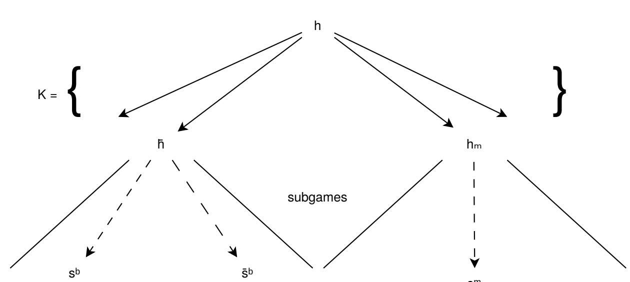

Figure 5.6: Case when player i controls h, $\mathrm{s^b}$ denotes $\boldsymbol{s}^{\bar{h}}$ and $\bar{\mathrm{s}}^{\mathrm{b}}$ profile $\bar{\boldsymbol{s}}^{\bar{h}}$

The first inequality in (5.1) is true by the induction step and the second holds by the definition of $h_{\max}$. In both cases the deviation of player i leads to a smaller or equal playoff. Hence $\boldsymbol{u}(\boldsymbol{s}^h) = \boldsymbol{u}(\boldsymbol{s}^{h_{\max}}) = \boldsymbol{u}(h_{\max}) = \boldsymbol{u}(h)$. This concludes the proof. □

Example 5.10 (Chess – continued)**.** Recall that in the model of chess, the payoffs were from $\{1, 0, -1\}$ and $u_1 = -u_2$ (i.e. chess is a zero-sum game). By Theorem 5.1, there is a SPE in pure strategies $\boldsymbol{s}^* = (s_1^*, s_2^*)$. However, then one of the following holds:

1. white has a winning strategy (if $u_1(\boldsymbol{s}^*) = 1$ and thus $u_2(\boldsymbol{s}^*) = -1$);
2. black has a winning strategy (if $u_2(\boldsymbol{s}^*) = 1$ and thus $u_1(\boldsymbol{s}^*) = -1$);
3. both players have strategies to force draw (if $u_1(\boldsymbol{s}^*) = 0$ and thus $u_2(\boldsymbol{s}^*) = 0$).

Now the question arises what is the right answer, but, in truth, nobody knows yet, as the tree of the game is too big. Even simplifying to trees only about 200 edges deep and with approximately 5 moves per node on average, we get a total count of 5^{200} nodes!

Recall that in the second step of Definition 5.5, the algorithm may choose **an arbitrary** $h_{\max} \in \operatorname{argmax}_{h' \in K} u_{\rho(h)}(h')$ and always obtain a SPE. Thus to compute all SPEs, the algorithm may systematically search through all possible choices of $h_{\max}$ throughout the induction.

Also one can realize that backward induction, see Definition 5.5, is too inefficient, as it unnecessarily searches through the whole tree. There are better methods mitigating this problem, e.g. α-β-pruning.

5.4.1 Criticism of Subgame Perfect Equilibria

Example 5.11. Consider the following game, called **centipede**:

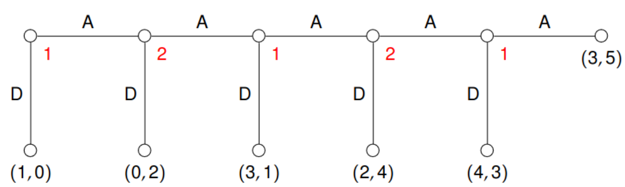

Figure 5.7: The centipede game

By backward induction, we can obtain that the SPE in pure strategies is (DDD, DD) – this should be at least a little bit weird. There are serious issues here:

- In a laboratory setting, people usually play A for several steps.
- There is a theoretical problem: Imagine, that you are player 2. What would you do when player 1 chooses A in the first step? The SPE analysis says that you should go down, but the same analysis also says that the situation you are in cannot appear :-)

6 Mixed and Behavioral Strategies

6.1 Introduction

Assume two players and a **finite** extensive-form game G

Definition 6.1. A **mixed strategy** σ_i of player i in G is a mixed strategy of player i in the corresponding strategic-form game.

In other words, a mixed strategy σ_i in G is a probability distribution of S_i (recall that S_i is the set of all pure strategies, i.e. functions of the form $s_i : H_i \to A$). As before, we denote by Σ_i the set of all mixed strategies of player i.

Definition 6.2. A **behavioral (mixed) strategy** of player i in G is a function $\beta_i : H_i \to \Delta(A)$ such that for every $h \in H_i$ and every $a \in A$: $\beta_i(h)(a) = 0$ if $a \notin \chi(h)$.

> **!** Important
>
> For mixed strategy, we randomize **once** for our behaviors everywhere, which we then follow deterministically.
> On the other hand, for the behavioral strategies we randomize for each choice of the action we take.

Given a profile $\boldsymbol{\beta} = (\beta_1, \beta_2)$ of behavioral strategies, we denote by $P_{\boldsymbol{\beta}}(z)$ the probability of reaching $z \in Z$ when $\boldsymbol{\beta}$ is used, i.e.,

$$P_{\boldsymbol{\beta}}(z) = \prod_{l=1}^{k} \beta_{\rho(h_{l-1})}(h_{l-1})(a_l),$$

where $h_0 a_1 h_1 a_2 h_2 \ldots a_k h_k$ is the unique path from h_0 to $h_k = z$. We also define the **expected payoff** under the behavioral strategy $\boldsymbol{\beta}$ as $u_i(\boldsymbol{\beta}) := \sum_{z \in Z} P_{\boldsymbol{\beta}(z)} u_i(z)$.

6.2 Examples

Example 6.1. Consider a game given by a tree in Figure 6.1.

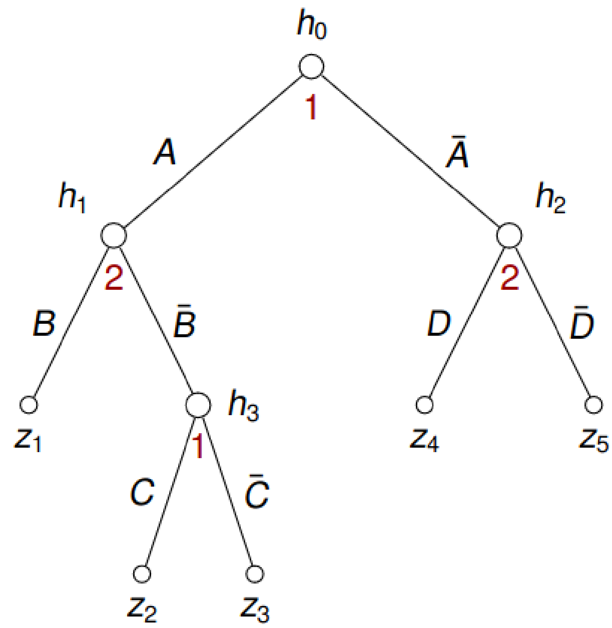

Figure 6.1: Mixed behavioral strategies

Then pure strategies of player 1 are $AC, A\bar{C}, \bar{A}C, \bar{A}\bar{C}$ and an example of a *mixed* strategy σ_1 of player 1 might be:

$$\sigma_1(AC) = \frac{1}{3}, \sigma_1(A\bar{C}) = \frac{1}{9}, \sigma_1(\bar{A}C) = \frac{1}{6}, \sigma_1(\bar{A}\bar{C}) = \frac{11}{18}.$$

On the other hand, an example of *behavioral* strategies of both players can be:

- player 1: $\beta_1(h_0)(A) = \frac{1}{3}$ and $\beta_1(h_3)(C) = \frac{1}{2}$;
- player 2: $\beta_2(h_1)(B) = \frac{1}{4}$ and $\beta_2(h_2)(D) = \frac{1}{5}$.

Then for example $P_{(\beta_1,\beta_2)}(z_2) = \frac{1}{3}\left(1 - \frac{1}{4}\right)\frac{1}{2} = \frac{1}{8}$.

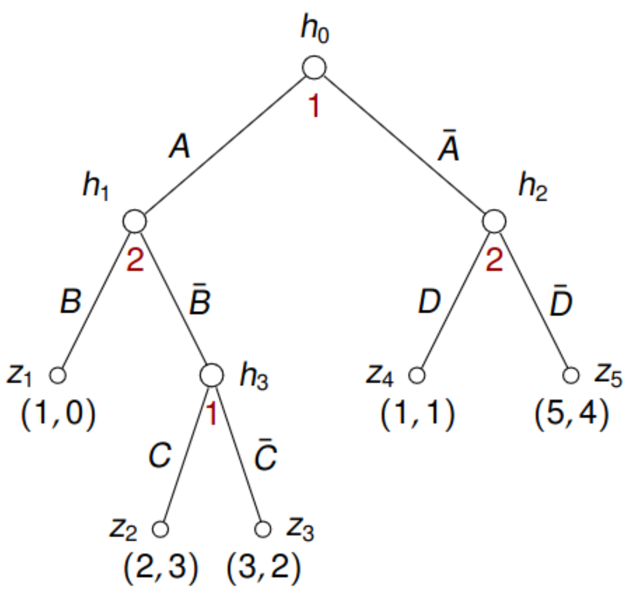

Figure 6.2: Game G with assigned payoffs

After assigning payoffs, see Figure 6.2, we can compute *expected payoff* of the behavioral strategy $\boldsymbol{\beta}$:

$$\begin{aligned} u_1(\boldsymbol{\beta}) &= P_{\boldsymbol{\beta}}(z_1) \cdot 1 + P_{\boldsymbol{\beta}}(z_2) \cdot 2 + P_{\boldsymbol{\beta}}(z_3) \cdot 3 + P_{\boldsymbol{\beta}}(z_4) \cdot 1 + P_{\boldsymbol{\beta}}(z_5) \cdot 5 \\ &= \frac{1}{3}\frac{1}{4}1 + \frac{1}{3}\frac{3}{4}\frac{1}{2}2 + \frac{1}{3}\frac{3}{4}\frac{1}{2}3 + \frac{2}{3}\frac{1}{5}1 + \frac{2}{3}\frac{4}{5}5 \approx 3.508. \end{aligned}$$

Each pure strategy can be seen as a behavioral strategy. Consider e.g. $s_1 : H_1 \to A$ defined by $s_1(h_0) = A$ and $s_1(h_3) = C$. The corresponding behavioral strategy β_1 would satisfy $\beta_1(h_0)(A) = \beta_1(h_3)(C) = 1$ (i.e. select actions chosen by s_1 with probability 1). Now given a behavioral strategy β_2 of player 2 defined by $\beta_2(h_1)(B) = \frac{1}{4}$ and $\beta_2(h_2)(D) = \frac{1}{5}$ we obtain

$$P_{(s_1,\beta_2)}(z_2) = P_{(\beta_1,\beta_2)}(z_2) = 1\left(1 - \frac{1}{4}\right)1 = \frac{3}{4}.$$

6.3 Equivalence

Let $\boldsymbol{\alpha} = (\alpha_1, \alpha_2)$ be a strategy profile where each α_i is either mixed or behavioral.

The game is then played as follows:

- if α_1 is mixed, select randomly a pure strategy β_1 according to α_1, else $\beta_1 := \alpha_1$;
- if α_2 is mixed, select randomly a pure strategy β_2 according to α_2, else $\beta_2 := \alpha_2$;
- play (β_1, β_2) and collect payoffs.

Denote the resulting payoffs by $u_1(\boldsymbol{\alpha})$ and $u_2(\boldsymbol{\alpha})$.

Lemma 6.1. *For every mixed/behavioral strategy α_1 of player 1 there is a behavioral/mixed strategy α_1' such that for every mixed/behavioral strategy α_2 we have that $u_i(\alpha_1, \alpha_2) = u_i(\alpha_1', \alpha_2)$ for $i \in \{1, 2\}$.*

Example 6.2. Consider the following game:

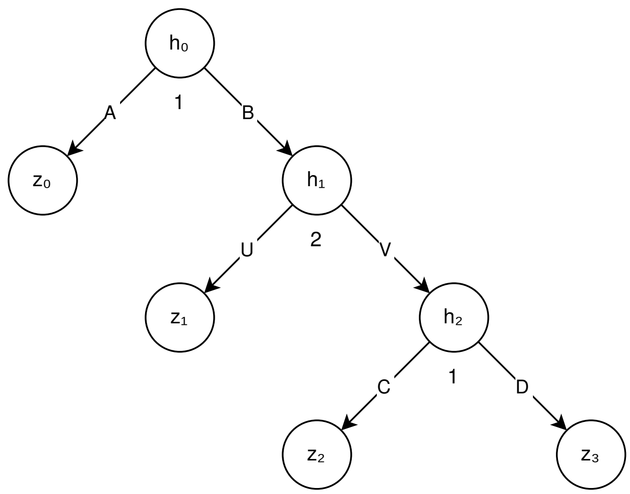

Let us first focus on "converting" *behavioral* strategies to *mixed* ones. Let $\boldsymbol{\alpha} = (\alpha_1, \alpha_2)$ be a behavioral strategy profile, that is

- $\alpha_1(h_0)(A) = p$ and $\alpha_1(h_2)(C) = q$;
- $\alpha_2(h_1)(U) = u$.

Then the equivalent mixed strategy α_1' of player 1 is:

$$\alpha_1'(AC) = pq, \qquad \alpha_1'(AD) = p(1-q),$$
$$\alpha_1'(BC) = (1-p)q, \qquad \alpha_1'(BD) = (1-p)(1-q),$$

because

$$P_{(\alpha_1,\alpha_2)}(z_2) = (1-p)u = (\overbrace{(1-p)q}^{\alpha_1'(BC)} + \overbrace{(1-p)(1-q)}^{\alpha_1'(BD)}) \cdot u = P_{(\alpha_1',\alpha_2)}(z_2)$$

and similarly for the other terminal nodes.

Let us now turn our attention to "converting" *mixed* strategies to *behavioral ones.* Let α_1 be a mixed strategy of player 1:

$$\alpha_1(AC) = e_{AC}, \qquad \alpha_1(AD) = e_{AD},$$
$$\alpha_1(BC) = e_{BC}, \qquad \alpha_1(BD) = e_{BD}.$$

Then we can construct an equivalent behavioral strategy α_1' as

$$\alpha_1'(h_0)(A) = e_{AC} + e_{AD}, \qquad \alpha_1'(h_0)(B) = e_{BC} + e_{BD},$$
$$\alpha_1'(h_2)(C) = \frac{e_{BC}}{e_{BC} + e_{BD}}, \qquad \alpha_1'(h_2)(D) = \frac{e_{BD}}{e_{BC} + e_{BD}},$$

where $\frac{e_{BC}}{e_{BC}+e_{BD}}$ can be interpreted as a conditional probability $P(C|BC \cup BD)$. We can then check our calculation:

$$P_{(\alpha_1,\alpha_2)}(z_3) = e_{BC} \cdot (1-u) = \overbrace{(e_{BC}+e_{BD})}^{\alpha_1'(h_0)B} \cdot \overbrace{\frac{e_{BC}}{e_{BC}+e_{BD}}}^{\alpha_1'(h_2)(C)} \cdot (1-u) = P_{(\alpha_1',\alpha_2)}(z_3)$$

7 Imperfect-Information Games

7.1 Introduction

Is it possible to model Matching pennies using extensive-form games?

	H	T
H	$1,-1$	$-1,1$
T	$-1,1$	$1,-1$

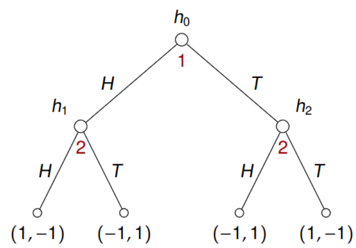

Figure 7.1: Matching pennies

The problem is that player 2 is "perfectly" informed about the choice of player 1. In particular, there are pure Nash equilibria (H, TH) and (T, TH) in the extensive-form game as opposed to the strategic-form, where there is none. Reversing the order of players does not help.

Thus, we need to extend the formalism to be able to hide some information about previous moves.

Matching pennies can be modeled using an **imperfect-information** extensive-form game:

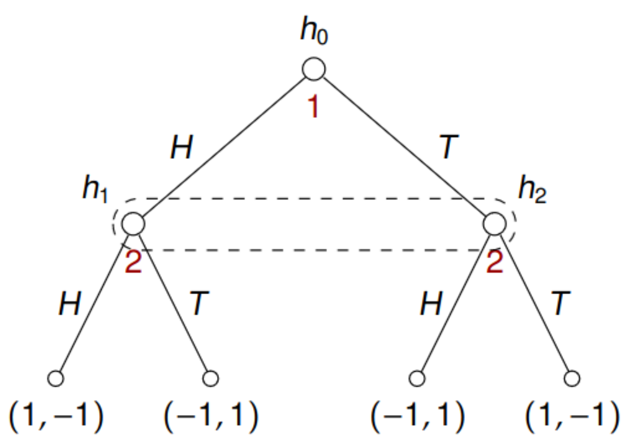

Figure 7.2: Imperfect-information Matching pennies

Here h_1 and h_2 belong to the same information set of player 2. As a result, player 2 is not able to distinguish between h_1 and h_2. So even though players do not move simultaneously, the information player 2 has about the current situation is the same as in the simultaneous case

> **Note**
>
> There must be the same set of actions from all nodes of the information set. Otherwise, the player would be able to deduce information about the current state.

> **Caution**
>
> **No history** is inherently encoded in the game (so if we want to provide past moves of the player as input for his decision process, we have to encode it in the notes themselves). Otherwise, the player could be able to deduce more information about the current state of the game.

7.2 Definition of Imperfect-Information Game

Definition 7.1. An **imperfect-information extensive-form** game is a tuple $G_{\text{imp}} = (G_{\text{perf}}, \boldsymbol{I})$ where

- $G_{\text{perf}} = (N, A, H, Z, \chi, \rho, \pi, h_0, \boldsymbol{u})$ is a *perfect-information extensive-form* game (called the **underlying game**),
- $\boldsymbol{I} = (I_1, \dots, I_n)$ where for each $i \in N = \{1, \dots, n\}$

  $$I_i = \{I_{i,1}, \dots, I_{i,k_i}\}$$

  is a collection of **information sets** for player i that satisfies

  - $\bigcup_{j=1}^{k_i} I_{i,j} = H_i$ and $I_{i,j} \cap I_{i,k} = \emptyset$ for $j \neq k$ (i.e. I_i is a partition of H_i, see Definition 5.1);
  - for all $h, h' \in I_{i,j}$, we have $\rho(h) = \rho(h')$ and $\chi(h) = \chi(h')$ (i.e., nodes from the same information set are owned by the same player and have the same sets of enabled actions).

Given $h \in H$, we denote by $I(h)$ the information set $I_{i,j}$ containing h. Also, given an information set $I_{i,j}$, we denote by $\chi(I_{i,j})$ the of all action enabled in some (and hence all, per Definition 7.1) nodes of $I_{i,j}$.

Now we define the set of pure, mixed, and behavioral strategies in G_{imp} as subsets of pure, mixed, and behavioral strategies, resp., in G_{perf} that respect the information sets. Let $G_{\text{imp}} = (G_{\text{perf}}, \boldsymbol{I})$ be an imperfect-information extensive-form game where $G_{\text{perf}} = (N, A, H, Z, \chi, \rho, \pi, h_0, \boldsymbol{u})$.

Definition 7.2. A **pure strategy** of player i in G_{imp} is a pure strategy s_i in G_{perf} such that for all $j = 1, \dots, k$ and all $h, h' \in I_{i,j}$ holds $s_i(h) = s_i(h')$.

💡 Tip

Note that each s_i can also be seen as a function $s_i : I_i \to A$ such that for every $I_{i,j} \in I_i$ we have that $s_i(I_{i,j}) \in \chi(I_{i,j})$.

As before, we denote by S_i the set of all pure strategies of player i in G_{imp}, and by $S = S_1 \times \cdots \times S_n$ the set of all pure strategy profiles. As in the perfect-information case we have a corresponding strategic-form game $\bar{G}_{\text{imp}} = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$.

7.3 Examples

Example 7.1 (Matching Pennies)**.** Recall the game of matching pennies:

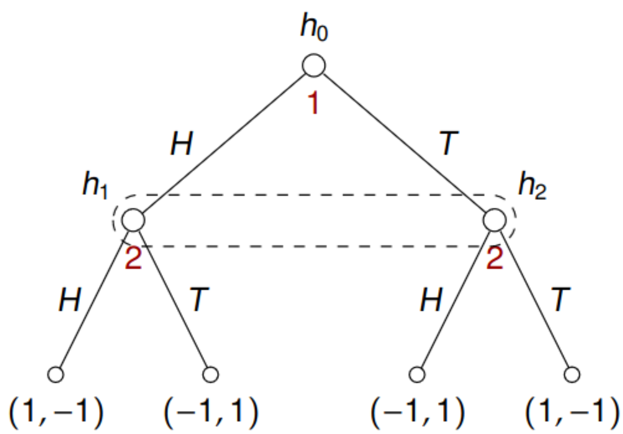

Figure 7.3: Matching pennies game

Here $I_1 = \{I_{1,1}\}$ where $I_{1,1} = \{h_0\}$ and $I_2 = \{I_{2,1}\}$ where $I_{2,1} = \{h_1, h_2\}$. An example of pure strategies might be:

- $s_1(I_{1,1}) = H$ which describes strategy $s_1(h_0) = H$;
- $s_2(I_{2,1}) = T$ which describes strategy $s_2(h_1) = s_2(h_2) = T$ (it is also sufficient to specify $s_2(h_1) = T$ since then by definition $s_2(h_2) = T$).

Thus we really only have strategies H, T for player 1 and H, T for player 2.

Example 7.2. Consider now a game given by:

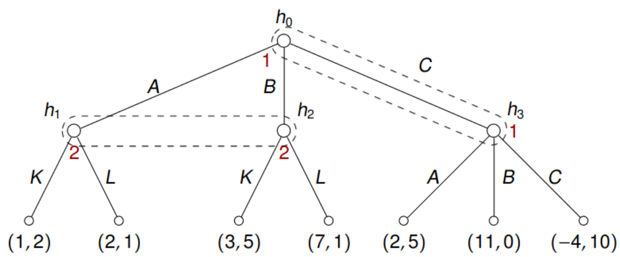

Note that $I_1 = \{I_{1,1}\}$ where $I_{1,1} = \{h_0, h_3\}$ and that $I_2 = \{I_{2,1}\}$ where $I_{2,1} = \{h_1, h_2\}$. The pure strategies in this example are

- $s(I_{1,1}) \in \{A, B, C\}$ for player 1;
- $s(I_{2,1}) \in \{K, L\}$ for player 2.

Playing A or B from h_3 is unreachable (or illegal/unfeasible) in this game with pure and mixed strategies.

7.4 Subgame-perfect Equilibria in Imperfect-Information Games

Now we shall turn our attention to the following game: h₂

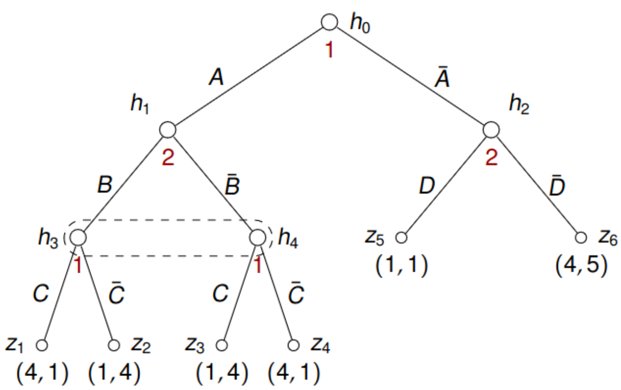

What do we designate as subgames to allow the backward induction, with which we could calculate subgame-perfect equilibria (or more precisely their equivalents here)? Only subtrees rooted in h_1, h_2, and h_0 (together will all subtrees rooted in *terminal* nodes) seem reasonable. Note that subtrees rooted in h_3 and h_4 cannot be considered as *"independent"* subgames because their individual solution cannot be combined to a single best response in the information set $\{h_3, h_4\}$.

Let $G_{\text{imp}} = (G_{\text{perf}}, \boldsymbol{I})$ be an imperfect-information extensive-form game where $G_{\text{perf}} = (N, A, H, Z, \chi, \rho, \pi, h_0, \boldsymbol{u})$ is the underlying perfect-information extensive-form game.

Let us denote by H_{proper} the set of all $h \in H$ that satisfy the following:
For every h' reachable from h (that includes the node h itself), we have that either all nodes of $I(h')$ are reachable from h, or no node of $I(h')$ is reachable from h.

Note

Intuitively, $h \in H_{\text{proper}}$ iff every information set $I_{i,j}$ is either completely contained in the subtree rooted in h, or no node of $I_{i,j}$ is contained in the subtree.

Definition 7.3. For every $h \in H_{\text{proper}}$ we define a subgame G^h_{imp} to be the imperfect information game $(G^h_{\text{perf}}, \boldsymbol{I}^h)$ where $\boldsymbol{I}^h$ is the restriction of $\boldsymbol{I}$ to H^h.

Tip

Note that as subgames of G_{imp} we consider only subgames of G_{perf} that respect the information sets, i.e. are rooted in nodes of H_{proper}.

Definition 7.4. A strategy profile $\boldsymbol{s} \in S$ is a *subgame-perfect equilibrium* (SPE) if $\boldsymbol{s}^h$ is a Nash equilibrium in every subgame G^h_{imp} of G_{imp} (here $h \in H_{\text{proper}}$).

Important

The way we generalized subgame-perfect equilibria is **not** the only one. But others are more complicated and use some kind of randomization.

7.4.1 Backwards Induction

Now we can generalize the backward induction to imperfect-information games, as we hypothesized, along the following lines:

1. As in the perfect-information case, the goal is to label each node $h \in H_{\text{proper}} \cup Z$ with a SPE $\boldsymbol{s}^h$ and a vector of payoffs $\boldsymbol{u}(h) = (u_1(h), \dots, u_n(h))$ for individual players according to $\boldsymbol{s}^h$.
2. Starting with terminal nodes, the labeling proceeds bottom up. Terminal nodes are labeled similarly as in the perfect-information case.
3. Consider $h \in H_{\text{proper}}$, let K be the set of all $h' \in \left(H_{\text{proper}} \cup Z\right) \setminus \{h\}$ that are h's **closest descendants out of** $H_{\text{proper}} \cup Z$, i.e. $h' \in K \iff h' \neq h$ is reachable from h and the unique path from h to h' visits only nodes of $\mathscr{H} \setminus H_{\text{proper}}$ (except the first and the last node). For every $h' \in K$ we already computed a SPE $\boldsymbol{s}^{h'}$ in $G^{h'}_{\text{imp}}$ and the vector of corresponding payoffs $\boldsymbol{u}(h')$.
4. Now consider all nodes of K as terminal nodes where each $h' \in K$ has payoffs $\boldsymbol{u}(h')$. This gives a new game in which we compute an equilibrium $\bar{\boldsymbol{s}}^h$ together with the vector $\boldsymbol{u}(h)$. The equilibrium $\boldsymbol{s}^h$ is then obtained by "concatenating" $\bar{\boldsymbol{s}}^h$ with all $\boldsymbol{s}^{h'}$, here $h' \in K$, in the subgames $G^{h'}_{\text{imp}}$ of G^h_{imp}.

7.4.2 Examples

Example 7.3 (Mutually Assured Destruction)**.** This example is an adaptation of the *Analysis of Cuban Missile Crisis of 1962* (as described in *Games for Business and Economics* by R. Gardner):

- The crisis started with the United States' discovery of Soviet nuclear missiles in Cuba.
- The USSR then backed down, agreeing to remove the missiles from Cuba, which suggests that US had a credible threat "if you don't back off we both pay dearly".

But could this be a credible threat? We shall model this situation as an extensive-form game:

- First, player 1 (US) chooses to either **i**gnore the incident (I), resulting in maintenance of status quo (payoffs $(0, 0)$), or **e**scalate the situation (E).
- Following escalation by player 1, player 2 (USSR) can **b**ack down (B), causing it to lose face (payoffs $(10, -10)$), or it can choose to proceed to a **n**uclear confrontation (N).
- Upon this choice, the players play a simultaneous-move game in which they can either **r**etreat (R), or choose **d**oomsday (D).
  - If both retreat, the payoffs are $(-5, 5)$, a small loss due to a mobilization process.
  - If either of them chooses doomsday, then the world destructs and payoffs are $(-100, -100)$.

This game can be re-written into the following tree:

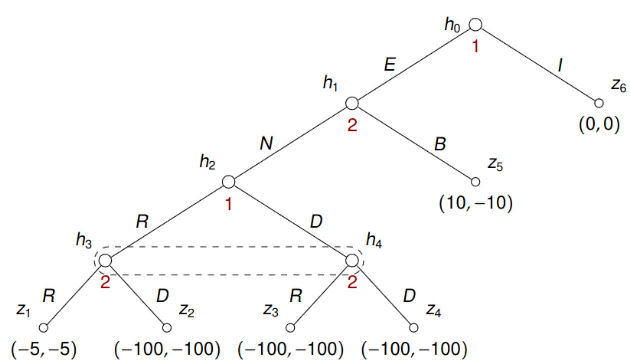

Figure 7.4: Mutually assured destruction

First and foremost, one can solve $G_{\text{imp}}^{h_2}$ (a strategic-form game). Then $G_{\text{imp}}^{h_1}$ by solving a game rooted in h_1 with terminal nodes h_2 and z_5 (payoffs in h_2 correspond to an equilibrium in $G_{\text{imp}}^{h_2}$). Finally, one solves G_{imp} by solving a game rooted in h_0 with terminal nodes h_1 and z_6 (payoffs in h_1 have been computed in the previous step). This produces 2 SPEs $\boldsymbol{s}_1 = ((I, R), (N, R))$ and $\boldsymbol{s}_2 = ((E, D), (B, D))$.

7.5 Mixed and Behavioral Strategies

Definition 7.5. A **mixed strategy** σ_i of player i in G_{imp} is a mixed strategy of player i in the corresponding strategic-form game $\bar{G}_{\text{imp}} = (N, (S_i)_{i \in N}, u_i)$. As before, we denote by Σ_i the set of all mixed strategies of player i.

Do not forget that in the corresponding game $\bar{G}_{\text{imp}}$ any strategy $s_i \in S_i$ iff s_i is a pure strategy that assigns the same action to all nodes of every information set. Hence each $s_i \in S_i$ can be seen as a function $s_i : I_i \to A$.

Definition 7.6. A **behavioral strategy** of player i in G_{imp} is a behavioral strategy β_i in G_{perf} such that for all $j = 1, \ldots, k_i$ and all $h, h' \in I_{i,j} : \beta_i(h) = \beta_i(h')$.

Here each β_i can be seen as a function $\beta_i : I_i \to \Delta(A)$ such that for all $I_{i,j} \in I_i$ we have $\text{supp}\left(\beta_i(I_{i,j})\right) \subseteq \chi(I_{i,j})$.

> **Note**
>
> For behavioral strategy, the distribution of probabilities always stays the same across the information set, unlike in the mixed strategy case, when the randomly selected pure strategy stays the same across the strategy set

Example 7.4. Consider a game

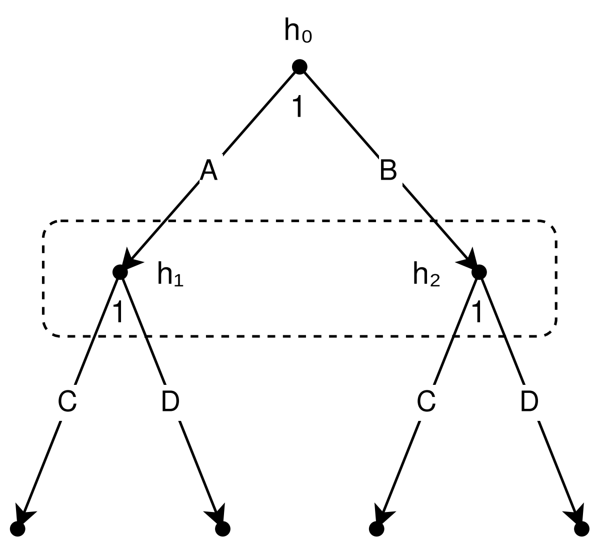

Figure 7.5: Behavioral ≠ mixed strategies

One mixed strategy is

$$\left(\frac{1}{2}(AC), 0(BD), 0(BC), \frac{1}{2}(BD)\right),$$

but this mixed strategy cannot be expressed using behavioral strategies.

> **Note**
>
> Here $ACC, ACD, ADC, ADD, BCC, BCD, BDC, BDD$ are pure strategies in perfect information. The information set $\{h_1, h_1'\}$ restricts us to AC, BC, AD, BD, where the second strategy refers to a strategy on the information set.
> Also, notice that *mixed* strategies don't really make sense in real life too much.

7.6 Perfect Recall

Example 7.5. Consider the following game of the so-called absent-minded driver:

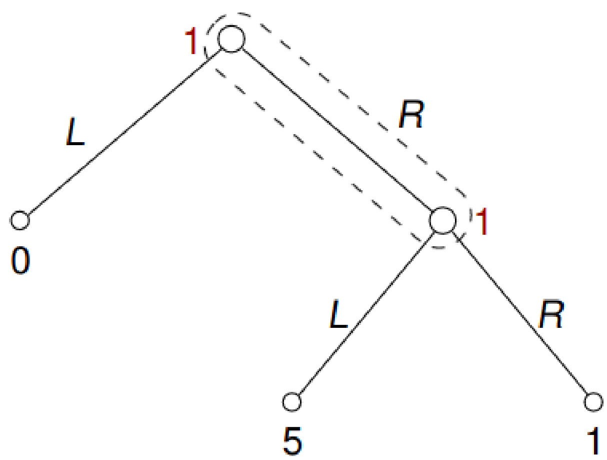

Note that there is only one player: A driver who has to take a turn at a particular junction. There are two identical junctions, the first one leads to a wrong neighborhood where the driver gets completely lost (payoff 0), and the second one leads home (payoff 5). If the driver misses both, there is a long way home (payoff 1). The problem is that after missing the first turn, the driver forgets that he missed the turn.

Notice that the behavioral strategy $\beta_1(I_{1,1})(L) = \frac{1}{2}$ has the expected payoff $\frac{3}{2}$. Also, it can be shown that no mixed strategy gives a larger payoff than 1 since no pure strategy ever reaches the terminal node with a payoff 5.

We say that player i has **perfect recall** in G_{imp} if the following holds:

- Every information set of player i (i.e., *his own*) intersects every path from the root h_0 to a terminal node at most once (so no absent-minded driver situation occurs, see Example 7.5).
- Every two paths from the root that end in the same information set of player i

- pass through the same information sets of player i,
- and in the same order,
- and in every such information set the two paths choose the same actions.

> ⚠ Warning
>
> These two paths, however, may pass through *different* information sets of other players and other players may choose different actions along each of the paths!
> I.e. each information set J of player i determines the sequence of information sets of player i and actions taken by player i along any path reaching J.

Theorem 7.1 (Kuhn, 1953)**.** *Assuming perfect recall, every mixed strategy can be translated to a behavioral strategy (and vice versa) so that the payoff for the resulting strategy is the same in any mixed profile.*

8 Repeated Games

Recall the Prisoner's dilemma game

	C	S
C	$(-5, -5)$	$(0, -20)$
S	$(-20, 0)$	$(-1, -1)$

Imagine that the criminals are being arrested repeatedly. Can they somewhat reflect upon their experience in order to play "better"? In this chapter, we consider strategic-form games played repeatedly:

- for finitely many rounds, the final payoff of each player will be the average of payoffs from all rounds;
- infinitely many rounds, here we consider a discounted sum of payoff and the long-run average payoff.

We will also analyze Nash and subgame-perfect equilibria.

> **!** Important
>
> We stick with pure strategies only!

8.1 Finitely Repeated Games

Let $G = (\{1, 2\}, (S_1, S_2), (u_1, u_2))$ be a *finite* strategic-form game of two players. We shall use S^t to denote $S^t = \underbrace{S \times \cdots \times S}_{t \text{ times}} = \prod_{i=1}^{t} S$.

Definition 8.1. A T-**stage game** $G_{T\text{-rep}}$ based on G proceeds in T **stages** so that in a stage $t \geq 1$, players choose a strategy profile $\boldsymbol{s}^t = (s_1^t, s_2^t)$. After T stages, both players collect the average payoff $\sum_{t=1}^{T} u_i(\boldsymbol{s}^t)/T$.

Definition 8.2. A **history of length** $0 \leq t \leq T$ is a sequence $h = \boldsymbol{s}^1 \cdots \boldsymbol{s}^t \in S^t$ of t strategy profiles. Denote by $H(t)$ the set of all histories of length t and let ϵ represent the empty history.

A **pure strategy** for player i in a T-stage game $G_{T\text{-rep}}$ is a function

$$\tau_i : \bigcup_{t=0}^{T-1} H(t) \to S_i,$$

which for every possible history chooses a next step for player i.

Every strategy profile $\boldsymbol{\tau} = (\tau_1, \tau_2)$ in $G_{T\text{-rep}}$ induces a sequence of pure strategy profiles $w_{\boldsymbol{\tau}} = \boldsymbol{s}^1 \cdots \boldsymbol{s}^T$ in G so that $s_i^t = \tau_i(\boldsymbol{s}^1 \cdots \boldsymbol{s}^{t-1})$. Given a pure strategy profile $\boldsymbol{\tau}$ in $G_{T\text{-rep}}$ such that $w_{\boldsymbol{\tau}} = \boldsymbol{s}^1 \cdots \boldsymbol{s}^T$, define the payoffs $u_i(\boldsymbol{\tau}) = \sum_{t=1}^{T} u_i(\boldsymbol{s}^t)/T$.

Example 8.1. Again remember the Prisoner's dilemma:

	C	S
C	$(-5, -5)$	$(0, -20)$
S	$(-20, 0)$	$(-1, -1)$

This time, we shall assume it is a 3-stage game. Examples of histories might be:

$$\epsilon, (C, S), (C, S)(S, S), (C, S)(S, S)(C, C).$$

As we have $G_{3\text{-rep}}$ this time, the last history is final, which induces that this history sequence can be obtained with

$$\tau_1(\epsilon) = C, \qquad \tau_1((C, S)) = S, \qquad \tau_1((C, S)(S, S)) = C,$$
$$\tau_2(\epsilon) = S, \qquad \tau_2((C, S)) = S, \qquad \tau_2((C, S)(S, S)) = C.$$

Thus $w_{\boldsymbol{\tau}} = (C, S)(S, S)(C, C)$ and

$$u_1(\boldsymbol{\tau}) = (0 + (-1) + (-5))/3 = -2$$
$$u_2(\boldsymbol{\tau}) = (-20 + (-1) + -(5))/3 = -\frac{26}{3}$$

8.1.1 Finitely Repeated Games in Extensive-Form

Every T-stage game $G_{T\text{-rep}}$ can be defined as an imperfect-information extensive-form game $G_{\text{imp}}^{\text{rep}} = (G_{\text{perf}}^{\text{rep}}, \boldsymbol{I})$ such that $G_{\text{perf}}^{\text{rep}} = (\{1, 2\}, A, H, Z, \chi, \rho, \pi, h_0, \boldsymbol{u})$ where

- $A = S_1 \cup S_2$;
- $H = (S_1 \times S_2)^{<T} \cup (S_1 \times S_2)^{<T} \cdot S_1$ (intuitively, elements of $(S_1 \times S_2)^{\leq k}$ are possible histories and $(S_1 \times S_2)^{<k} \cdot S_1$ is used to simulate a simultaneous play of G letting player 1 choose first and player 2 second);
- $Z = (S_1 \times S_2)^T$;

- $\chi(\epsilon) = S_1$ and $\chi(h \cdot s_1) = S_2$ for $s_1 \in S_1$, and $\chi(h \cdot (s_1, s_2)) = S_1$ for $(s_1, s_2) \in S$;
- $\rho(\epsilon) = 1$ and $\rho(h \cdot s_1) = 2$ and $\rho(h \cdot (s_1, s_2)) = 1$;
- $\pi(\epsilon, s_1) = s_1$ and $\pi(h \cdot s_1, s_2) = h \cdot (s_1, s_2)$ and $\pi(h \cdot (s_1, s_2), s_1') = h \cdot (s_1, s_2) \cdot s_1'$;
- $h_0 = \epsilon$ and $u_i((s_1^1, s_2^1)(s_1^2, s_2^2) \cdots (s_1^T, s_2^T)) = \sum_{t=1}^{T} u_i(s_1^t, s_2^t)/T$.

Now the collection of information sets is defined as follows: Let $h \in H_1$ be a node of player 1, then

- there is exactly one information set of player 1 containing h as the only element,
- there is exactly one information set of player 2 containing all nodes of the form $h \cdot s_1$ where $s_1 \in S_1$.

> Tip
>
> Intuitively, in every round, player 1 has complete information about the results of past plays, thus player 1 chooses a pure strategy $s_1 \in S_1$, but player 2 is **not** informed about s_1, but still has complete information about results of all previous rounds. Hence player 2 chooses a pure strategy s_2 and both players are informed about the result.

8.1.2 Equilibria

Definition 8.3. A strategy profile $\boldsymbol{\tau} = (\tau_1, \tau_2)$ in a T-stage game $G_{T\text{-rep}}$ is a Nash equilibrium if for every $i \in \{1, 2\}$ and every τ_i' we have

$$u_i(\tau_1, \boldsymbol{\tau}_{-i}) \geq u_i(\tau_i', \boldsymbol{\tau}_{-i}).$$

To define a subgame-perfect equilibrium we use the following notation. Given a history $h = \boldsymbol{s}^1 \cdots \boldsymbol{s}^t$ and a strategy τ_i of player i, we define strategy τ_i^h in $(T-t)$-stage game based on G by

$$\tau_i^h(\bar{\boldsymbol{s}}^1 \cdots \bar{\boldsymbol{s}}^{\bar{t}}) = \tau_i(\overbrace{\boldsymbol{s}^1 \cdots \boldsymbol{s}^t}^{h} \bar{\boldsymbol{s}}^1 \cdots \bar{\boldsymbol{s}}^{\bar{t}})$$

for every sequence $\bar{\boldsymbol{s}}^1 \cdots \bar{\boldsymbol{s}}^{\bar{t}}$ (i.e. τ_i^h behaves as τ_i after h).

> Tip
>
> This can be sort of interpreted as a closure (from functional programming languages terminology).

Definition 8.4. A strategy profile $\boldsymbol{\tau} = (\tau_1, \tau_2)$ in a T-stage game $G_{T\text{-rep}}$ is a subgame-perfect Nash equilibrium (SPE) is for every history h the profile (τ_1^h, τ_2^h) is a Nash equilibrium in the $(T - |h|)$-stage game based on G.

Example 8.2. Consider now a T-stage game based on the Prisoner's dilemma:

	C	S
C	$(-5,-5)$	$(0,-20)$
S	$(-20,0)$	$(-1,-1)$

Surely, playing (C,C) (the Nash equilibrium in the strategic-form game) for every $t \le T$ gives us a SPE in the T-stage. But is this the only SPE in this game?

Theorem 8.1. *Let G be an arbitrary finite strategic-form game. If G has a unique Nash equilibrium, then playing this equilibrium every time is the unique SPE (but not necessarily a unique Nash equilibrium) in the T-stage game based on G.*

Proof. By backward induction, players have to play the Nash equilibrium in the last stage. As the behavior in the last stage does not depend on the behavior in the $(T-1)$-th stage, they have to play the NE also in the $(T-1)$-th stage. Then the same holds in the $(T-2)$-th stage, etc. □

8.1.2.1 Nash Equilibria

As Theorem 8.1 states, there is only one subgame-perfect equilibrium in a T-stage game based on Prisoner's dilemma, but what are other Nash equilibria (if there are any)?

	C	S
C	$(-5,-5)$	$(0,-20)$
S	$(-20,0)$	$(-1,-1)$

To simplify our discussion, we use the following notation: $X - YZ$, where $X, Y, Z \in \{C, S\}$ denotes the following strategy:

- in the first phase, play X;
- in the second phase, play Y if the opponent plays C in the first phase, otherwise play Z.

It can be shown that there are 4 Nash equilibria, as there are exactly four profiles that lead to $(C,C)(C,C)$ history – in these profiles, each player plays either $C - CC$ or $C - CS$.

Now we shall turn our attention to the underlying strategic-form game for a moment. Here, the strategy C strictly dominates S. But in the 2-stage game based on the Prisoner's dilemma, if player 2 plays $S - CS$, then the best responses of player 1 are $S - CC$ and $S - CS$. On the other hand, if player 2 plays $S - CS$, then the best responses are $C - SC$ and $C - CC$, so there is no strictly dominant strategy for player 1 (which would be among the best responses for all strategies of player 2).

> **Note**
>
> The strategy $S - CS$ is usually called "tit-for-tat".

8.1.2.2 Subgame-Perfect Equilibria

Let $\boldsymbol{s} = (s_1, s_2)$ be a Nash equilibrium in G. Define a strategy profile $\boldsymbol{\tau} = (\tau_1, \tau_2)$ in $G_{T\text{-rep}}$ where

- τ_1 chooses s_1 in every stage;
- τ_2 chooses s_2 in every stage.

Lemma 8.1. *Using $\boldsymbol{s}$ and $\boldsymbol{\tau}$ as we have defined above, $\boldsymbol{\tau}$ is a subgame-perfect equilibrium in $G_{T\text{-rep}}$ for every $T \geq 1$.*

Proof. Apparently, changing τ_i in some stage(s) may only result in the same or worse payoff for player i, since the other player always plays s_2 independent of the choices of player 1. □

The Lemma 8.1 may be generalized by allowing players to play different equilibria in particular stages, i.e., consider a sequence of Nash equilibria $\boldsymbol{s}^1, \boldsymbol{s}^2, \ldots, \boldsymbol{s}^T$ in G and assume that in stage l player i plays s_i^l.

But this still does **not** cover all possible subgame-perfect equilibria in finitely repeated games! Consider the following game G:

	m	f	r
M	$(4,4)$	$(-1,5)$	$(0,0)$
F	$(5,-1)$	$(1,1)$	$(0,0)$
R	$(0,0)$	$(0,0)$	$(3,3)$

The Nash equilibria in this strategic form game G are (F, f) and (R, r). Now consider a 2-stage game $G_{2\text{-rep}}$ and strategies τ_1, τ_2 where

- τ_1: chooses M in stage 1. In stage 2 plays R if (M, m) was played in the first stage, and plays F otherwise;
- τ_2: chooses m in stage 2. In stage 2 plays r if (M, m) was played in the first stage, and plays f otherwise.

Although both players **do not** play Nash equilibrium in the first stage, it still **is subgame-perfect equilibrium**. The idea is that both players agree to play a Pareto optimal profile. If both comply, then a favorable Nash equilibrium is played in the second stage. If one of them betrays then a "punishing" Nash equilibrium is played.

8.2 Infinitely Repeated Games

Definition 8.5. Let $G = (\{1,2\}, (S_1, S_2), (u_1, u_2))$ be a *finite* strategic-form game of two players. An **infinitely-repeated game** G_{irep} based on G proceeds in **stages** so that in each stage, say t, players choose a strategy profile $\boldsymbol{s}^t = (s_1^t, s_2^t)$.

Recall that a history of length $t \geq 0$ is a sequence $h = \boldsymbol{s}^1 \cdots \boldsymbol{s}^t \in S^t$ of t strategy profiles. Denote again by $H(t)$ the set of all histories of length t.

Definition 8.6. A **pure strategy** for player i in the infinitely repeated game G_{irep} is a function

$$\tau_i : \bigcup_{t=0}^{\infty} H(t) \to S_i,$$

which for every possible history chooses a next step for player i.

Every pure strategy profile $\boldsymbol{\tau} = (\tau_1, \tau_2)$ in G_{irep} induces a sequence of pure strategy profiles $w_{\boldsymbol{\tau}} = \boldsymbol{s}^1 \boldsymbol{s}^2 \cdots$ in G so that $s_i^t = \tau_i(\boldsymbol{s}^1 \cdots \boldsymbol{s}^{t-1})$. Here for $t = 1$ we have that $\boldsymbol{s}^1 \cdots \boldsymbol{s}^{t-1} = \epsilon$, which again denotes the empty history.

8.2.1 Discounted payoff

Let $\boldsymbol{\tau} = (\tau_1, \tau_2)$ be a pure strategy profile in G_{irep} such that $w_{\boldsymbol{\tau}} = \boldsymbol{s}^1 \boldsymbol{s}^2 \cdots$

Definition 8.7. Given $0 < \delta < 1$, we define δ-discounted payoff of player i by

$$u_i^{\delta}(\boldsymbol{\tau}) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t \cdot u_i(\boldsymbol{s}^{t+1}).$$

Given a strategic-form game G and $0 < \delta < 1$, we denote by G_{irep}^{δ} the infinitely repeated game based on G together with the δ-discounted payoffs.

Definition 8.8. A strategy profile $\boldsymbol{\tau} = (\tau_1, \tau_2)$ is a Nash equilibrium in G_{irep}^{δ} if for both $i \in \{1, 2\}$ and for every τ_i' and every $\boldsymbol{\tau}_{-i}$ we have that

$$u_i^{\delta}(\tau_i, \boldsymbol{\tau}_{-i}) \geq u_i^{\delta}(\tau_i', \boldsymbol{\tau}_{-i}).$$

Given a history $h = \boldsymbol{s}^1 \cdots \boldsymbol{s}^t$ and a strategy τ_i of player i, we define a strategy τ_i^h in the infinitely repeated game G_{irep} by

$$\tau_i^h(\bar{\boldsymbol{s}}^1 \cdots \bar{\boldsymbol{s}}^{\bar{t}}) = \tau_i(\boldsymbol{s}^1 \cdots \boldsymbol{s}^t \bar{\boldsymbol{s}}^1 \cdots \bar{\boldsymbol{s}}^{\bar{t}})$$

for every sequence $\bar{\boldsymbol{s}}^1 \cdots \bar{\boldsymbol{s}}^{\bar{t}}$ (i.e. τ_i^h behaves as τ_i after h).

Definition 8.9. Now $\boldsymbol{\tau} = (\tau_1, \tau_2)$ is a subgame-perfect equilibrium in G_{irep}^{δ} if for every history h we have that (τ_1^h, τ_2^h) is a Nash equilibrium.

Note

Note that (τ_1^h, τ_2^h) must be a NE also for all histories of h that are **not** visited when the profile (τ_1, τ_2) is used.

Example 8.3. Consider the infinitely repeated game G_{irep} based on the Prisoner's dilemma:

	C	S
C	$(-5, -5)$	$(0, -20)$
S	$(-20, 0)$	$(-1, -1)$

Our goal now will to be calculate the Nash and subgame-perfect equilibria in G_{irep}^{δ} for a given discount δ. Consider a pure strategy profile (τ_1, τ_2) where $\tau_1(\boldsymbol{s}^1 \cdots \boldsymbol{s}^T) = C$ for all $T \geq 1$ and $i \in \{1, 2\}$. Just as this was a subgame-perfect equilibrium in the finitely repeated Prisoner's dilemma, it stays a SPE in the infinitely repeated Prisoner's dilemma too (as it maximizes the payoff of each stage separately – as if it was a one-shot game).

8.2.2 Grim Trigger and Simple Folk Theorem

Now consider the infinitely repeated Prisoner's dilemma from Example 8.3 and the **grim trigger** profile $\boldsymbol{\tau} = (\tau_1, \tau_2)$ where

$$\tau_i(\boldsymbol{s}^1 \cdots \boldsymbol{s}^T) = \begin{cases} S, & T = 0, \\ S, & \boldsymbol{s}^l = (S, S) \quad \forall l \in \{1, \dots, T\}, \\ C, & \text{otherwise.} \end{cases}$$

Notice that (S, S) provides better payoffs than NE in G in each stage of G_{irep}. Thus it seems logical to remain **s**ilent as long, as the other player cooperates, but then keep **c**onfessing if he ever betrays us. Suppose that player i starts with this strategy and considers deviating in stage k to receive a payoff of 0 instead of -1. Thereafter, his opponent chooses C, and so he will also choose C in the remainder of the game. The use of the *grim trigger* strategies therefore defines a Nash equilibrium iff the equilibrium payoff of -1 is at least as large as the payoff from deviating to C in stage k and ever after (deviation to S

from C is surely suboptimal) [1]:

$$
\begin{aligned}
(1-\delta)\sum_{t=0}^{\infty}\delta^t(-1) &\geq (1-\delta)\left(\sum_{t=0}^{k-1}\delta^t(-1) + \underbrace{\delta^k\cdot 0}_{0} + \sum_{t=k+1}^{\infty}\delta^t(-5)\right) \\
\sum_{t=0}^{\infty}\delta^t(-1) &\geq \sum_{t=0}^{k-1}\delta^t(-1) + \sum_{t=k+1}^{\infty}\delta^t(-5) \\
\sum_{t=k}^{\infty}\delta^t(-1) &\geq \sum_{t=k+1}^{\infty}\delta^t(-5) \\
&\iff \\
\sum_{t=0}^{\infty}\delta^t(-1) &\geq \sum_{t=1}^{\infty}\delta^t(-5) \\
\frac{-1}{1-\delta} &\geq 5 + \sum_{t=0}^{\infty}\delta^t(-5) = 5 + \frac{-5}{1-\delta} \\
\frac{4}{1-\delta} &\geq 5 \\
\delta &\geq \frac{1}{5}.
\end{aligned}
$$

In general, let $G = (\{1,2\}, (S_1, S_2), (u_1, u_2))$ be two-player strategic form game where u_1, u_2 are bounded on $S = S_1 \times S_2$ (but S may be infinite) and let $\boldsymbol{s}^*$ be a Nash equilibrium in G. Let $\boldsymbol{s}$ be a strategy profile in G satisfying $u_i(\boldsymbol{s}) > u_i(\boldsymbol{s}^*)$ for all $i \in N$. Consider the following **grim trigger for $\boldsymbol{s}$ using $\boldsymbol{s}^*$** strategy profile $\boldsymbol{\tau} = (\tau_1, \tau_2)$ in G_{irep} where

$$
\tau_i(\boldsymbol{s}^1 \cdots \boldsymbol{s}^T) = \begin{cases} s_i, & T = 0, \\ s_i, & \boldsymbol{s}^l = \boldsymbol{s} \quad \forall l \in \{1, \dots, T\}, \\ s_i^*, & \text{otherwise.} \end{cases}
$$

Then for

$$
\delta \geq \max_{i \in \{1,2\}} \frac{\max_{s_i' \in S_i} u_i(s_i', \boldsymbol{s}_{-i}) - u_i(\boldsymbol{s})}{\max_{s_i' \in S_i} u_i(s_i', \boldsymbol{s}_{-i}) - u_i(\boldsymbol{s}^*)}
$$

we have that (τ_1, τ_2) is a subgame-perfect equilibrium in G_{irep}^{δ} and $u_i^{\delta}(\boldsymbol{\tau}) = u_i(\boldsymbol{s})$.

8.2.2.1 Examples

Example 8.4. Consider the infinitely repeated game G_{irep} based on the following game G:

	m	f	r
M	$(4,4)$	$(-1,5)$	$(3,0)$

[1] Mostly adapted from this document.

	m	f	r
F	$(5,-1)$	$(1,1)$	$(0,0)$
R	$(0,3)$	$(0,0)$	$(2,2)$

There is one Nash equilibrium in G, that is (F, f). Consider the grim trigger for (M, m) using (F, f), i.e., the profile $\boldsymbol{\tau} = (\tau_1, \tau_2)$ in G_{irep} where

- τ_1: plays M in a given stage if (M, m) was played so far in all previous stages, and plays F otherwise;
- τ_2: plays m in a given stage if (M, m) was played so far in all previous stages, and plays f otherwise;

Then this strategy profile $\boldsymbol{\tau}$ is a subgame-perfect equilibrium in G_{irep} for all $\delta \geq \frac{1}{4}$. Also, $u_i(\boldsymbol{\tau}) = 4$ for all $i \in \{1, 2\}$. There is also a grim trigger for (R, r) using (F, f), which is a SPE in G_{irep} for $\delta \geq \frac{1}{2}$.

Example 8.5 (Tacit Collusion)**.** Consider the Cournot duopoly, see Section 3.3.3, game model $G = (N, (S_i)_{i\in N}, (u_i)_{i\in N})$

- $N = \{1, 2\}$;
- $S_i = [0, \kappa]$;
- $u_1(q_1, q_2) = q_1(\kappa - q - 1 - q_2) - q_1c_1 = (\kappa - c_1)q_1 - q_1^2 - q_1q_2$ and
  $u_2(q_1, q_2) = q_2(\kappa - q - 1 - q_2) - q_2c_2 = (\kappa - c_2)q_2 - q_2^2 - q_1q_2$.

For simplicity we assume that $c_1 = c_2 = c$ and denote $\theta = \kappa - c$.

If the firms sign a *binding contract* to produce only $\theta/4$, their profit would be $\theta^2/8$ which is higher than the profit $\theta^2/9$ for playing the Nash equilibrium $(\theta/3, \theta/3)$ in G. However, such contracts are forbidden in many countries (including US). Is then still possible that the firms will behave selfishly (i.e. only maximizing their profits) and still obtain such payoffs? In other words, is there a SPE in the infinitely repeated game based on G (with a discount factor δ) which gives the payoffs $\theta^2/8$?

Consider the grim trigger profile for $(\theta/4, \theta/4)$ using $(\theta/3, \theta/3)$ Nash equilibrium. Then player i will

- produce $q_i = \theta/4$ whenever all profiles in history are $(\theta/4, \theta/4)$;
- whenever one of the players deviates, produce $\theta/3$ from that moment on.

Assume that $\kappa = 100$ and $c = 10$ (which gives $\theta = 90$), this is a SPE in G_{irep} for $\delta \geq 0.5294\cdots$ and it results in $(\theta/4, \theta/4)(\theta/4, \theta/4)\cdots$ with discounted payffos $\theta^2/8$.

8.3 Long-Run Average Payoff

> **!** Important
>
> In this section, we assume that all payoffs in the game G are positive and that S is finite!

Let $\boldsymbol{\tau} = (\tau_1, \tau_2)$ be a strategy profile in the infinitely repeated game G_{irep} such that $w_{\boldsymbol{\tau}} = \boldsymbol{s}^1\boldsymbol{s}^2\cdots$

Definition 8.10. We define a **long-run average payoff** for player i by

$$u_i^{\text{avg}}(\boldsymbol{\tau}) = \limsup_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} u_i(\boldsymbol{s}^t).$$

The long-run average payoff $u_i^{\text{avg}}(\boldsymbol{\tau})$ is **well-defined** if the limit $u_i^{\text{avg}}(\boldsymbol{\tau}) = \lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} u_i(\boldsymbol{s}^t)$ exists.

> Caution
>
> Here, lim sup is necessary in general, because τ_i may cause non-existence in the limit (but in $\mathbb{R}$ with the usual choice of metrics, we can always select a convergent subsequence).

Given a strategic-form game G, we denote by $G_{\text{irep}}^{\text{avg}}$ the infinitely repeated game based on G together with the long-run average payoff.

Definition 8.11. A strategy profile $\boldsymbol{\tau}$ is a Nash equilibrium if $u_i^{\text{avg}}(\boldsymbol{\tau})$ is *well-defined* for all $i \in N$, and for every i and every τ_i' we have that

$$u_i^{\text{avg}}(\tau_i, \boldsymbol{\tau}_{-i}) \geq u_i^{\text{avg}}(\tau_i', \boldsymbol{\tau}_{-i}).$$

> Note
>
> Note that we demand the existence of the defining limit of $u_i^{\text{avg}}(\tau_i, \boldsymbol{\tau}_{-i})$ but the limit does **not** have to exist for $u_i^{\text{avg}}(\tau_i', \boldsymbol{\tau}_{-i})$.

Moreover, $\boldsymbol{\tau} = (\tau_1, \tau_2)$ is a subgame-perfect equilibrium in $G_{\text{irep}}^{\text{avg}}$ if for every history h that (τ_1^h, τ_2^h) is a Nash equilibrium.

Example 8.6. Consider the infinitely repeated game based on the Prisoner's dilemma with long-run average payoff:

	C	S
C	$(-5, -5)$	$(0, -20)$
S	$(-20, 0)$	$(-1, -1)$

The grim trigger profile $\boldsymbol{\tau} = (\tau_1, \tau_2)$ where

$$\tau_i(\boldsymbol{s}^1 \cdots \boldsymbol{s}^T) = \begin{cases} S, & T = 0, \\ S, & \boldsymbol{s}^l = (S, S) \quad \forall l \in \{1, \dots, T\}, \\ C, & \text{otherwise.} \end{cases}$$

is a SPE that gives the long-run average payoff -1 to each player.

The intuition behind the grim trigger works just like for the discounted payoff – whenever a player i deviates, the player $-i$ starts playing C for which the best response of player i is also C. So we obtain $(S,S)\cdots(S,S)(X,Y)(C,C)(C,C)\cdots$ (here (X,Y) is either (C,S) or (S,C) depending on who deviates). Apparently, the long-run average payoff is -5 for both players, which is worse than -1.

However, other payoffs can be supported by Nash equilibrium. Consider e.g. a strategy profile $\boldsymbol{\tau} = (\tau_1, \tau_2)$ such that

- both players **cyclically** play as follows:
  - 9 times (S,S);
  - once (S,C);
- if one of the players deviates, then, from that moment on, both play (C,C) forever.

Then $\boldsymbol{\tau} = (\tau_1, \tau_2)$ is also a SPE. Apparently, $u_1^{\text{avg}}(\boldsymbol{\tau}) = \frac{9}{10}(-1) + \frac{-20}{10} = -\frac{29}{10}$ and $u_2^{\text{avg}}(\boldsymbol{\tau}) = \frac{9}{10}(-1) + 0 = -\frac{9}{10}$. Hence player 2 gets a better payoff than from the "best" profile (S,S).

8.4 Folk Theorems

The previous examples suggest that other (possibly all?) convex combinations of payoffs may be obtained by means of Nash equilibria. This observation forms a basis for a bunch of theorems, collectively referred to as *Folk Theorems*.

> **Note**
>
> No author is listed since these theorems had been known in the games community long before they were formalized.

In this section, we prove several versions of *Folk Theorem* concerning achievable payoffs for repeated games. We consider the following variants:

- Long-run average payoffs and SPE;
- Long-run average payoffs and Nash equilibria.

> **Tip**
>
> Similar theorems can also be proven for the discounted payoff.

8.4.1 Feasible Payoffs

Definition 8.12. We say that a vector of payoffs $\boldsymbol{v} = (v_1, v_2) \in \mathbb{R}^2$ is **feasible** if it is a convex combination of payoffs for pure strategy profiles in G with **rational** coefficients, i.e. if there are rational numbers $\beta_{\boldsymbol{s}}$, here $\boldsymbol{s} \in S$, satisfying $\beta_{\boldsymbol{s}} \geq 0$ and $\sum_{\boldsymbol{s} \in S} \beta_{\boldsymbol{s}} = 1$ such that for both $i \in \{1, 2\}$ holds

$$v_i = \sum_{\boldsymbol{s} \in S} \beta_{\boldsymbol{s}} \cdot u_i(\boldsymbol{s}).$$

We assume that there is $m \in \mathbb{N}$ such that each $\beta_{\boldsymbol{s}}$ can be written in the form $\beta_{\boldsymbol{s}} = \gamma_{\boldsymbol{s}}/m$.

The following theorems can be extended to a notion of feasible payoffs using *arbitrary, possibly irrational,* coefficients $\beta_{\boldsymbol{s}}$ in the convex combination. Roughly speaking, this follows from the fact that each real number can be approximated with rational numbers up to an arbitrary error. However, the proofs are technically more involved.

Theorem 8.2. *Let $\boldsymbol{s}^*$ be a pure strategy Nash equilibrium in G and let $\boldsymbol{v} = (v_1, v_2)$ be a **feasible** vector of payoffs satisfying $v_i \geq u_i(\boldsymbol{s}^*)$ for both $i \in \{1, 2\}$. Then there is a strategy profile $\boldsymbol{\tau} = (\tau_1, \tau_2)$ in G_{irep} such that*

- *$\boldsymbol{\tau}$ is a subgame-perfect equilibrium in $G_{\text{irep}}^{\text{avg}}$;*
- *$u_i^{\text{avg}}(\boldsymbol{\tau}) = v_i$ for $i \in \{1, 2\}$.*

Proof. Consider a strategy profile $\boldsymbol{\tau}$ in G_{irep} which gives the following behavior:

1. unless one of the players deviates, the players play **cyclically** all profiles $\boldsymbol{s} \in S$ so that each $\boldsymbol{s}$ is always played for $\gamma_{\boldsymbol{s}}$ rounds;
2. whenever one of the players deviates, then, from that moment on, each player i plays $\boldsymbol{s}_i^*$.

Trivially, $\boldsymbol{u}^{\text{avg}}(\boldsymbol{\tau}) = \boldsymbol{v}$. We shall now verify that $\boldsymbol{\tau}$ is a subgame-perfect equilibrium. Consider a fixed history h, we show that $\boldsymbol{\tau}^h = (\tau_1^h, \tau_2^h)$ is a Nash equilibrium in $G_{\text{irep}}^{\text{avg}}$.

- If h does not contain deviation from the cyclic behavior (1.), then $\boldsymbol{\tau}^h$ continues according to (1.), thus $u_i^{\text{avg}}(\boldsymbol{\tau}^h) = v_i$.
- If h contains a deviation from (1.), then

  $$w_{\boldsymbol{\tau}^h} = \boldsymbol{s}^* \boldsymbol{s}^* \cdots$$

  and thus $u_i^{\text{avg}}(\boldsymbol{\tau}^h) = u_i(\boldsymbol{s}^*)$.
- Now if a player i deviates from τ_i^h to $\bar{\tau}_i^h$ in $G_{\text{irep}}^{\text{avg}}$, then

  $$w_{(\bar{\tau}_i^h, \boldsymbol{\tau}_{-i}^h)} = \alpha(s_i^1, \boldsymbol{s}'_{-i})(s_i^2, \boldsymbol{s}_{-i}^*)(s_i^3, \boldsymbol{s}_{-i}^*) \cdots$$

  where α is a sequence of profiles following the cyclic behavior (1.), $s_i^1, s_i^2, \ldots$ are strategies of S_i and $\boldsymbol{s}'_{-i}$ is a strategy of S_{-i}. However, then $u_i^{\text{avg}}(\bar{\tau}_i^h, \boldsymbol{\tau}_{-i}^h) \leq u_i(\boldsymbol{s}^*) \leq v_i$ since $\boldsymbol{s}^*$ is a Nash equilibrium and thus $u_i(s_i^k, \boldsymbol{s}_{-i}^*) \leq u_i(\boldsymbol{s}^*)$ for all $k \geq 1$.

□

Note

Intuitively said, player $-i$ punishes player i by playing $\boldsymbol{s}^*_{-i}$.

8.4.2 Individual Rationality

Definition 8.13. A tuple $\boldsymbol{v} = (v_1, v_2) \in \mathbb{R}^2$ is **individually rational** if for both $i \in \{1, 2\}$ holds

$$v_i \geq \min_{\boldsymbol{s}_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_i, \boldsymbol{s}_{-i}).$$

That is, v_i is least as large as the value that player i may secure by playing best responses to the most hostile behavior of player $-i$.

Tip

Note that here, player $-i$ chooses his "most hostile behavior" before player i, who then reacts to it maximizing his payoff in this unfavorable situation, i.e. he minimizes the harm done by player $-i$ to him. Hence player $-i$ is able to hurt player i less than if the order of maximization/minimization was the other way around.

Example 8.7. Consider a game given by the following table:

	L	R
U	$(-2, 2)$	$(1, -2)$
M	$(1, -2)$	$(-2, 2)$
D	$(0, 1)$	$(2, 3)$

Here any $\boldsymbol{v} = (v_1, v_2)$ such that $v_1 \geq 1$ and $v_2 \geq 2$ is individually rational.

Theorem 8.3. *Let $\boldsymbol{v} = (v_1, v_2)$ be a **feasible** and **individually rational** vector of payoffs. Then there is a strategy profile $\boldsymbol{\tau} = (\tau_1, \tau_2)$ in G_{irep} such that*

- *$\boldsymbol{\tau}$ is a Nash equilibrium in $G^{\text{avg}}_{\text{irep}}$;*
- *$\boldsymbol{u}^{\text{avg}}(\boldsymbol{\tau}) = \boldsymbol{v}$.*

Proof. We will use a slightly modified strategy profile $\boldsymbol{\tau} = (\tau_1, \tau_2)$ in G_{irep} from the proof of Theorem 8.2:

- unless one of the players deviates, the players play **cyclically** all profiles $\boldsymbol{s} \in S$ so that each $\boldsymbol{s}$ is always played for $\gamma_{\boldsymbol{s}}$ rounds;
- whenever a player i deviates, the opponent $-i$ starts playing a strategy $\boldsymbol{s}^{\min}_{-i} \in \operatorname{argmin}_{\boldsymbol{s}_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_i, \boldsymbol{s}_{-i})$.

It is trivial to see that $\boldsymbol{u}^{\text{avg}}(\boldsymbol{\tau}) = \boldsymbol{v}$. Now if a player i deviates, then his long-run average payoff cannot be higher than $\min_{\boldsymbol{s}_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_i, \boldsymbol{s}_{-i}) \leq v_i$, so $\boldsymbol{\tau}$ is a Nash equilibrium. □

Theorem 8.4. *If a strategy profile* $\boldsymbol{\tau} = (\tau_1, \tau_2)$ *is a Nash equilibrium in* $G^{\text{avg}}_{\text{irep}}$*, then* $\left(u_1^{\text{avg}}(\boldsymbol{\tau}), u_2^{\text{avg}}(\boldsymbol{\tau})\right)$ *is individually rational.*

Proof. Suppose that $\left(u_1^{\text{avg}}(\boldsymbol{\tau}), u_2^{\text{avg}}(\boldsymbol{\tau})\right)$ is not individually rational. Without loss of generality, assume that $u_1^{\text{avg}}(\boldsymbol{\tau}) < \min_{s_2 \in S_2} \max_{s_1 \in S_1} u_1(s_1, s_2)$. Now let us consider a new strategy $\bar{\tau}_1$ such that for every history h the pure strategy $\bar{\tau}_1(h)$ is a best response to $\tau_2(h)$. But then, for every history h, we have

$$u_1(\bar{\tau}_1(h), \tau_2(h)) \geq \min_{s_2 \in S_2} \max_{s_1 \in S_1} u_1(s_1, s_2) > u_1^{\text{avg}}(\boldsymbol{\tau}).$$

So clearly $u_1^{\text{avg}}(\bar{\tau}_1, \tau_2) > u_1^{\text{avg}}(\boldsymbol{\tau})$ which contradicts the fact that $\boldsymbol{\tau} = (\tau_1, \tau_2)$ is a Nash equilibrium. □

Note that if irrational convex combinations are allowed in the definition of feasibility, see Definition 8.12, then vectors of payoffs for Nash equilibria in $G^{\text{avg}}_{\text{irep}}$ are exactly feasible and individually rational vectors of payoffs. Indeed, the coefficients $\beta_{\boldsymbol{s}}$ in the definition of feasibility are exactly frequencies with which the individual profiles of S are played in the Nash equilibria.

8.4.3 Summary

We have proven that "any reasonable" (i.e. feasible and individually rational) vector of payoffs can be justified as payoffs for a Nash equilibrium in $G^{\text{avg}}_{\text{irep}}$ (where the future has "an infinite weight"). Concerning subgame-perfect equilibria, we have proven that any feasible vector of payoffs dominating a Nash equilibrium in G can be justified as payoffs for subgame-perfect equilibrium in $G^{\text{avg}}_{\text{irep}}$.

> **Note**
>
> This result can be generalized to arbitrary feasible and *strictly* individually rational payoffs by means of a more demanding construction.

For discounted payoffs, one can prove that an arbitrary feasible vector of payoffs dominating a Nash equilibrium in G can be approximated using payoffs for subgame-perfect equilibrium in G^{δ}_{irep} as δ goes to 1, and even this result can be extended to feasible and strictly individually rational payoffs.

> **Tip**
>
> For a very detailed discussion of Folk Theorems see *"A Course in Game Theory"* by M. J. Osborne and A. Rubinstein.

8.5 Summary of Extensive-Form Games

We have considered extensive-form games (i.e., games on trees)

- with perfect information;
- with imperfect information.

We have considered pure strategies, mixed and behavioral strategies (see Kuhn's Theorem 7.1). We have also studied Nash equilibria and subgame-perfect equilibria in pure strategies.

For finite perfect-information games, we have shown that

- there always exists a pure strategy SPE;
- SPE can be computed using backward induction in polynomial time.

On the other hand, for imperfect information games, the following holds:

- the backward induction can be used to propagate values through "perfect information nodes", but "imperfect information parts" have to be solved by different means;
- solving imperfect information games is at least as hard as solving games in strategic form; however, even in the zero-sum case, most decision problems are NP-hard.

Finally, we discussed repeated games. We considered both, finitely as well as infinitely repeated games. For finitely repeated games we considered the average payoff and discussed the existence of pure strategy Nash and subgame-perfect equilibria with respect to the existence of Nash equilibria in the original strategic-form game.

For infinitely repeated games we considered both

- **discounted payoff**: we have formulated and applied a simple folk theorem: "grim trigger" strategy profiles can be used to implement any vector of payoffs strictly dominating payoffs for a Nash equilibrium in the original strategic-form game;
- **long-run average payoff**: we have proven that all feasible and individually rational vectors of payoffs can be achieved by Nash equilibria (a variant of the grim trigger).

9 Games of Incomplete Information

9.1 Auctions

Let us consider the (general) problem: How to allocate (discrete) resources among selfish agents in a multi-agent system?

Auctions provide a general solution to this problem. As such, auctions have been heavily used in real life, in consumer, corporate as well as government settings:

- eBay, art auctions, wine auctions, etc.
- advertising (Google adWords)
- governments selling public resources: electromagnetic spectrum, oil leases, etc.
- …

Auctions also provide a theoretical framework for understanding resource allocation problems among self-interested agents. Formally, an auction is any protocol that allows agents to indicate their interest in one or more resources and that uses these indications to determine the resource allocation and payments of the agents.

Auctions may be used in various settings depending on the complexity of the resource allocation problem:

- *Single-item auctions*: Here n bidders (players) compete for a single indivisible item that can be allocated to just one of them. Each bidder has his own private value of the item in case he wins (gets zero if he loses). Typically (but not always) the highest bid wins. How much should he pay?
- *Multiunit auctions*: Here a fixed number of identical units of a homogenous commodity are sold. Each bidder submits both a number of units he demands and a unit price he is willing to pay. Here also the highest bidders typically win, but it is unclear how much should they pay (pay-as-bid vs uniform pricing).
- *Combinatorial auctions*: Here bidders compete for a set of distinct goods. Each player has a valuation function that assigns values to *subsets* of the set (some goods are useful only in groups etc.). Who wins and what he pays?

> **Tip**
>
> We shall mostly concentrate on single-item auctions.

9.1.1 Single-Item Auctions

There are many single-item auctions, but we consider the following well-known versions:

- *open auctions:*
  - **The English Auction**: Often occurs in movies; bidders are sitting in a room (by computer or a phone) and the price of the item goes up as long as someone is willing to bid it higher. Once the last increase is no longer challenged, the last bidder to increase the price wins the auction and pays the price for the item.
  - **The Dutch Auction**: Opposite of the English auction; here, the price starts at a prohibitively high value and the auctioneer gradually drops the price. Once a bidder shouts "buy", the auction ends and the bidder gets the item at the price.
- *sealed-bid auctions:*
  - ***k*-th price Sealed-Bid Auction**: Each bidder writes down his bid and places it in an envelope; the envelopes are opened simultaneously. The highest bidder wins and then pays the ***k*-th maximum bid**. (In reverse auction it is the k-th minimum). The most prominent special cases are **The First-Price Auction** and **The Second-Price Auction**.

Observe that

- the English auction is essentially equivalent to the second price auction if the increments in every round are very small.

> **Note**
>
> There exists a "continuous" version, called the Japanese auction, where the price continuously increases. Each bidder may drop out at any time. The last one who stays gets the item for the current price (which is the dropping price of the "second highest bid").

- similarly, the Dutch auction is equivalent to the first price auction. Note that the bidder with the highest bid stops the decrement of the price and buys at the current price which corresponds to his bid.

But now the question arises, which auction (and if any) is better? The goal of the bidder is clear – to get the item at as low a price as possible (i.e. they maximize the difference between their private value and the price they pay). We consider only self-interested non-communicating bidders, that are rational and intelligent. There are also at least two goals that may be pursued by the auctioneer (in various settings):

- revenue maximization;
- incentive compatibility: we want the bidder to spontaneously bid their true value of the item.

> **Tip**
>
> The *incentive compatibility* objective means that we do not want it to be possible to strategically manipulate the auction by lying.

9.1.2 Auctions as Games

Consider *single-item sealed-bid auctions* as strategic-form games: Let $G = (N, (B_i)_{i\in N}, (u_i)_{i\in N})$ be a strategic-form game where

- the set of players N is the set of bidders;
- $B_i = [0, \infty)$ where each $b_i \in B_i$ corresponds to the bid b_i (we follow the standard notation and use b_i to denote pure strategies (bids)).
- to define u_i, we assume that each bidder has his own private value v_i of the item, then given bids $\boldsymbol{b} = (b_1, \dots, b_n)$:
  - **First Price:**
$$u_i(\boldsymbol{b}) = \begin{cases} v_i - b_i, & b_i > \max_{j\neq i} b_j, \\ 0, & \text{otherwise;} \end{cases}$$
  - **Second Price:**
$$u_i(\boldsymbol{b}) = \begin{cases} v_i - \max_{j\neq i} b_j, & b_i > \max_{j\neq i} b_j, \\ 0, & \text{otherwise.} \end{cases}$$

Is this model realistic? Not really – usually, the bidders are not perfectly informed about the private values of the other bidders.

> **Tip**
>
> Although imperfect-information extensive-form games would solve this issue, the construction would be awkward, to say the least (and would miss the point of the problem).

9.2 Incomplete-Information Games

Definition 9.1. A **(strict) incomplete information** game is a tuple $G = (N, (A_i)_{i\in N}, (T_i)_{i\in N}, (u_i)_{i\in N})$ where

- $N = \{1, \dots, n\}$ is a set of players;
- each A_i is a set of **actions** available to player i; we denote by $A = \prod_{i=1}^n A_i$ the set of all possible **action profiles** $\boldsymbol{a} = (a_1, \dots, a_n)$;
- each T_i is the set of all **possible types** of player i and we denote by $T = \prod_{i=1}^n T_i$ the set of all possible **type profiles** $\boldsymbol{t} = (t_1, \dots, t_n)$;

- u_i is a type-dependent payoff function

$$u_i : A_1 \times \cdots \times A_n \times T_i \to \mathbb{R};$$

  given a profile of actions $\boldsymbol{a} = (a_1, \dots, a_n)$ and a type $t_i \in T_i$, we write $u_i(\boldsymbol{a}; t_i) = u_i(a_1, \dots, a_n; t_i)$ to denote the corresponding payoff.

A pure strategy of player i is a function $s_i : T_i \to A_i$. As before, we denote by S_i the set of all pure strategies of player i, and by S the set of all pure strategy profiles $S = \prod_{i=1}^n S_i$.

9.2.1 Dominance

Definition 9.2. A pure strategy s_i **very weakly dominates** s_i' if for every $t_i \in T_i$ the following holds: For all $\boldsymbol{a}_{-i} \in A_{-i}$ we have that

$$u_i(s_i(t_i), \boldsymbol{a}_{-i}; t_i) \geq u_i(s_i'(t_i), \boldsymbol{a}_{-i}; t_i).$$

Definition 9.3. A pure strategy s_i **weakly dominates** s_i' if for every $t_i \in T_i$ the following holds: For all $\boldsymbol{a}_{-i} \in A_{-i}$ we have that

$$u_i(s_i(t_i), \boldsymbol{a}_{-i}; t_i) \geq u_i(s_i'(t_i), \boldsymbol{a}_{-i}; t_i)$$

and the inequality is strict for at least one $\boldsymbol{a}_{-i}$.

> Tip
>
> Such $\boldsymbol{a}_{-i}$ satisfying the strict inequality may be different for different types t_i.

Definition 9.4. A pure strategy s_i **strictly dominates** s_i' if for every $t_i \in T_i$ the following holds: For all $\boldsymbol{a}_{-i} \in A_{-i}$ we have that

$$u_i(s_i(t_i), \boldsymbol{a}_{-i}; t_i) > u_i(s_i'(t_i), \boldsymbol{a}_{-i}; t_i).$$

Definition 9.5. A pure strategy s_i is **(very weakly, weakly, strictly) dominant** if it (very weakly, weakly, strictly) dominates all other pure strategies $s_i' \in S_i$.

In order to generalize Nash equilibria to incomplete-information games, we use the following notation: Given a pure strategy profile $\boldsymbol{s} = (s_1, \dots, s_n) \in S$ and a profile of types $\boldsymbol{t} = (t_1, \dots, t_n) \in T$, for every player i we write

$$\boldsymbol{s}_{-i}(\boldsymbol{t}_{-i}) = (s_1(t_1), \dots, s_{i-1}(t_{i-1}), s_{i+1}(t_{i+1}), \dots, s_n(t_n)).$$

Definition 9.6 (Ex-post Nash equilibrium[1])**.** A strategy profile $\boldsymbol{s} = (s_1, \dots, s_n)$ is an **ex-post Nash equilibrium** if no player can increase their *ex-post* expected utility $u_i(s_i, \boldsymbol{s}_{-i}; t_i, \boldsymbol{t}_{-i})$ by unilaterally changing their strategy, i.e. for every player i and every type profile $\boldsymbol{t} \in T$ and every strategy $s_i' \in S_i$ of player i it holds that

$$u_i(\boldsymbol{s}; \boldsymbol{t}) \geq u_i(s_i', \boldsymbol{s}_{-i}; \boldsymbol{t}).$$

Note that in the *ex-post* expected utility, player i knows types of other players!

[1] The definition, and much more, can be found in this awesome document by prof. Greenwald.

Warning

In the lectures, it was defined as follows:
A strategy profile $\boldsymbol{s} = (s_1, \dots, s_n)$ is an **ex-post Nash equilibrium** if for **every** type profile $\boldsymbol{t} = (t_1, \dots, t_n) \in T$, we have that $\boldsymbol{s}(\boldsymbol{t}) = (s_1(t_1), \dots, s_n(t_n))$ is a Nash equilibrium in the strategic form game defined by the t_i's.
Formally, $\boldsymbol{s} = (s_1, \dots, s_n)$ is an **ex-post Nash equilibrium** if for all $i \in N$ and all $\boldsymbol{t} = (t_1, \dots, t_n) \in T$ and all $a_i \in A_i$ the following holds

$$u_i(\boldsymbol{s}(\boldsymbol{t}); t_i) \geq u_i(a_i, \boldsymbol{s}_{-i}(\boldsymbol{t}_{-i}); t_i).$$

Example 9.1. Consider *single-item sealed-bid auctions* as strict incomplete-information games: $G = (N, (B_i)_{i \in N}, (V_i)_{i \in N}, (u_i)_{i \in N})$ where

- the set of players N is the set of bidders;
- $B_i = [0, \infty)$ where each action $b_i \in B_i$ corresponds to the bid b_i;
- $V_i = [0, \infty)$ where each type $v_i \in V_i$ corresponds to the private value v_i;
- let $v_i \in V_i$ be the type of player i (i.e. his private value), then given an action profile $\boldsymbol{b} = (b_1, \dots, b_n)$ (i.e. bids) we define
  - **First Price**
    $$u_i(\boldsymbol{b}; v_i) = \begin{cases} v_i - b_i, & b_i > \max_{j \neq i} b_j, \\ 0, & \text{otherwise,} \end{cases}$$
  - **Second Price**
    $$u_i(\boldsymbol{b}; v_i) = \begin{cases} v_i - \max_{j \neq i} b_j, & b_i > \max_{j \neq i} b_j, \\ 0, & \text{otherwise.} \end{cases}$$

Are there dominant strategies? Are there ex-post-Nash equilibria?

Note

Note that when there is a tie (i.e. there are $k \neq l$ such that $b_l = b_k = \max_j b_j$), then all players get 0.

9.2.2 Second-Price Auction

For every i, we denote by v_i the pure strategy s_i for player i defined by $s_i(v_i) = v_i$.

Tip

Intuitively, such a strategy is *truth-telling*, which means that the player bids his own private value truthfully.

Theorem 9.1. *Assume Second-Price Auction. Then for every player* i *we have that* v_i *is a weakly dominant strategy. So* $\boldsymbol{v}$ *is also an* ex-post Nash equilibrium.

Proof. Let us fix the private value v_i and the bid $b_i \in B_i$ of player i such $b_i \neq v_i$. We show that for all bids of opponents $\boldsymbol{b}_{-i} \in B_{-i}$:

$$u_i(v_i, \boldsymbol{b}_{-i}; v_i) \geq u_i(b_i, \boldsymbol{b}_{-i}; v_i)$$

with the strict inequality for at least one $\boldsymbol{b}_{-i}$. Intuitively, assume that player i bids b_i against $\boldsymbol{b}_{-i}$ of his opponents and compares his payoff with the payoff he obtains by playing v_i against $\boldsymbol{b}_{-i}$.

There are two cases to consider; $b_i < v_i$ and $b_i > v_i$:

- **Case** $b_i < v_i$: We distinguish three cases depending on $\boldsymbol{b}_{-i}$

  1. If $b_i > \max_{j\neq i} b_j$, then
  $$u_i(b_i, \boldsymbol{b}_{-i}; v_i) = v_i - \max_{j\neq i} b_j = u_i(v_i, \boldsymbol{b}_{-i}; v_i);$$
  intuitively, player i wins and pays the price $\max_{j\neq i} b_j < b_i$; however, then bidding $v_i > b_i$, player i still wins and pays $\max_{j\neq i} b_j$ as well;
  2. If there is $k \neq i$ such that $b_k > \max_{j\neq k} b_j$, then
  $$u_i(b_i, \boldsymbol{b}_{-i}; v_i) = 0 \leq u_i(v_i, \boldsymbol{b}_{-i}; v_i);$$
  moreover, if $b_i < b_k < v_i$, then we get a strict inequality (required by Definition 9.3)
  $$u_i(b_i, \boldsymbol{b}_{-i}; v_i) = 0 < v_i - b_k = u_i(v_i, \boldsymbol{b}_{-i}; v_i);$$
  intuitively, if another player k wins, then player i gets 0 and increasing b_i to v_i surely does not hurt. Moreover, if $b_i < b_k < v_i$, then increasing b_i to v_i strictly increases the payoff of player i;
  3. If there are $k \neq l$ such that $b_k = b_l = \max_j b_j$, then
  $$u_i(b_i, \boldsymbol{b}_{-i}; v_i) = 0 \leq u_i(v_i, \boldsymbol{b}_{-i}; v_i);$$
  intuitively, there is a tie in $(b_i, \boldsymbol{b}_{-i})$ and all players get 0.

- **Case** $b_i > v_i$: We distinguish four sub-cases depending on $\boldsymbol{b}_{-i}$

  1. If $b_i > \max_{j\neq i} b_j > v_i$, then
  $$u_i(b_i, \boldsymbol{b}_{-i}; v_i) = v_i - \max_{j\neq i} b_j < 0 = u_i(v_i, \boldsymbol{b}_{-i}; v_i),$$
  so in this case, the inequality is strict;
  2. If $b_i > v_i \geq \max_{j\neq i} b_j$, then
  $$u_i(b_i, \boldsymbol{b}_{-i}; v_i) = v_i - \max_{j\neq i} b_j = u_i(v_i, \boldsymbol{b}_{-i}; v_i);$$
  note that this case also covers $v_i = \max_{j\neq i} b_j$ where decreasing b_i to v_i causes a tie with zero payoff for player i (and everybody else);

3. If there is $k \neq i$ such that $b_k > \max_{j \neq k} b_j > v_i$, then
$$u_i(b_i, \boldsymbol{b}_{-i}; v_i) = 0 = u_i(v_i, \boldsymbol{b}_{-i}; v_i);$$
4. If there are $k \neq l$ such that $b_k = b_l = \max_j b_j > v_i$, then
$$u_i(b_i, \boldsymbol{b}_{-i}; v_i) = 0 = u_i(v_i, \boldsymbol{b}_{-i}; v_i).$$

□

9.2.3 First-Price Auction

Consider now the First-Price Auction. Here the highest bidder wins and pays his bid. Let us impose a (reasonable) assumption that no player bids more than his private value. We shall now show that there are no dominant strategies.

Assume the opposite, i.e. that there is a very weakly dominant strategy s_i of player i.

> Counter-example idea
>
> Intuitively, if player i wins against some bids of his opponents, then his bid is strictly larger than the bids of all his opponents. Hence he can slightly decrease his bid and still win.

Formally, assume that all other players bid 0, i.e. $b_j = 0$ for $j \neq i$, and let $v_i > 0$. If $s_i(v_i) > 0$, then
$$u_i(s_i(v_i), \boldsymbol{b}_{-i}; v_i) = v_i - s_i(v_i) < v_i - \frac{1}{2}s_i(v_i) = u_i\left(\frac{1}{2}s_i(v_i), \boldsymbol{b}_{-i}; v_i\right).$$
On the other hand, if $s_i(v_i) = 0$, then
$$u_i(s_i(v_i), \boldsymbol{b}_{-i}; v_i) = 0 < v_i - \frac{1}{2}v_i = u_i\left(\frac{1}{2}v_i, \boldsymbol{b}_{-i}; v_i\right).$$

Hence s_i cannot be very weakly dominant (and thus neither weakly, nor strictly). Similarly, we will show, using again a counter-example, that there is no ex-post Nash equilibrium. Therefore, let us assume the pure strategy profile $\boldsymbol{s} = (s_1, \ldots, s_n)$ is an ex-post Nash equilibrium. Consider $0 < v_1 < \cdots < v_{n-1}$ and define
$$M = \max\{s_i(v_i) |\, i \in \{1, \ldots, n-1\}\}.$$
Let $v_n = M + 1$. If player n wins, i.e. $s_n(v_n) > M$, then
$$\begin{aligned} u_n(s_n(v_n), \boldsymbol{s}_{-i}(\boldsymbol{v}_{-i}); v_n) &= v_n - s_n(v_n) \\ &< v_n - (s_n(v_n) - \varepsilon) \\ &= u_n(s_n(v_n) - \varepsilon, \boldsymbol{s}_{-i}(\boldsymbol{v}_{-i}); v_n) \end{aligned}$$
for $\varepsilon \in (0, s_n(v_n) - M)$ (i.e. player n may help himself by lowering his bid slightly). If player n does **not** win, i.e. $s_n(v_n) \leq M < M + 1 = v_n$, then for $\varepsilon \in (0, 1)$, e.g. $\varepsilon = \frac{1}{2}$, we get
$$u_n(s_n(v_n), \boldsymbol{s}_{-i}(\boldsymbol{v}_{-i}); v_n) = 0 < \frac{1}{2} = u_n(v_n - \varepsilon, \boldsymbol{s}_{-i}(\boldsymbol{v}_{-i}); v_n),$$
i.e. player n may help himself by playing $v_n - \frac{1}{2}$. This concludes the counter-example.

9.2.4 Summary

To summarize, for Second-Price Auctions, there is an ex-post Nash equilibrium in weakly dominant strategies and it is incentive-compatible (players are self-motivated to bid their private value). On the other hand, for First-Price Auctions, there are no ex-post Nash equilibria or dominant strategies.

9.3 Bayesian Games

Definition 9.7. A **Bayesian game** $G = (N, (A_i)_{i\in N}, (T_i)_{i\in N}, (u_i)_{i\in N}, P)$ where $(N, (A_i)_{i\in N}, (T_i)_{i\in N}, (u_i)_{i\in N})$ is a strict incomplete-information game and P is **distribution on types**, i.e.

- $N = \{1, \dots, n\}$ is the set of players;
- A_i is a set of **actions** available to player i;
- T_i is a set if **types** available to player i (recall that $T = \prod_{i=1}^{n} T_i$ is the set of all type profiles and $A = \prod_{i=1}^{n} A_i$ is the set of all action profiles);
- u_i is a type-dependent payoff function

$$u_i : A_1 \times \cdots \times A_n \times T_i \to \mathbb{R};$$

- P is a **(joint) prior distribution over** T called **common prior**.

> **Note**
>
> Formally, P is a probability measure over an appropriate measurable space on T. However, for simplicity, we will not go into measure theory and only consider two special cases: finite T (in which case $P : T \to [0, 1]$ so that $\sum_{\boldsymbol{t}\in T} P(\boldsymbol{t}) = 1$) and $T_i = \mathbb{R}$ for all i (in which case we assume that P is determined by a (joint) density function p on $\mathbb{R}^n$).

A play then proceeds as follows:

1. a type profile $\boldsymbol{t} = (t_1, \dots, t_n) \in T$ is randomly chosen according to P;
2. then each player learns **his type** t_i (and it is common knowledge that every player knows his own type but not the types of other players);
3. each player i chooses his action based on t_i;
4. each player receives his payoff $u_i(\boldsymbol{a}; t_i)$.

A **pure strategy** for player i is again a function $s_i : T_i \to A_i$. As before, we use S to denote the set of all pure strategy profiles. We also assume that u_i depends only on t_i and **not** on $\boldsymbol{t}_{-i}$. This is called the **private values** model and can be used to model auctions. This model can be extended to **common values** by using $u_i(\boldsymbol{a}; \boldsymbol{t})$.

Moreover, we assume the *common prior* P. This means that all players have the *same* beliefs about the type profiles. This assumption is rather strong. More general models allow players to have

- their own individual beliefs about types;

- ... beliefs about beliefs about types;
- .... beliefs about beliefs about beliefs about types;
- .... (we get an infinite hierarchy).

But there is a general result of Harsanyi saying that this complicated hierarchy is not necessary – it is possible to extend the type space in such a way that each player's "extended type" describes his original type as well as all of his beliefs.

Example 9.2. Assume that player 1 may suspect that player 2 is angry with him/her but cannot be sure. In other words, there are two types of player 2 giving two different games. So formally, we have Bayesian game $G = (N, (A_i)_{i\in N}, (T_i)_{i\in N}, (u_i)_{i\in N}, P)$ where

- $N = \{1, 2\}$;
- $A_1 = A_2 = \{O, F\}$;
- $T_1 = \{t_1\}$ and $T_2 = \{t_2^1, t_2^2\}$;
- the payoffs are given by Table 9.1;
- $P(t_1, t_2^1) = P(t_1, t_2^2) = \frac{1}{2}$.

Table 9.1: Rows are chosen by player 1 with type t_1

Table 9.2: The corresponding strategic-form game for $t_2 = t_2^1$

	F	O
F	$(2, 1)$	$(0, 0)$
O	$(0, 0)$	$(1, 2)$

Table 9.3: The corresponding strategic-form game for $t_2 = t_2^2$

	F	O
F	$(2, 0)$	$(0, 2)$
O	$(0, 1)$	$(1, 0)$

Example 9.3 (Bayesian sealed-bid Single-item auction)**.** Consider *single-item sealed-bid auctions* as Bayesian games: $G = (N, (B_i)_{i\in N}, (V_i)_{i\in N}, (u_i)_{i\in N}, P)$ where

- the set of players N is the set of bidders;
- $B_i = [0, \infty)$ where each action $b_i \in B_i$ corresponds to the bid b_i;
- $V_i = [0, \infty)$ where each type $v_i \in V_i$ corresponds to the private value v_i;
- let $v_i \in V_i$ be the type of player i (i.e. his private value), then given an action profile $\boldsymbol{b} = (b_1, \dots, b_n)$ (i.e. bids) we define
  - **First Price**
$$u_i(\boldsymbol{b}; v_i) = \begin{cases} v_i - b_i, & b_i > \max_{j\neq i} b_j, \\ 0, & \text{otherwise,} \end{cases}$$
  - **Second Price**
$$u_i(\boldsymbol{b}; v_i) = \begin{cases} v_i - \max_{j\neq i} b_j, & b_i > \max_{j\neq i} b_j, \\ 0, & \text{otherwise;} \end{cases}$$

- P is a probability distribution of the private values such that $P(\boldsymbol{v} \in [0, \infty)^n) = 1$. For example, we may (and will) assume that v_i is chosen independently and uniformly from $[0, v_{\max}]$, where $v_{\max}$ is a fixed given number. Then $P \sim \mathrm{Unif}([0, v_{\max}]^n)$.

For now, let us assume that each player has only finitely many types, i.e. T is finite. Given a type profile $\boldsymbol{t} = (t_1, \dots, t_n)$ we denote by $P(\boldsymbol{t}_{-i}|t_i)$ the *conditional probability* that the opponents of player i have types $\boldsymbol{t}_{-i}$ conditioned on player i having type t_i, i.e.

$$P(\boldsymbol{t}_{-i}|t_i) = \frac{P(t_i, \boldsymbol{t}_{-i})}{\sum_{\boldsymbol{t}'_{-i} \in T_{-i}} P(t_i, \boldsymbol{t}'_{-i})}.$$

> **Tip**
>
> Intuitively, $P(\boldsymbol{t}_{-i}|t_i)$ is the maximum information player i may squeeze out of P about possible types of others once he learns his own type t_i.

Given a pure strategy profile $\boldsymbol{s} = (s_1, \dots, s_n)$ and type $t_i \in T_i$ of player i the **expected payoff** for player i is

$$u_i(\boldsymbol{s}; t_i) = \sum_{\boldsymbol{t}_{-i} \in T_{-i}} u_i(\boldsymbol{s}(t_i, \boldsymbol{t}_{-i}); t_i) P(\boldsymbol{t}_{-i}|t_i).$$

> **Note**
>
> This is the conditional expected value of the measurable function $u_i \circ \boldsymbol{s} \circ (\boldsymbol{t}_{-i} \mapsto (t_i, \boldsymbol{t}_{-i}))$ given the random vector t_{-i} from the distribution $P(\cdot|t_i)$:
>
> $$u_i(\boldsymbol{s}; t_i) = \mathbb{E}\left(u_i(\boldsymbol{s}(t_i, \mathrm{t}_{-i}); t_i)\right).$$

Example 9.4. Let us continue with the Bayesian Battle of Sexes example, see Example 9.2, and recall $P(t_1, t_2^1) = P(t_1, t_2^2) = \frac{1}{2}$ and

Table 9.4: Rows are chosen by player 1 with type t_1

Table 9.5: The corresponding strategic-form game for $t_2 = t_2^1$

	F	O
F	$(2, 1)$	$(0, 0)$
O	$(0, 0)$	$(1, 2)$

Table 9.6: The corresponding strategic-form game for $t_2 = t_2^2$

	F	O
F	$(2, 0)$	$(0, 2)$
O	$(0, 1)$	$(1, 0)$

Consider strategies s_1 of player 1 and s_2 of player 2 defined by

- $s_1(t_1) = F$;
- $s_2(t_2^1) = F$ and $s_2(t_2^2) = O$.

Then

- $u_1(\boldsymbol{s}; t_1) = \frac{1}{2} \cdot 2 + \frac{1}{2} \cdot 0 = 1$;
- $u_2(\boldsymbol{s}; t_2^1) = 1$ and $u_2(\boldsymbol{s}; t_2^2) = 2$.

Example 9.5 (First-Price auction)**.** Consider the first-price auction as a Bayesian game where the types of players are chosen uniformly and independently from $[0, v_{\max}]$. Consider a pure strategy profile $\boldsymbol{v} = (v_1/2, \dots, v_n/2)$ (i.e. each player i plays $v_i/2$). Then

$$\begin{aligned} u_i(\boldsymbol{v}; v_i) &= P(\text{player } i \text{ wins}) \cdot v_i/2 + P(\text{player } i \text{ loses}) \cdot 0 \\ &= P(\text{all player except } i \text{ bid less than } v_i) \cdot v_i/2 \\ &\overset{\text{iid}}{=} \left(\frac{v_i}{v_{\max}}\right)^{n-1} \cdot v_i/2 \\ &= \frac{v_i^n}{2v_{\max}^{n-1}}. \end{aligned}$$

9.3.1 Risk Aversion

We assume that players maximize their expected payoff. Such players are called **risk neutral**. In general, there are three kinds of players that can be described using the following experiment. A player can choose between two possibilities: either get 50$ surely, or get 100$ with probability $\frac{1}{2}$ and 0 with probability $\frac{1}{2}$, then

- a *risk-neutral* player has no preference;
- a *risk-averse* player prefers the first alternative;
- a *risk-seeking* player prefers the second alternative.

9.3.2 Dominance and Nash Equilibria

A pure strategy s_i weakly dominates $s_i' \neq s_i$ if for every $t_i \in T_i$ satisfying $s_i(t_i) \neq s_i'(t_i)$ and every $\boldsymbol{s}_{-i} \in S_{-i}$ it holds that

$$u_i(s_i, \boldsymbol{s}_{-i}; t_i) \geq u_i(s_i', \boldsymbol{s}_{-i}; t_i)$$

and there exists at least one $\boldsymbol{s}_{-i}$ such that the inequality is strict.

> Tip
>
> The other modes of dominance are defined analogously. Dominant strategies are, too, defined as usual.

Definition 9.8. A pure strategy profile $\boldsymbol{s} = (s_1, \ldots, s_n) \in S$ in the Bayesian game is a **pure strategy Bayesian Nash equilibrium** (BNE) if for each player i and each type $t_i \in T_i$ of player i and every strategy $s_i' \in S_i$ we have that

$$u_i(\boldsymbol{s}; t_i) \geq u_i(s_i', \boldsymbol{s}_{-i}; t_i).$$

Example 9.6. Let us return to the Battle of Sexes as a Bayesian game, i.e. $P(t_1, t_2^1) = P(t_1, t_2^2) = \frac{1}{2}$ and

Table 9.7: Rows are chosen by player 1 with type t_1

Table 9.8: The corresponding strategic-form game for $t_2 = t_2^1$

	F	O
F	$(2, 1)$	$(0, 0)$
O	$(0, 0)$	$(1, 2)$

Table 9.9: The corresponding strategic-form game for $t_2 = t_2^2$

	F	O
F	$(2, 0)$	$(0, 2)$
O	$(0, 1)$	$(1, 0)$

We will use the following notation in this example: $(X, (Y, Z))$ means that player 1 players $X \in \{F, O\}$, and player 2 plays $Y \in \{F, O\}$ if their type is t_2^1 and $Z \in \{F, O\}$ otherwise. It is easy to check that $(F, (F, O))$ is a pure strategy Bayesian Nash equilibrium (BNE). Even though O is preferred by player 2, the outcome (O, O) cannot occur with a positive probability in any BNE, as

- to ever meet at the opera, player 1 needs to play O;
- the unique best response of player 2 to O is (O, F);
- but $(O, (O, F))$ is not a Bayesian Nash equilibrium, because
  - the excepted payoff of player 1 at $(O, (O, F))$ is $\frac{1}{2}$;
  - on the other hand, the expected payoff of player 1 at $(F, (F, O))$ is 1.

9.3.3 Auctions as Bayesian Games

Consider the second-price sealed-bid auction as a Bayesian game where the types of players are chosen according to an arbitrary distribution.

Theorem 9.2. *In a second-price sealed-bid auction, with any probability distribution P, the truth revealing profile of bids, i.e. $\boldsymbol{v} = (v_1, \ldots, v_n)$, is a weakly dominant strategy profile.*

Proof. The same exact proof as for the strict incomplete-information games can be repeated, as we did not assume the players to have a common prior. □

Now consider the first-price sealed-bid auction as a Bayesian game with some prior distribution P. Note that bidding truthfully does **not** have to be a dominant strategy. For example, if player i knows that (with high probability) his value v_i is much larger than $\max_{j \neq i} v_j$, he will not *waste* money and bid less than v_i. So is there a pure strategy Bayesian Nash equilibrium?

Theorem 9.3. *Assume that for all players i the type of player i is chosen independently and uniformly from* $[0, v_{\max}]$*. Consider a pure strategy profile* $\mathbf{s} = (s_1, \dots, s_n) \in S$ *where* $s_i(v_i) = \frac{n-1}{n} v_i$ *for every player i and every private value* v_i*. Then* $\mathbf{s}$ *is a Bayesian Nash equilibrium.*

Hence we can see that there are, in some sense, "optimal" strategies for players in both types of auctions, but then a question arises which one is better for the auctioneer (in terms of expected revenue)?

Consider the first and second price sealed-bid auctions. For simplicity, assume that the type (their private value) of each player is chosen independently and uniformly from $[0, 1]$, then

- in the first-price auction, players bid $\frac{n-1}{n} v_i$ Thus the probability distribution (*the cumulative distribution function*) of the revenue is

$$F(x) = P\left(\max_j \frac{n-1}{n} v_j \le x\right) = P\left(\max_j v_j \le \frac{nx}{n-1}\right) = \left(\frac{nx}{n-1}\right)^n.$$

  It is then straightforward to show that the expected maximum bid in the first-price auction (i.e. the revenue) is $\frac{n-1}{n+1}$;
- in the second-price auction, players bid v_i. However, the revenue is the expected second-largest value. Thus the distribution of the revenue[2] is

$$F(x) = \underbrace{P\left(\max_j v_j \le x\right)}_{\text{tie or too small bids}} + \underbrace{\sum_{i=1}^{n} P\left(v_i > x \ \land \ \forall j \ne i : v_j \le x\right)}_{\text{one of the players bids enough}}.$$

  Amazingly, this also gives the expectation $\frac{n-1}{n+1}$.

This result is a special case of a rather general **revenue equivalence theorem**, first proven by Vickrey (1961) and then generalized by Myerson (1981).

> **Note**
>
> Both Vickrey and Myerson were awarded Nobel Price in economics for their contribution to the auction theory.

Theorem 9.4 (Revenue Equivalence)**.** *Assume that each of n risk-neutral players has independent private values drawn from a common cumulative distribution function* $F(x)$ *which is continuous and strictly increasing on an interval* $[v_{\min}, v_{\max}]$ *(the probability* $v_i \notin [v_{\min}, v_{\max}]$ *is zero). Then any efficient auction mechanism in which any player with value* $v_{\min}$ *has expected payoff zero yields the same expected revenue.*

Here *efficient* means that the auction has a symmetric and increasing Bayesian Nash equilibrium in each of v_i and always allocates the item to the player with the highest bid.

[2] For more information about the expected revenues, see this document.

Note

The Theorem 9.4 is not necessary for the exam.