

IA168 Algorithmic Game Theory

Tomáš Brázdil


Organization of This Course

Sources:

- ▶ Lectures (slides, notes)
 - ▶ based on several sources
 - ▶ slides are prepared for lectures, some stuff on greenboard (⇒ attend the lectures)
- ▶ Books:
 - ▶ Nisan/Roughgarden/Tardos/Vazirani, **Algorithmic Game Theory**, Cambridge University, 2007.
Available online for free:
http://www.cambridge.org/journals/nisan/downloads/Nisan_Non-printable.pdf
 - ▶ Tadelis, **Game Theory: An Introduction**, Princeton University Press, 2013

(I use various resources, so please, attend the lectures)

- ▶ Oral exam
- ▶ Homework



What is Algorithmic Game Theory?

First, what is the game theory?

According to the Oxford dictionary it is "the branch of mathematics concerned with the analysis of strategies for dealing with competitive situations where the outcome of a participant's choice of action depends critically on the actions of other participants"

According to Myerson it is "the study of mathematical models of conflict and cooperation between intelligent rational decision-makers"






What does the "algorithmic" mean?

- ▶ It means that we are "concerned with the computational questions that arise in game theory, and that enlighten game theory. In particular, questions about finding efficient algorithms to 'solve' games."

Let's have a look at some examples

Prisoner's Dilemma

Prisoners' dilemma

		prisoner B			
		confess		remain silent	
prisoner A	confess	 5 years 5 years	 0 year 20 years		
	remain silent	 20 years 0 year	 1 year 1 year		

© 2006 Encyclopædia Britannica, Inc.

- ▶ Two suspects of a serious crime are arrested and imprisoned.
- ▶ Police has 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).

Sentence depends on the behavior of both suspects.

The problem: What would the suspects do?

Prisoner's Dilemma – Solution(?)

	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 5 years sentence.

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

Are there always "dominant" strategies?

Nash equilibria – Battle of Sexes



- ▶ A couple agreed to meet this evening, but cannot recall if they will be attending the opera or a football match.
- ▶ The husband would like to go to the football game. The wife would like to go to the opera. Both would prefer to go to the same place rather than different ones.

If they cannot communicate, where should they go?

Nash equilibria – Battle of Sexes

Battle of Sexes can be modeled as a game of two players (Wife, Husband) with the following payoffs:

	<i>O</i>	<i>F</i>
<i>O</i>	2,1	0,0
<i>F</i>	0,0	1,2


Apparently, no strategy of any player is dominant. A “solution”?

Note that whenever *both* players play *O*, then neither of them wants to *unilaterally* deviate from his strategy!

(O, O) is an example of a *Nash equilibrium* (as is (F, F))

Mixed Equilibria – 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 is an example of *zero-sum* games: whatever one of the players wins, the other one loses.
- ▶ What is an optimal behavior here? Is there a Nash equilibrium?

Use *mixed strategies*: Each player plays each pure strategy with probability $1/3$. The expected payoff of each player is 0 (even if one of the players changes his strategy, he still gets 0!).

How to algorithmically solve games in mixed strategies? (we shall use probability theory and linear programming)


Philosophical Issues in Games

I UNDERSTAND THAT SCISSORS CAN BEAT PAPER, AND I GET HOW ROCK CAN BEAT SCISSORS, BUT THERE'S NO WAY PAPER CAN BEAT ROCK. PAPER IS SUPPOSED TO MAGICALLY WRAP AROUND ROCK LEAVING IT IMMOBILE? WHY CAN'T PAPER DO THIS TO SCISSORS? SCREW SCISSORS, WHY CAN'T PAPER DO THIS TO PEOPLE? WHY AREN'T SHEETS OF COLLEGE RULED NOTEBOOK PAPER CONSTANTLY SUFFOCATING STUDENTS AS THEY ATTEMPT TO TAKE NOTES IN CLASS? I'LL TELL YOU WHY, BECAUSE PAPER CAN'T BEAT ANYBODY, A ROCK WOULD TEAR IT UP IN TWO SECONDS. WHEN I PLAY ROCK PAPER SCISSORS, I ALWAYS CHOOSE ROCK. THEN WHEN SOMEBODY CLAIMS TO HAVE BEATEN ME WITH THEIR PAPER I CAN PUNCH THEM IN THE FACE WITH MY ALREADY CLENCHED FIST AND SAY, OH SORRY, I THOUGHT PAPER WOULD PROTECT YOU.

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:




How to "solve" such games?

What is their relationship to the strategic form games?

Chance and Imperfect Information

Some decisions in the game tree may be by chance and 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 opponent’s cards hidden).



Again, how to solve such games?

Games of Incomplete Information

According to a study by the Institute of incomplete information 9 out of every 10.

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

Example: 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 the player 2 **does not know** u_1 .

How to deal with such a game? Assume the “worst” private value? What if we have a partial knowledge about the private values?

Inefficiency of Equilibria

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

	C	S
C	-5, -5	0, -20
S	-20, 0	-1, -1

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

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.


Inefficiency of Equilibria – 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 measure the ratio between average travel time in these two cases.

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



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 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, 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.

Games, the Internet and E-commerce: An extremely active research area at the intersection of CS and Economics

Basic idea: “The internet is a HUGE experiment in interaction between agents (both human and automated)”

How do we set up the rules of this game to harness “socially optimal” results?

Summary and Brief Overview

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 (in particular, Lemke-Howson algorithm for computing Nash Eq.)
- ▶ 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.

Static Games of Complete Information

Strategic-Form Games

Solution concepts

Static Games of Complete Information – Intuition

Proceed in two steps:

1. Each player *simultaneously and independently* chooses a *strategy*. 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 1

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 common knowledge).

Strategic-Form Games

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

Definition 2

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:

- ▶ $N = \{1, 2, \dots, n\}$ 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* is a vector of strategies of all players $(s_1, \dots, s_n) \in S_1 \times \dots \times S_n$.

We denote the set of all strategy profiles by $S = S_1 \times \dots \times S_n$.

- ▶ $u_i : S \rightarrow \mathbb{R}$ is a function associating each strategy profile $s = (s_1, \dots, s_n) \in S$ with the *payoff* $u_i(s)$ to player i , for every player $i \in N$.

Definition 3

A *zero-sum* game G is one in which for all $s = (s_1, \dots, s_n) \in S$ we have $u_1(s) + u_2(s) + \dots + u_n(s) = 0$.

Example: Prisoner's Dilemma

- ▶ $N = \{1, 2\}$
- ▶ $S_1 = S_2 = \{S, C\}$
- ▶ u_1, u_2 are defined as follows:
 - ▶ $u_1(C, C) = -5, u_1(C, S) = 0, u_1(S, C) = -20,$
 $u_1(S, S) = -1$
 - ▶ $u_2(C, C) = -5, u_2(C, S) = -20, u_2(S, C) = 0,$
 $u_2(S, S) = -1$

(Is it zero sum?)

We usually write payoffs in the following form:

	C	S
C	-5, -5	0, -20
S	-20, 0	-1, -1

or as two matrices:

	C	S
C	-5	0
S	-20	-1

	C	S
C	-5	-20
S	0	-1

Example: Cournot Duopoly

- ▶ 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 $\kappa - q_1 - q_2$ (here κ is a positive constant)
- ▶ Firms 1 and 2 have per item production costs c_1 and c_2 , resp.

Question: How these firms are going to behave?

We may model the situation using a strategic-form game.

Strategic-form game model $(N, (S_i)_{i \in N}, (u_i)_{i \in N})$

- ▶ $N = \{1, 2\}$
- ▶ $S_i = [0, \infty)$
- ▶ $u_1(q_1, q_2) = q_1(\kappa - q_1 - q_2) - q_1 c_1$
 $u_2(q_1, q_2) = q_2(\kappa - q_1 - q_2) - q_2 c_2$

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 term *equilibrium* for any one of the strategy profiles that emerges as one of the solution concepts' predictions.

(I follow the approach of Steven Tadelis here, it is not completely standard)

Example 4

Nash equilibrium is a solution concept. That is, we “solve” games by finding Nash equilibria and declare them to be reasonable outcomes.

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 a common knowledge among them.
4. **Self-enforcement**: Any prediction (or equilibrium) of a solution concept must be *self-enforcing*.

Here 4. 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.

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?): Solution concept should apply to a wide variety of games.
E.g. We prove 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.

The basic notion for evaluating "social outcome" is the following

Definition 5

A strategy profile $s \in S$ *Pareto dominates* a strategy profile $s' \in S$ if $u_i(s) \geq u_i(s')$ for all $i \in N$, and $u_i(s) > u_i(s')$ for at least one $i \in N$.

A strategy profile $s \in S$ is *Pareto optimal* if it is not Pareto dominated by any other strategy profile.

We will see more measures of social outcome later.

Solution Concepts – Pure Strategies

We will consider the following solution concepts:

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

For now, let us concentrate on

pure strategies only!

I.e., no mixed strategies are allowed. We will generalize to mixed setting later.

- ▶ Let $N = \{1, \dots, n\}$ be a finite set and for each $i \in N$ let X_i be a set. Let $X := \prod_{i \in N} X_i = \{(x_1, \dots, x_n) \mid x_j \in X_j, j \in N\}$.
 - ▶ For $i \in N$ we define $X_{-i} := \prod_{j \neq i} X_j$, i.e.,

$$X_{-i} = \{(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n) \mid x_j \in X_j, \forall j \neq i\}$$

- ▶ An element of X_{-i} will be denoted by

$$x_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$$

We slightly abuse notation and write (x_i, x_{-i}) to denote $(x_1, \dots, x_i, \dots, x_n) \in X$.

Strict Dominance in Pure Strategies

Definition 6

Let $s_i, s'_i \in S_i$ be strategies of player i . Then s'_i is *strictly dominated* by s_i (write $s_i > s'_i$) if for any possible combination of the other players' strategies, $s_{-i} \in S_{-i}$, we have

$$u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i}) \quad \text{for all } s_{-i} \in S_{-i}$$

Claim 1

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

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

Strictly Dominant Strategy Equilibrium in Pure Str.

Definition 7

$s_i \in S_i$ is *strictly dominant* if every other pure strategy of player i is strictly dominated by s_i .

Observe that every player has at most one strictly dominant strategy, and that strictly dominant strategies do not have to exist.

Claim 2

Any rational player will play the strictly dominant strategy (if it exists).

Definition 8

A strategy profile $s \in S$ is a *strictly dominant strategy equilibrium* if $s_i \in S_i$ is strictly dominant for all $i \in N$.

Corollary 9

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

Is the strictly dominant strategy equilibrium always Pareto optimal?

Examples

In the Prisoner's dilemma:

	<i>C</i>	<i>S</i>
<i>C</i>	-5, -5	0, -20
<i>S</i>	-20, 0	-1, -1

(*C, C*) is the strictly dominant strategy equilibrium (the only profile that is not Pareto optimal!).

In the Battle of Sexes:

	<i>O</i>	<i>F</i>
<i>O</i>	2, 1	0, 0
<i>F</i>	0, 0	1, 2

no strictly dominant strategies exist.

Indiana Jones and the Last Crusade

(Taken from Dixit & Nalebuff's "The Art of Strategy" and a lecture of Robert Marks)

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.

Indy Goofed

- ▶ Although this scene adds excitement, it is somewhat embarrassing that such a distinguished professor as Dr. Indiana Jones would overlook his dominant strategy.
- ▶ He should have given the water to his father without testing it first.
 - ▶ 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.

Iterated Strict Dominance in Pure Strategies

We know that 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 a 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 a 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.

The previous reasoning yields the **Iterated Elimination of Strictly Dominated Strategies (IESDS)**:

Define a sequence $D_i^0, D_i^1, D_i^2, \dots$ of strategy sets of player i .
(Denote by G_{DS}^k the game obtained from G by restricting to $D_i^k, i \in N$.)

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. 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 $k = 0, 1, 2, \dots$

Definition 10

A strategy profile $s = (s_1, \dots, s_n) \in S$ 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.

IESDS Examples

In the Prisoner's dilemma:

	<i>C</i>	<i>S</i>
<i>C</i>	-5, -5	0, -20
<i>S</i>	-20, 0	-1, -1

(*C, C*) is the only one surviving the first round of IESDS.

In the Battle of Sexes:

	<i>O</i>	<i>F</i>
<i>O</i>	2, 1	0, 0
<i>F</i>	0, 0	1, 2

all strategies survive all rounds (i.e. IESDS \equiv anything may happen, sorry)

A Bit More Interesting Example

	<i>L</i>	<i>C</i>	<i>R</i>
<i>L</i>	4,3	5,1	6,2
<i>C</i>	2,1	8,4	3,6
<i>R</i>	3,0	9,6	2,8


IESDS on greenboard!

Political Science Example: Median Voter Theorem

Hotelling (1929) and Downs (1957)

- ▶ $N = \{1, 2\}$
- ▶ $S_i = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$ (political and ideological spectrum)
- ▶ 10 voters belong to each position
(Here 10 means ten percent in the real-world)
- ▶ Voters vote for the closest candidate. If there is a tie, then $\frac{1}{2}$ go to each candidate
- ▶ Payoff: The number of voters for the candidate, each candidate (selfishly) strives to maximize this number

Political Science Example: Median Voter Theorem



Candidate A



Candidates must choose to position themselves at one of the ten ideological locations. Voters are evenly distributed along the ideological spectrum, i.e. 10% at each location.



Candidate B

- ▶ 1 and 10 are the (only) strictly dominated strategies \Rightarrow
 $D_1^1 = D_2^1 = \{2, \dots, 9\}$
- ▶ in G_{DS}^1 , 2 and 9 are the (only) strictly dominated strategies \Rightarrow
 $D_1^2 = D_2^2 = \{3, \dots, 8\}$
- ▶ ...
- ▶ only 5, 6 survive IESDS

Belief & Best Response

IESDS eliminated apparently unreasonable behavior (leaving "reasonable" behavior implicitly untouched).

What if we rather want to actively preserve reasonable behavior?
What is reasonable? what we believe is reasonable :-).

Intuition:

- ▶ 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 reasonableness of the belief)

Let us formalize this type of reasoning

Belief & Best Response

Definition 11

A *belief* of player i is a pure strategy profile $s_{-i} \in S_{-i}$ of his opponents.

Definition 12

A strategy $s_i \in S_i$ of player i is a *best response* to a belief $s_{-i} \in S_{-i}$ if

$$u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i}) \text{ for all } s'_i \in S_i$$

Claim 3

A rational player who believes that his opponents will play $s_{-i} \in S_{-i}$ always chooses a best response to $s_{-i} \in S_{-i}$.

Definition 13

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

A rational player never plays any strategy that is never best response.

Best Response vs Strict Dominance

Proposition 1

If s_i is strictly dominated for player i , then it is never best response.

The opposite does not have to be true in pure strategies:

	X	Y
A	1, 1	1, 1
B	2, 1	0, 1
C	0, 1	2, 1

Here A is never best response but is strictly dominated neither by B, nor by C.

Elimination of Stupid Strategies = Rationalizability

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

Define a sequence $R_i^0, R_i^1, R_i^2, \dots$ of strategy sets of player i .
(Denote by G_{Rat}^k the game obtained from G by restricting to $R_i^k, i \in N$.)

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. 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 $k = 0, 1, 2, \dots$

Definition 14

A strategy profile $s = (s_1, \dots, s_n) \in S$ is a *rationalizable equilibrium* if each s_i is rationalizable.

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

(Warning: For some reasons, rationalizable strategies are almost always defined using mixed strategies!)

Rationalizability Examples

In the Prisoner's dilemma:

	<i>C</i>	<i>S</i>
<i>C</i>	-5, -5	0, -20
<i>S</i>	-20, 0	-1, -1

(*C*, *C*) is the only rationalizable equilibrium.

In the Battle of Sexes:

	<i>O</i>	<i>F</i>
<i>O</i>	2, 1	0, 0
<i>F</i>	0, 0	1, 2

all strategies are rationalizable.

Cournot Duopoly

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

- ▶ $N = \{1, 2\}$
- ▶ $S_i = [0, \infty)$
- ▶ $u_1(q_1, q_2) = q_1(\kappa - q_1 - q_2) - q_1 c_1 = (\kappa - c_1)q_1 - q_1^2 - q_1 q_2$
 $u_2(q_1, q_2) = q_2(\kappa - q_2 - q_1) - q_2 c_2 = (\kappa - c_2)q_2 - q_2^2 - q_2 q_1$

Assume for simplicity that $c_1 = c_2 = c$ and denote $\theta = \kappa - c$.

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

Solve $\frac{\delta u_1}{\delta q_1} = \theta - 2q_1 - q_2 = 0$, 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 $q_2 \geq 0$, we obtain that q_1 is never best response iff $q_1 > \theta/2$.

Similarly q_2 is never best response iff $q_2 > \theta/2$.

Thus $R_1^1 = R_2^1 = [0, \theta/2]$.

Cournot Duopoly

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

▶ $N = \{1, 2\}$

▶ $S_i = [0, \infty)$

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

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

Assume for simplicity that $c_1 = c_2 = c$ and denote $\theta = \kappa - c$.

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 resp. 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 $q_2 \in [0, \theta/4)$

Thus $R_1^2 = R_2^2 = [\theta/4, \theta/2]$.

....

Cournot Duopoly (cont.)

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

- ▶ $N = \{1, 2\}$
- ▶ $S_i = [0, \infty)$
- ▶ $u_1(q_1, q_2) = q_1(\kappa - q_1 - q_2) - q_1 c_1 = (\kappa - c_1)q_1 - q_1^2 - q_1 q_2$
 $u_2(q_1, q_2) = q_2(\kappa - q_2 - q_1) - q_2 c_2 = (\kappa - c_2)q_2 - q_2^2 - q_2 q_1$

Assume for simplicity that $c_1 = c_2 = c$ and denote $\theta = \kappa - c$.

In general, after $2k$ iterations we have $R_i^{2k} = R_i^{2k} = [\ell_k, r_k]$ where

- ▶ $r_k = (\theta - \ell_{k-1})/2$ for $k \geq 1$
- ▶ $\ell_k = (\theta - r_k)/2$ for $k \geq 1$ and $\ell_0 = 0$

Solving the recurrence we obtain

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

Hence, $\lim_{k \rightarrow \infty} \ell_k = \lim_{k \rightarrow \infty} r_k = \theta/3$ and thus $(\theta/3, \theta/3)$ is the only rationalizable equilibrium.

Cournot Duopoly (cont.)

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

▶ $N = \{1, 2\}$

▶ $S_i = [0, \infty)$

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

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

Assume for simplicity that $c_1 = c_2 = c$ and denote $\theta = \kappa - c$.

Are $q_i = \theta/3$ Pareto optimal? NO!

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

but

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

IESDS vs Rationalizability in Pure Strategies

Theorem 15

Assume that S is finite. Then for all k we have that $R_i^k \subseteq D_i^k$. That is, in particular, all rationalizable strategies survive IESDS.

The opposite inclusion does not have to be true in pure strategies:

	X	Y
A	1, 1	1, 1
B	2, 1	0, 1
C	0, 1	2, 1

Recall that A is never best response but is strictly dominated by neither B , nor C . That is, A survives IESDS but is not rationalizable.

Proof of Theorem 15

By induction on k . For $k = 0$ we have that $R_i^0 = S_i = D_i^0$ by definition. 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 $s_j \in R_i^{k+1}$. Then there must be $s_{-j} \in R_{-j}^k$ such that

s_j is a best response to s_{-j} in G_{Rat}^k

(This follows from the fact that s_j has not been eliminated in G_{Rat}^k .)

But then s_j is a best response to s_{-j} in G_{Rat}^{k-1} as well!

Indeed, let s'_j be a best response to s_{-j} in G_{Rat}^{k-1} . Then $s'_j \in R_i^k$ since s'_j is not eliminated in G_{Rat}^{k-1} . But then $u_i(s_j, s_{-j}) \geq u_i(s'_j, s_{-j})$ since s_j is a best response to s_{-j} in G_{Rat}^k . Thus s_j is a best response to s_{-j} in G_{Rat}^{k-1} .

By the same reason, s_j is a best response to s_{-j} in G_{Rat}^{k-2} .

By the same reason, s_j is a best response to s_{-j} in G_{Rat}^{k-3} .

...

By the same reason, s_j is a best response to s_{-j} in $G_{Rat}^0 = G$.

However, then s_j is a best response to s_{-j} in G_{DS}^k .

(This follows from the fact that the “best response” relationship of s_j and s_{-j} is preserved by removing arbitrarily many other strategies.)

Thus s_j is not strictly dominated in G_{DS}^k and $s_j \in D_i^{k+1}$. □

Pinning Down Beliefs – Nash Equilibria

Criticism of previous approaches:

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

Typical example is Battle of Sexes:

	<i>O</i>	<i>F</i>
<i>O</i>	2,1	0,0
<i>F</i>	0,0	1,2

Here all strategies are equally reasonable according to the above concepts.

But are all strategy profiles really equally reasonable?

Pinning Down Beliefs – Nash Equilibria

	O	F
O	2,1	0,0
F	0,0	1,2

Assume that each player has a belief about strategies of other players.

By Claim 3, each player plays a best response to his beliefs.

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!**

(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

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 beliefs of his opponents.

A usual definition is following:

Definition 16

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

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*) \quad \text{for all } s_i \in S_i \text{ and all } i \in N$$

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

Nash Equilibria Examples

In the Prisoner's dilemma:

	<i>C</i>	<i>S</i>
<i>C</i>	-5, -5	0, -20
<i>S</i>	-20, 0	-1, -1

(C, C) is the only Nash equilibrium.

In the Battle of Sexes:

	<i>O</i>	<i>F</i>
<i>O</i>	2, 1	0, 0
<i>F</i>	0, 0	1, 2

only (O, O) and (F, F) are Nash equilibria.

In Cournot Duopoly, $(\theta/3, \theta/3)$ is the only Nash equilibrium.

(Best response relations: $q_1 = (\theta - q_2)/2$ and $q_2 = (\theta - q_1)/2$ are both satisfied only by $q_1 = q_2 = \theta/3$)

Example: Stag Hunt

Story:

- ▶ 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



- ▶ Hunting stag is much more demanding and forces of both players need to be joined (hare can be hunted individually)

Strategy-form game model: $N = \{1, 2\}$, $S_1 = S_2 = \{S, H\}$, the payoff:

	S	H
S	5,5	0,3
H	3,0	3,3

Two NE: (S, S), and (H, H), where the former Pareto dominates the latter! Which one is more reasonable?

Example: Stag Hunt

Strategy-form game model: $N = \{1, 2\}$, $S_1 = S_2 = \{S, H\}$, the payoff:

	S	H
S	5,5	0,3
H	3,0	3,3

Two NE: (S, S) , and (H, H) , where the former Pareto dominates the latter! Which one is more reasonable?

If each player believes that the other one will go for hare, then (H, H) is a reasonable outcome \Rightarrow a society of individualists who do not cooperate at all.

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.

This is supposed to explain that in 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).

Example: Stag Hunt

Strategy-form game model: $N = \{1, 2\}$, $S_1 = S_2 = \{S, H\}$, the payoff:

	S	H
S	5,5	0,3
H	3,0	3,3

Two NE: (S, S) , and (H, H) , where the former Pareto dominates the latter! Which one is more reasonable?

Another point of view: (H, H) is less risky

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) (?)

Nash Equilibria vs Previous Concepts

Theorem 17

1. *If 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.*

Proof: Homework!

Corollary 18

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

Interpretations of Nash Equilibria

Except the two definitions, usual interpretations are following:

- ▶ When the goal is to give advice to all of the players in a game (i.e., to advise each player 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.
- ▶ 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.

Static Games of Complete Information

Mixed Strategies

Let's Mix It

As pointed out before, neither of the solution concepts has to exist in pure strategies

Example: Rock-Paper-scissors

	<i>R</i>	<i>P</i>	<i>C</i>
<i>R</i>	0,0	-1,1	1,-1
<i>P</i>	1,-1	0,0	-1,1
<i>C</i>	-1,1	1,-1	0,0

There are no strictly dominant pure strategies

No strategy is strictly dominated (IESDS removes nothing)

Each strategy is a best response to some strategy of the opponent (rationalizability removes nothing)

No pure Nash equilibria: 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

Probability Distributions

Definition 19

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

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

We denote by $\text{supp}(\sigma)$ the *support* of σ , that is the set of all $a \in A$ satisfying $\sigma(a) > 0$.

Example 20

Consider $A = \{a, b, c\}$ and a function $\sigma : A \rightarrow [0, 1]$ such that $\sigma(a) = \frac{1}{4}$, $\sigma(b) = \frac{3}{4}$, and $\sigma(c) = 0$. Then $\sigma \in \Delta(A)$ and $\text{supp}(\sigma) = \{a, b\}$.

Mixed Strategies

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

From now on, **assume that all S_i are finite!**

Definition 21

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 $\Sigma := \Sigma_1 \times \cdots \times \Sigma_n$, the set of all *mixed strategy profiles*.

Recall that by Σ_{-i} we denote the set $\Sigma_1 \times \cdots \times \Sigma_{i-1} \times \Sigma_{i+1} \times \cdots \times \Sigma_n$

Elements of Σ_{-i} are denoted by $\sigma_{-i} = (\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_n)$.

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).

For example, in rock-paper-scissors, the pure strategy R corresponds

to σ_i which satisfies $\sigma_i(X) = \begin{cases} 1 & X = R \\ 0 & \text{otherwise} \end{cases}$

Mixed Strategies

Sometimes we assume $S_i = \{1, \dots, m_i\}$, here $m_i \in \{1, 2, \dots\}$, for all $i \in N$.

Then every mixed strategy σ_i is a vector $\sigma_i = (\sigma_i(1), \dots, \sigma_i(m_i))^T \in [0, 1]^{m_i}$ so that

$$\sigma_i(1) + \dots + \sigma_i(m_i) = 1$$

Mixed Strategy Profiles

Let $\sigma = (\sigma_1, \dots, \sigma_n)$ be a mixed strategy profile.

Intuitively, we assume that each player i *randomly* chooses his pure strategy according to σ_i and *independently* of his opponents.

Thus for $\mathbf{s} = (s_1, \dots, s_n) \in S = S_1 \times \dots \times S_n$ we have that

$$\sigma(\mathbf{s}) := \prod_{i=1}^n \sigma_i(s_i)$$

is the probability that the players choose the pure strategy profile \mathbf{s} according to the mixed strategy profile σ , and

$$\sigma_{-i}(\mathbf{s}_{-i}) := \prod_{k \neq i}^n \sigma_k(s_k)$$

is the probability that the opponents of player i choose $\mathbf{s}_{-i} \in S_{-i}$ when they play according to the mixed strategy profile $\sigma_{-i} \in \Sigma_{-i}$.

(We abuse notation a bit here: σ denotes two things, a vector of mixed strategies as well as a probability distribution on S (the same for σ_{-i})

Mixed Strategies – Example

	<i>R</i>	<i>P</i>	<i>C</i>
<i>R</i>	0,0	-1,1	1,-1
<i>P</i>	1,-1	0,0	-1,1
<i>C</i>	-1,1	1,-1	0,0

An example of a mixed strategy σ_1 : $\sigma_1(R) = \frac{1}{2}$, $\sigma_1(P) = \frac{1}{3}$, $\sigma_1(C) = \frac{1}{6}$.

Sometimes we write σ_1 as $(\frac{1}{2}(R), \frac{1}{3}(P), \frac{1}{6}(C))$, or only $(\frac{1}{2}, \frac{1}{3}, \frac{1}{6})$ if the order of pure strategies is fixed.

Consider a mixed strategy profile (σ_1, σ_2) where $\sigma_1 = (\frac{1}{2}(R), \frac{1}{3}(P), \frac{1}{6}(C))$ and $\sigma_2 = (\frac{1}{3}(R), \frac{2}{3}(P), 0(C))$.

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

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

Expected Payoff

... but now what is the suitable notion of payoff?

Definition 22

The *expected payoff* of player i under a mixed strategy profile $\sigma \in \Sigma$ is

$$u_i(\sigma) := \sum_{s \in S} \sigma(s) u_i(s) \quad \left(= \sum_{s \in S} \prod_{k=1}^n \sigma_k(s_k) u_i(s) \right)$$

I.e., it is the "weighted average" of what player i wins under each pure strategy profile s , weighted by the probability of that profile.

Assumption: Every rational player strives to maximize his own expected payoff.

(This assumption is not always completely convincing ...)

Expected Payoff – Example

Matching Pennies:

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

Each player secretly turns a penny to 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.

Consider $\sigma_1 = (\frac{1}{3}(H), \frac{2}{3}(T))$ and $\sigma_2 = (\frac{1}{4}(H), \frac{3}{4}(T))$

$$\begin{aligned}u_1(\sigma_1, \sigma_2) &= \sum_{(X,Y) \in \{H,T\}^2} \sigma_1(X)\sigma_2(Y)u_1(X, Y) \\ &= \frac{1}{3} \frac{1}{4} 1 + \frac{1}{3} \frac{3}{4} (-1) + \frac{2}{3} \frac{1}{4} (-1) + \frac{2}{3} \frac{3}{4} 1 = \frac{1}{6}\end{aligned}$$

$$\begin{aligned}u_2(\sigma_1, \sigma_2) &= \sum_{(X,Y) \in \{H,T\}^2} \sigma_1(X)\sigma_2(Y)u_2(X, Y) \\ &= \frac{1}{3} \frac{1}{4} (-1) + \frac{1}{3} \frac{3}{4} 1 + \frac{2}{3} \frac{1}{4} 1 + \frac{2}{3} \frac{3}{4} (-1) = -\frac{1}{6}\end{aligned}$$

"Decomposition" of Expected Payoff

Consider the matching pennies example from the previous slide:

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

together with some mixed strategies σ_1 and σ_2 .

We prove the following important property of the expected payoff:

$$u_1(\sigma_1, \sigma_2) = \sum_{X \in \{H, T\}} \sigma_1(X) u_1(X, \sigma_2)$$

An intuition behind this equality is following:

- ▶ $u_1(\sigma_1, \sigma_2)$ is the expected payoff of player 1 in the following experiment: Both players simultaneously and independently choose their pure strategies X, Y according to σ_1, σ_2 , resp., and then player 1 collects his payoff $u_1(X, Y)$.
- ▶ $\sum_{X \in \{H, T\}} \sigma_1(X) u_1(X, \sigma_2)$ is the expected payoff of player 1 in the following: Player 1 chooses his *pure* strategy X and then uses it against the mixed strategy σ_2 of player 2. Then player 2 chooses Y according to σ_2 *independently of* X , and player 1 collects the payoff $u_1(X, Y)$.

As Y does not depend on X in neither experiment, we obtain the above equality of expected payoffs.

"Decomposition" of Expected Payoff

Consider the matching pennies example from the previous slide:

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

together with some mixed strategies σ_1 and σ_2 .

A formal proof is straightforward:

$$\begin{aligned}u_1(\sigma_1, \sigma_2) &= \sum_{(X,Y) \in \{H,T\}^2} \sigma_1(X)\sigma_2(Y)u_1(X, Y) \\&= \sum_{X \in \{H,T\}} \sum_{Y \in \{H,T\}} \sigma_1(X)\sigma_2(Y)u_1(X, Y) \\&= \sum_{X \in \{H,T\}} \sigma_1(X) \sum_{Y \in \{H,T\}} \sigma_2(Y)u_1(X, Y) \\&= \sum_{X \in \{H,T\}} \sigma_1(X)u_1(X, \sigma_2)\end{aligned}$$

(In the last equality we used the fact that X is identified with a mixed strategy assigning one to X .)

"Decomposition" of Expected Payoff

Consider the matching pennies example from the previous slide:

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

together with some mixed strategies σ_1 and σ_2 .

Similarly,

$$\begin{aligned}u_1(\sigma_1, \sigma_2) &= \sum_{(X,Y) \in \{H,T\}^2} \sigma_1(X)\sigma_2(Y)u_1(X, Y) \\&= \sum_{X \in \{H,T\}} \sum_{Y \in \{H,T\}} \sigma_1(X)\sigma_2(Y)u_1(X, Y) \\&= \sum_{Y \in \{H,T\}} \sum_{X \in \{H,T\}} \sigma_1(X)\sigma_2(Y)u_1(X, Y) \\&= \sum_{Y \in \{H,T\}} \sigma_2(Y) \sum_{X \in \{H,T\}} \sigma_1(X)u_1(X, Y) \\&= \sum_{Y \in \{H,T\}} \sigma_2(Y)u_1(\sigma_1, Y)\end{aligned}$$

Expected Payoff – "Decomposition" in General

Lemma 23

For every mixed strategy profile $\sigma \in \Sigma$ and every $k \in N$ we have

$$u_i(\sigma) = \sum_{s_k \in S_k} \sigma_k(s_k) \cdot u_i(s_k, \sigma_{-k}) = \sum_{s_{-k} \in S_{-k}} \sigma_{-k}(s_{-k}) \cdot u_i(\sigma_k, s_{-k})$$

Proof:

$$\begin{aligned} u_i(\sigma) &= \sum_{s \in S} \sigma(s) u_i(s) = \sum_{s \in S} \prod_{\ell=1}^n \sigma_\ell(s_\ell) u_i(s) \\ &= \sum_{s \in S} \sigma_k(s_k) \prod_{\ell \neq k} \sigma_\ell(s_\ell) u_i(s) \\ &= \sum_{s_k \in S_k} \sum_{s_{-k} \in S_{-k}} \sigma_k(s_k) \prod_{\ell \neq k} \sigma_\ell(s_\ell) u_i(s_k, s_{-k}) \\ &= \sum_{s_k \in S_k} \sum_{s_{-k} \in S_{-k}} \sigma_k(s_k) \sigma_{-k}(s_{-k}) u_i(s_k, s_{-k}) \end{aligned}$$

Proof of Lemma 23 (cont.)

The first equality:

$$\begin{aligned}u_i(\sigma) &= \sum_{\mathbf{s}_k \in \mathcal{S}_k} \sum_{\mathbf{s}_{-k} \in \mathcal{S}_{-k}} \sigma_k(\mathbf{s}_k) \sigma_{-k}(\mathbf{s}_{-k}) u_i(\mathbf{s}_k, \mathbf{s}_{-k}) \\&= \sum_{\mathbf{s}_k \in \mathcal{S}_k} \sigma_k(\mathbf{s}_k) \sum_{\mathbf{s}_{-k} \in \mathcal{S}_{-k}} \sigma_{-k}(\mathbf{s}_{-k}) u_i(\mathbf{s}_k, \mathbf{s}_{-k}) \\&= \sum_{\mathbf{s}_k \in \mathcal{S}_k} \sigma_k(\mathbf{s}_k) u_i(\mathbf{s}_k, \sigma_{-k})\end{aligned}$$

The second equality:

$$\begin{aligned}u_i(\sigma) &= \sum_{\mathbf{s}_k \in \mathcal{S}_k} \sum_{\mathbf{s}_{-k} \in \mathcal{S}_{-k}} \sigma_k(\mathbf{s}_k) \sigma_{-k}(\mathbf{s}_{-k}) u_i(\mathbf{s}_k, \mathbf{s}_{-k}) \\&= \sum_{\mathbf{s}_{-k} \in \mathcal{S}_{-k}} \sum_{\mathbf{s}_k \in \mathcal{S}_k} \sigma_k(\mathbf{s}_k) \sigma_{-k}(\mathbf{s}_{-k}) u_i(\mathbf{s}_k, \mathbf{s}_{-k}) \\&= \sum_{\mathbf{s}_{-k} \in \mathcal{S}_{-k}} \sigma_{-k}(\mathbf{s}_{-k}) \sum_{\mathbf{s}_k \in \mathcal{S}_k} \sigma_k(\mathbf{s}_k) u_i(\mathbf{s}_k, \mathbf{s}_{-k}) \\&= \sum_{\mathbf{s}_{-k} \in \mathcal{S}_{-k}} \sigma_{-k}(\mathbf{s}_{-k}) u_i(\sigma_k, \mathbf{s}_{-k})\end{aligned}$$

Expected Payoff – Pure Strategy Bounds

Corollary 24

For all $i, k \in N$ and $\sigma \in \Sigma$ we have that

- ▶ $\min_{s_k \in S_k} u_i(s_k, \sigma_{-k}) \leq u_i(\sigma) \leq \max_{s_k \in S_k} u_i(s_k, \sigma_{-k})$
- ▶ $\min_{s_{-k} \in S_{-k}} u_i(\sigma_k, s_{-k}) \leq u_i(\sigma) \leq \max_{s_{-k} \in S_{-k}} u_i(\sigma_k, s_{-k})$

Proof.

We prove $u_i(\sigma) \leq \max_{s_k \in S_k} u_i(s_k, \sigma_{-k})$ the rest is similar. Define $B := \max_{s_k \in S_k} u_i(s_k, \sigma_{-k})$. Then

$$\begin{aligned} u_i(\sigma) &= \sum_{s_k \in S_k} \sigma_k(s_k) \cdot u_i(s_k, \sigma_{-k}) \\ &\leq \sum_{s_k \in S_k} \sigma_k(s_k) \cdot B \\ &= B \end{aligned}$$

□

Solution Concepts

We revisit the following solution concepts in mixed strategies:

- ▶ strict dominant strategy equilibrium
- ▶ IESDS equilibrium
- ▶ rationalizable equilibria
- ▶ Nash equilibria

From now on, when I say a *strategy* I 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 $|u_i| = \sum_{s \in S} |u_i(s)|$ and $|u_i(s)|$ is the length of a binary encoding of $u_i(s)$ (we assume that rational numbers are encoded as quotients of two binary integers)

Note that, in particular, $|G| > |S|$.

Strict Dominance in Mixed Strategies

Definition 25

Let $\sigma_i, \sigma'_i \in \Sigma_i$ be (mixed) strategies of player i . Then σ'_i is *strictly dominated* by σ_i (write $\sigma'_i < \sigma_i$) if

$$u_i(\sigma_i, \sigma_{-i}) > u_i(\sigma'_i, \sigma_{-i}) \quad \text{for all } \sigma_{-i} \in \Sigma_{-i}$$

Example 26

	X	Y
A	3	0
B	0	3
C	1	1

Is there a strictly dominated strategy?

Question: Is there a game with at least one strictly dominated strategy but without strictly dominated *pure* strategies?

Strictly Dominant Strategy Equilibrium

Definition 27

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

Definition 28

A strategy profile $\sigma \in \Sigma$ is a *strictly dominant strategy equilibrium* if $\sigma_i \in \Sigma_i$ is strictly dominant for all $i \in N$.

Proposition 2

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

Proof.

Let $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*) \in \Sigma_i$ be the strictly dominant strategy equilibrium.

By Corollary 24, for every $i \in N$ and $\sigma_{-i} \in \Sigma_{-i}$, there must exist $s_i \in S_i$ such that $u_i(\sigma_i^*, \sigma_{-i}) \leq u_i(s_i, \sigma_{-i})$.

But then $\sigma_i^* = s_i$ since σ_i^* is strictly dominant.



Computing Strictly Dominant Strategy Equilibrium

How to decide whether there is a strictly dominant strategy equilibrium $s = (s_1, \dots, s_n) \in S$?

I.e. whether for a given $s_i \in S_i$, all $\sigma_i \in \Sigma_i \setminus \{s_i\}$ and all $\sigma_{-i} \in \Sigma_{-i}$:

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

There are some serious issues here:

Obviously there are uncountably many possible σ_i and σ_{-i} .

$u_i(\sigma_i, \sigma_{-i})$ is nonlinear, and for more that two players even $u_i(s_i, \sigma_{-i})$ is nonlinear in probabilities assigned to pure strategies.

Computing Strictly Dominant Strategy Equilibrium

First, we prove the following useful proposition using Lemma 23:

Lemma 29

σ'_i strictly dominates σ_i **iff** for all pure strategy profiles $s_{-i} \in S_{-i}$:

$$u_i(\sigma'_i, s_{-i}) > u_i(\sigma_i, s_{-i}) \quad (1)$$

Proof.

' \Rightarrow ' direction is trivial, let us prove ' \Leftarrow '. Assume that (1) is true for all pure strategy profiles $s_{-i} \in S_{-i}$. Then, by Lemma 23,

$$u_i(\sigma_i, \sigma_{-i}) = \sum_{s_{-i} \in S_{-i}} \sigma_{-i}(s_{-i}) u_i(\sigma_i, s_{-i}) < \sum_{s_{-i} \in S_{-i}} \sigma_{-i}(s_{-i}) u_i(\sigma'_i, s_{-i}) = u_i(\sigma'_i, \sigma_{-i})$$

holds for all mixed strategy profiles $\sigma_{-i} \in \Sigma_{-i}$. □

In other words, it suffices to check the strict dominance only with respect to all *pure* profiles of opponents.

Computing Strictly Dominant Strategy Equilibrium

How to decide whether for a given $s_i \in S_i$, all $\sigma_i \in \Sigma_i \setminus \{s_i\}$ and all $s_{-i} \in S_{-i}$ we have $u_i(s_i, s_{-i}) > u_i(\sigma_i, s_{-i})$.

Lemma 30

$u_i(s_i, s_{-i}) > u_i(\sigma_i, s_{-i})$ for all $\sigma_i \in \Sigma_i \setminus \{s_i\}$ and all $s_{-i} \in S_{-i}$ **iff**
 $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$ for all $s'_i \in S_i \setminus \{s_i\}$ and all $s_{-i} \in S_{-i}$.

Proof.

' \Rightarrow ' direction is trivial, let us prove ' \Leftarrow '. Assume $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$ for all $s'_i \in S_i \setminus \{s_i\}$ and all $s_{-i} \in S_{-i}$. Given $\sigma_i \in \Sigma_i \setminus \{s_i\}$, we have by Lemma 23,

$$u_i(\sigma_i, s_{-i}) = \sum_{s'_i \in S_i} \sigma_i(s'_i) u_i(s'_i, s_{-i}) < \sum_{s'_i \in S_i} \sigma_i(s'_i) u_i(s_i, s_{-i}) = u_i(s_i, s_{-i})$$

The inequality follows from our assumption and the fact that $\sigma_i(s'_i) > 0$ for at least one $s'_i \neq s_i$ (due to $\sigma_i \in \Sigma_i \setminus \{s_i\}$). \square

Thus it suffices to check whether $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$ for all $s'_i \in S_i$ and all $s_{-i} \in S_{-i}$. This can easily be done in time polynomial w.r.t. $|G|$.

IESDS in Mixed Strategies

Define a sequence $D_i^0, D_i^1, D_i^2, \dots$ of strategy sets of player i .
(Denote by G_{DS}^k the game obtained from G by restricting the pure strategy sets to $D_i^k, i \in N$.)

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. 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 $k = 0, 1, 2, \dots$

Definition 31

A strategy profile $s = (s_1, \dots, s_n) \in S$ is an *IESDS equilibrium* if each s_i survives IESDS.

Note that in step 2 it is not sufficient to consider pure strategies.
Consider the following zero sum game:

	X	Y
A	3	0
B	0	3
C	1	1

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

However, there are uncountably many mixed strategies that may dominate a given pure strategy ...

But $u_i(\sigma) = u_i(\sigma_1, \dots, \sigma_n)$ is linear in each σ_k (if σ_{-k} is kept fixed)!

Indeed, assuming w.l.o.g. that $S_k = \{1, \dots, m_k\}$,

$$u_i(\sigma) = \sum_{s_k \in S_k} \sigma_k(s_k) \cdot u_i(s_k, \sigma_{-k}) = \sum_{\ell=1}^{m_k} \sigma_k(\ell) \cdot u_i(\ell, \sigma_{-k})$$

is the scalar product of the vector $\sigma_k = (\sigma_k(1), \dots, \sigma_k(m_k))$ with the vector $(u_i(1, \sigma_{-k}), \dots, u_i(m_k, \sigma_{-k}))$, which is linear.

So to decide strict dominance, we use linear programming ...

Intermezzo: Linear Programming

Linear programming is a technique for 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:

$$\text{maximize } \sum_{j=1}^m c_j x_j \quad (\text{objective function})$$

$$\text{subject to } \sum_{j=1}^m a_{ij} x_j \leq b_i \quad 1 \leq i \leq n \quad (\text{constraints})$$

$$x_j \geq 0 \quad 1 \leq j \leq m$$

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 which maximizes the *objective function* $\sum_{j=1}^m c_j x_j$.

Intermezzo: Complexity of Linear Programming

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 32 (Khachiyan, Doklady Akademii Nauk SSSR, 1979)

There is an algorithm which for any linear program computes an optimal solution in polynomial time.

The algorithm uses 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.

For more info see

http://en.wikipedia.org/wiki/Linear_programming#Solvers_and_scripting_.28programming.29_languages

IESDS Algorithm – Strict Dominance Step

So how do we use linear programming to decide strict dominance in step 2 of IESDS procedure?

I.e. whether for a given s_i there exists σ_i such that for all σ_{-i} we have

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

Recall that by Lemma 29 we have that σ_i strictly dominates s_i **iff** for all *pure strategy profiles* $s_{-i} \in S_{-i}$:

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

In other words, it suffices to check the strict dominance only with respect to all *pure* profiles of opponents.

IESDS Algorithm – Strict Dominance Step

Recall that $u_i(\sigma_i, \mathbf{s}_{-i}) = \sum_{\mathbf{s}'_i \in \mathcal{S}_i} \sigma_i(\mathbf{s}'_i) u_i(\mathbf{s}'_i, \mathbf{s}_{-i})$.

So to decide whether $\mathbf{s}_i \in \mathcal{S}_i$ is strictly dominated by some mixed strategy σ_i , it suffices to solve the following system:

$$\begin{aligned} \sum_{\mathbf{s}'_i \in \mathcal{S}_i} x_{\mathbf{s}'_i} \cdot u_i(\mathbf{s}'_i, \mathbf{s}_{-i}) &> u_i(\mathbf{s}_i, \mathbf{s}_{-i}) && \mathbf{s}_{-i} \in \mathcal{S}_{-i} \\ x_{\mathbf{s}'_i} &\geq 0 && \mathbf{s}'_i \in \mathcal{S}_i \\ \sum_{\mathbf{s}'_i \in \mathcal{S}_i} x_{\mathbf{s}'_i} &= 1 \end{aligned}$$

(Here each variable $x_{\mathbf{s}'_i}$ corresponds to the probability $\sigma_i(\mathbf{s}'_i)$ assigned by the strictly dominant strategy σ_i to \mathbf{s}'_i)

Unfortunately, this is a "strict linear program" ... How to deal with the strict inequality?

IESDS Algorithm – Complexity

Introduce a new variable y to be **maximized** under the following constraints:

$$\sum_{s'_i \in S_i} x_{s'_i} \cdot u_i(s'_i, s_{-i}) \geq u_i(s_i, s_{-i}) + y \quad s_{-i} \in S_{-i}$$

$$x_{s'_i} \geq 0 \quad s'_i \in S_i$$

$$\sum_{s'_i \in S_i} x_{s'_i} = 1$$

$$y \geq 0$$

Now s_i is strictly dominated **iff** a solution maximizing y satisfies $y > 0$

The size of the above program is polynomial in $|G|$.

So the step 2 of IESDS can be executed in polynomial time.

As every iteration of IESDS removes at least one pure strategy, IESDS runs in time polynomial in $|G|$.

IESDS in Mixed Strategie – Example

	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 the linear program for deciding whether C is strictly dominated: The program maximizes y under the following constraints:

$$3x_A + 0x_B + x_C \geq 1 + y$$

$$0x_A + 3x_B + x_C \geq 1 + y$$

$$x_A, x_B, x_C \geq 0$$

$$x_A + x_B + x_C = 1$$

$$y \geq 0$$

The maximum $y = \frac{1}{2}$ is attained at $x_A = \frac{1}{2}$ and $x_B = \frac{1}{2}$.

Definition 33

A strategy $\sigma_i \in \Sigma_i$ of player i is a *best response* to a strategy profile $\sigma_{-i} \in \Sigma_{-i}$ of his opponents if

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i}) \quad \text{for all } \sigma'_i \in \Sigma_i$$

We denote by $BR_i(\sigma_{-i}) \subseteq \Sigma_i$ the set of all best responses of player i to the strategy profile of opponents $\sigma_{-i} \in \Sigma_{-i}$.

Best Response – Example

Consider a game with the following payoffs of 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).

Compute $BR_1(q)$.

Rationalizability in Mixed Strategies (Two Players)

For simplicity, we temporarily switch to **two-player** setting $N = \{1, 2\}$.

Definition 34

A *(mixed) belief* of player $i \in \{1, 2\}$ is a mixed strategy σ_{-i} of his opponent.

(A general definition works with so called *correlated beliefs* that are arbitrary distributions on S_{-i} , the notion of the expected payoff needs to be adjusted, we are not going in this direction)

Assumption: Any rational player with a belief σ_{-i} always plays a best response to σ_{-i} .

Definition 35

A strategy $\sigma_i \in \Sigma_i$ of player $i \in \{1, 2\}$ is *never best response* if it is not a best response to any belief σ_{-i} .

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

Rationalizability in Mixed Strategies (Two Players)

Define a sequence $R_i^0, R_i^1, R_i^2, \dots$ of strategy sets of player i .
(Denote by G_{Rat}^k the game obtained from G by restricting the pure strategy sets to $R_i^k, i \in N$.)

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 $k = 0, 1, 2, \dots$

Definition 36

A strategy profile $s = (s_1, \dots, s_n) \in S$ is a *rationalizable equilibrium* if each s_i is rationalizable.

Rationalizability vs IESDS (Two Players)

	X	Y
A	3	0
B	0	3
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)

What strategies of player 1 are never best responses?

What strategies of player 1 are strictly dominated?

Observation: The set of strictly dominated strategies coincides with the set of never best responses!

... and this holds in general for two player games:

Theorem 37

Assume $N = \{1, 2\}$. A pure strategy s_i is never best response to any belief $\sigma_{-i} \in \Sigma_{-i}$ **iff** s_i is strictly dominated by a strategy $\sigma_i \in \Sigma_i$.

It follows that a strategy of S_i survives IESDS **iff** it is rationalizable.

(The theorem is true also for an arbitrary number of players but correlated beliefs need to be used.)

Mixed Nash Equilibrium

Definition 38

A mixed-strategy profile $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*) \in \Sigma$ is a **(mixed) Nash equilibrium** if σ_i^* is a best response to σ_{-i}^* for each $i \in N$, that is

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*) \quad \text{for all } \sigma_i \in \Sigma_i \text{ and all } i \in N$$

An interpretation: each σ_{-i}^* can be seen as a belief of player i against which he plays a best response σ_i^* .

Given a mixed strategy profile of opponents $\sigma_{-i} \in \Sigma_{-i}$, we denote by $BR_i(\sigma_{-i})$ the set of all $\sigma_i \in \Sigma_i$ that are best responses to σ_{-i} .

Then σ^* is a Nash equilibrium iff $\sigma_i^* \in BR_i(\sigma_{-i}^*)$ for all $i \in N$.

Theorem 39 (Nash 1950)

Every finite game in strategic form has a Nash equilibrium.

This is THE fundamental theorem of game theory.

Example: Matching Pennies

	<i>H</i>	<i>T</i>
<i>H</i>	1, -1	-1, 1
<i>T</i>	-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?

$$u_1(H, q) = 2q - 1 \text{ and } u_1(T, q) = 1 - 2q$$

Then

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

We obtain the best-response correspondence BR_1 :

$$BR_1(q) = \begin{cases} p = 0 & \text{if } q < \frac{1}{2} \\ p \in [0, 1] & \text{if } q = \frac{1}{2} \\ p = 1 & \text{if } q > \frac{1}{2} \end{cases}$$

Example: Matching Pennies

	<i>H</i>	<i>T</i>
<i>H</i>	1, -1	-1, 1
<i>T</i>	-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.

Similarly for player 2 :

$$u_2(p, H) = 1 - 2p \text{ and } u_2(p, T) = 2p - 1$$

$$u_2(p, q) = qu_2(p, H) + (1 - q)u_2(p, T) = q(1 - 2p) + (1 - q)(2p - 1)$$

We obtain best-response relation BR_2 :

$$BR_2(p) = \begin{cases} q = 1 & \text{if } p < \frac{1}{2} \\ q \in [0, 1] & \text{if } p = \frac{1}{2} \\ q = 0 & \text{if } p > \frac{1}{2} \end{cases}$$

The only "intersection" of BR_1 and BR_2 is the only Nash equilibrium

$$\sigma_1 = \sigma_2 = \left(\frac{1}{2}, \frac{1}{2}\right).$$

Static Games of Complete Information
Mixed Strategies
Computing Nash Equilibria – Support Enumeration

Computing Mixed Nash Equilibria

Lemma 40

$\sigma^* = (\sigma_1^*, \dots, \sigma_n^*) \in \Sigma$ is a Nash equilibrium **iff** there exist $w_1, \dots, w_n \in \mathbb{R}$ such that the following holds:

- ▶ For all $i \in N$ and all $s_i \in \text{supp}(\sigma_i^*)$ we have $u_i(s_i, \sigma_{-i}^*) = w_i$.
- ▶ For all $i \in N$ and all $s_i \notin \text{supp}(\sigma_i^*)$ we have $u_i(s_i, \sigma_{-i}^*) \leq w_i$.

Here, the right hand side implies $u_i(\sigma^*) = w_i$.

Proof.

The fact that the right hand side implies $u_i(\sigma^*) = w_i$ follows immediately from Lemma 23:

$$\begin{aligned} u_i(\sigma^*) &= \sum_{s_i \in \mathcal{S}_i} \sigma^*(s_i) u_i(s_i, \sigma_{-i}^*) = \sum_{s_i \in \text{supp}(\sigma_i^*)} \sigma^*(s_i) u_i(s_i, \sigma_{-i}^*) \\ &= \sum_{s_i \in \text{supp}(\sigma_i^*)} \sigma^*(s_i) w_i = w_i \sum_{s_i \in \text{supp}(\sigma_i^*)} \sigma^*(s_i) = w_i \end{aligned}$$

Computing Mixed Nash Equilibria

Lemma 41

$\sigma^* = (\sigma_1^*, \dots, \sigma_n^*) \in \Sigma$ is a Nash equilibrium **iff** there exist $w_1, \dots, w_n \in \mathbb{R}$ such that the following holds:

- ▶ For all $i \in N$ and all $s_i \in \text{supp}(\sigma_i^*)$ we have $u_i(s_i, \sigma_{-i}^*) = w_i$.
- ▶ For all $i \in N$ and all $s_i \notin \text{supp}(\sigma_i^*)$ we have $u_i(s_i, \sigma_{-i}^*) \leq w_i$.

Here, the right hand side implies $u_i(\sigma^*) = w_i$.

Proof. (Cont.)

" \Leftarrow ": Use the first equality of Lemma 23 to obtain for every $i \in N$ and every $\sigma'_i \in \Sigma_i$

$$\begin{aligned} u_i(\sigma'_i, \sigma_{-i}^*) &= \sum_{s_i \in \mathcal{S}_i} \sigma'_i(s_i) u_i(s_i, \sigma_{-i}^*) \leq \\ &\leq \sum_{s_i \in \mathcal{S}_i} \sigma'_i(s_i) w_i = \sum_{s_i \in \mathcal{S}_i} \sigma'_i(s_i) u_i(\sigma^*) = u_i(\sigma^*) \end{aligned}$$

Thus σ^* is a Nash equilibrium.

Computing Mixed Nash Equilibria

Lemma 42

$\sigma^* = (\sigma_1^*, \dots, \sigma_n^*) \in \Sigma$ is a Nash equilibrium **iff** there exist $w_1, \dots, w_n \in \mathbb{R}$ such that the following holds:

- ▶ For all $i \in N$ and all $s_i \in \text{supp}(\sigma_i^*)$ we have $u_i(s_i, \sigma_{-i}^*) = w_i$.
- ▶ For all $i \in N$ and all $s_i \notin \text{supp}(\sigma_i^*)$ we have $u_i(s_i, \sigma_{-i}^*) \leq w_i$.

Here, the right hand side implies $u_i(\sigma^*) = w_i$.

Proof (Cont.)

Idea for " \Rightarrow ": Let $w_i := u_i(\sigma^*)$.

Clearly, every $i \in N$ and $s_i \in S_i$ satisfy $u_i(s_i, \sigma_{-i}^*) \leq u_i(\sigma^*) = w_i$.

By Corollary 24, there is *at least one* $s_i \in \text{supp}(\sigma_i^*)$ satisfying $u_i(s_i, \sigma_{-i}^*) = u_i(\sigma^*) = w_i$.

Now if there is $s'_i \in \text{supp}(\sigma_i^*)$ such that

$$u_i(s'_i, \sigma_{-i}^*) < u_i(\sigma^*) \quad (= u_i(s_i, \sigma_{-i}^*))$$

then increasing the probability $\sigma_i^*(s_i)$ and decreasing (in proportion) $\sigma_i^*(s'_i)$ strictly increases $u_i(\sigma^*)$, a contradiction with σ^* being NE.

Example: Matching Pennies

	<i>H</i>	<i>T</i>
<i>H</i>	1, -1	-1, 1
<i>T</i>	-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.

There are no pure strategy equilibria.

There are no equilibria where only player 1 randomizes:

Indeed, assume that (p, H) is such an equilibrium. Then by Lemma 42,

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

a contradiction. Also, (p, T) cannot be an equilibrium.

Similarly, there is no NE where only player 2 randomizes.

Example: Matching Pennies

	<i>H</i>	<i>T</i>
<i>H</i>	1, -1	-1, 1
<i>T</i>	-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.

Assume that both players randomize, i.e., $p, q \in (0, 1)$.

The expected payoffs of playing pure strategies for player 1:

$$u_1(H, q) = 2q - 1 \text{ and } u_1(T, q) = 1 - 2q$$

Similarly for player 2 :

$$u_2(p, H) = 1 - 2p \text{ and } u_2(p, T) = 2p - 1$$

By Lemma 42, Nash equilibria must satisfy:

$$2q - 1 = 1 - 2q \quad \text{and} \quad 1 - 2p = 2p - 1$$

That is $p = q = \frac{1}{2}$ is the only Nash equilibrium.

Example: Battle of Sexes

	<i>O</i>	<i>F</i>
<i>O</i>	2, 1	0, 0
<i>F</i>	0, 0	1, 2

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 q).

Compute all Nash equilibria.

There are two pure strategy equilibria $(2, 1)$ and $(1, 2)$, no Nash equilibrium where only one player randomizes.

Now assume that

- ▶ 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)

where $p, q \in (0, 1)$.

By Lemma 42, any Nash equilibrium must satisfy:

$$2q = 1 - q \quad \text{and} \quad p = 2(1 - p)$$

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

An Algorithm?

What did we do in the previous examples?

We went through all support combinations for both players.

(pure, one player mixing, both mixing)

For each pair of supports 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 computation of Nash equilibria to *linear equations*.

Support Enumeration (Idea)

Recall Lemma 42: $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*) \in \Sigma$ is a Nash equilibrium **iff** there exist $w_1, \dots, w_n \in \mathbb{R}$ such that the following holds:

- ▶ For all $i \in N$ and all $s_i \in \text{supp}(\sigma_i^*)$ we have $u_i(s_i, \sigma_{-i}^*) = w_i$.
- ▶ For all $i \in N$ and all $s_i \notin \text{supp}(\sigma_i^*)$ we have $u_i(s_i, \sigma_{-i}^*) \leq w_i$.

Suppose that we somehow know the supports $\text{supp}(\sigma_1^*), \dots, \text{supp}(\sigma_n^*)$ for some Nash equilibrium $\sigma_1^*, \dots, \sigma_n^*$ (which itself is unknown to us).

Now we may consider all $\sigma_i^*(s_i)$'s and all w_i 's as variables and use the above conditions to design a system of inequalities capturing Nash equilibria with the given support sets $\text{supp}(\sigma_1^*), \dots, \text{supp}(\sigma_n^*)$.

Support Enumeration

To simplify notation, assume that for every 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 .

Fix supports $supp_i \subseteq S_i$ for every $i \in N$ and consider the following system of constraints with variables

$\sigma_1(1), \dots, \sigma_1(m_1), \dots, \sigma_n(1), \dots, \sigma_n(m_n), w_1, \dots, w_n$:

1. For all $i \in N$ and all $k \in supp_i$ we have

$$(u_i(k, \sigma_{-i}) =) \quad \sum_{s \in S \wedge s_i = k} \left(\prod_{j \neq i} \sigma_j(s_j) \right) u_i(s) = w_i$$

2. For all $i \in N$ and all $k \notin supp_i$ we have

$$(u_i(k, \sigma_{-i}) =) \quad \sum_{s \in S \wedge s_i = k} \left(\prod_{j \neq i} \sigma_j(s_j) \right) u_i(s) \leq w_i$$

3. For all $i \in N$: $\sigma_i(1) + \dots + \sigma_i(m_i) = 1$.
4. For all $i \in N$ and all $k \in supp_i$: $\sigma_i(k) \geq 0$.
5. For all $i \in N$ and all $k \notin supp_i$: $\sigma_i(k) = 0$.

Support Enumeration

Consider the system of constraints from the previous slide.

The following lemma follows immediately from Lemma 42.

Lemma 43

Let $\sigma^ \in \Sigma$ be a strategy profile.*

- ▶ *If σ^* is a Nash equilibrium and $\text{supp}(\sigma_i^*) = \text{supp}_i$ for all $i \in N$, then assigning $\sigma_i(k) := \sigma_i^*(k)$ and $w_i := u_i(\sigma^*)$ solves the system.*
- ▶ *If $\sigma_i(k) := \sigma_i^*(k)$ and $w_i := u_i(\sigma^*)$ solves the system, then σ^* is a Nash equilibrium with $\text{supp}(\sigma_i^*) \subseteq \text{supp}_i$ for all $i \in N$.*

Support Enumeration (Two Players)

The constraints are *non-linear* in general, but *linear* for two player games! Let us stick to two players.

How to find $supp_1$ and $supp_2$? ... Just guess!

Input: A two-player strategic-form 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 .

Output: A Nash equilibrium σ^* .

Algorithm: For all possible $supp_1 \subseteq S_1$ and $supp_2 \subseteq S_2$:

- ▶ Check if the corresponding system of linear constraints (from the previous slide) has a feasible solution $\sigma^*, w_1^*, \dots, w_n^*$.
- ▶ If so, STOP: the feasible solution σ^* is a Nash equilibrium satisfying $u_i(\sigma^*) = w_i^*$.

Question: How many possible subsets $supp_1, supp_2$ are there to try?

Answer: $2^{(m_1+m_2)}$

So, unfortunately, the algorithm requires worst-case exponential time.

Remarks on Support Enumeration

- ▶ The algorithm combined with Theorem 39 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 $w_1 + \dots + w_n$. More precisely, the algorithm can be modified as follows:

- ▶ Initialize $W := -\infty$ (W stores the current maximum welfare)
- ▶ For all possible $supp_1 \subseteq S_1$ and $supp_2 \subseteq S_2$:
 - ▶ Find the maximum value $\max(\sum w_i)$ of $w_1 + \dots + w_n$ so that the constraints are satisfiable (using linear programming).
 - ▶ Put $W := \max\{W, \max(\sum w_i)\}$.
- ▶ Return W .

Remarks on Support Enumeration (Cont.)

Similar trick works for any notion of "good" NE that can be expressed using a linear objective function and (additional) linear constraints in variables $\sigma_i(j)$ and w_j .

(e.g., maximize payoff of player 1, minimize payoff of player 2 and keep probability of playing the strategy 1 below 1/2, etc.)

Complexity Results – (Two Players)

Theorem 44

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

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

Membership to NP follows from the support enumeration:

For example, for 1., it suffices to guess supports $supp_1, supp_2$ and add $w_1 \geq v$ to the constraints; the resulting NE σ^* satisfies $u_1(\sigma^*) \geq v$.

Complexity Results (Two Players)

Theorem 45

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

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

NP-hardness can be proved using reduction from SAT

(The reduction is not difficult but we are not going into it.)

It is presented in "New Complexity Results about Nash Equilibria" by V. Conitzer and T. Sandholm (pages 6–8)

The Reduction (It's Short and Sweet)

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 variable⁶), and C its set of clauses. The function $v : L \rightarrow V$ 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, \dots, 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.

... But 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 complexity of the sample equilibrium problem.

We use complexity classes of *function problems* such as FP, FNP, etc.

The support enumeration gives a deterministic algorithm which runs in exponential time. Can we do better?

In what follows we show that

- ▶ the sample equilibrium problem can be solved in polynomial time for zero-sum two-player games,
(Using a beautiful characterization of all Nash equilibria)
- ▶ the sample equilibrium problem belongs to the complexity class PPAD (which is a subclass of FNP) for two-player games.
(... to be defined later)

Is there a better characterization of Nash equilibria than Lemma 42 ?

Definition 46

$\sigma_i^* \in \Sigma_i$ is a *maxmin* strategy of player i if

$$\sigma_i^* \in \operatorname{argmax}_{\sigma_i \in \Sigma_i} \min_{\sigma_{-i} \in \Sigma_{-i}} u_i(\sigma_i, \sigma_{-i})$$

(Intuitively, a *maxmin* strategy σ_i^* maximizes player 1's worst-case payoff in the situation where player 2 strives to cause the greatest harm to player 1.)

(Since u_i is continuous and Σ_{-i} compact, $\min_{\sigma_{-i} \in \Sigma_{-i}} u_i(\sigma_i, \sigma_{-i})$ is well defined and continuous on Σ_i , which implies that there is at least one maxmin strategy.)

Lemma 47

σ_i^* is maxmin iff

$$\sigma_i^* \in \operatorname{argmax}_{\sigma_i \in \Sigma_i} \min_{s_{-i} \in S_{-i}} u_i(\sigma_i, s_{-i})$$

Proof.

By Corollary 24, for every $\sigma \in \Sigma$ we have $u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma_i, s_{-i})$ for some $s_{-i} \in S_{-i}$.

Thus $\min_{\sigma_{-i} \in \Sigma_{-i}} u_i(\sigma_i, \sigma_{-i}) = \min_{s_{-i} \in S_{-i}} u_i(\sigma_i, s_{-i})$. Hence,

$$\operatorname{argmax}_{\sigma_i \in \Sigma_i} \min_{\sigma_{-i} \in \Sigma_{-i}} u_i(\sigma_i, \sigma_{-i}) = \operatorname{argmax}_{\sigma_i \in \Sigma_i} \min_{s_{-i} \in S_{-i}} u_i(\sigma_i, s_{-i})$$

□

Question: Assume a strategy profile where both players play their maxmin strategies? Does it have to be a Nash equilibrium?

Zero-Sum Games: von Neumann's Theorem

Assume that G is zero sum, i.e., $u_1 = -u_2$.

Then $\sigma_2^* \in \Sigma_2$ is maxmin of player 2 **iff**

$$\sigma_2^* \in \operatorname{argmin}_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2) \quad (= \operatorname{argmin}_{\sigma_2 \in \Sigma_2} \max_{s_1 \in S_1} u_1(s_1, \sigma_2))$$

(Intuitively, maxmin of player 2 minimizes the payoff of player 1 when player 1 plays his best responses. Such strategy of player 2 is often called minmax.)

Theorem 48 (von Neumann)

Assume a two-player **zero-sum** game. Then

$$\max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2) = \min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2)$$

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

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

Proof of Theorem 48 (Homework)

Homework: Prove von Neumann's Theorem in 4 easy steps:

1. Prove this inequality:

$$\max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2) \leq \min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2)$$

2. Prove that (σ_1^*, σ_2^*) is a Nash equilibrium iff

$$\min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1^*, \sigma_2) \geq u_1(\sigma_1^*, \sigma_2^*) \geq \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2^*)$$

Hint: One of the inequalities is trivial and the other one almost.

3. Use 1. and 2. together with Theorem 39 to prove

$$\max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2) \geq \min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2)$$

4. Use the above to prove the rest of the theorem.

Hint: Use the characterization of NE from 2., do not forget that you already have $\max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2) = \min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2)$

You may already have proved one of the implications when proving 3.

Zero-Sum Two-Player Games – Computing NE

Assume $S_1 = \{1, \dots, m_1\}$ and $S_2 = \{1, \dots, m_2\}$.

We want to compute

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

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

maximize: v

subject to:
$$\sum_{k=1}^{m_1} \sigma_1(k) \cdot u_1(k, \ell) \geq v \quad \ell = 1, \dots, m_2$$

$$\sum_{k=1}^{m_1} \sigma_1(k) = 1$$

$$\sigma_1(k) \geq 0 \quad k = 1, \dots, m_1$$

Lemma 49

$\sigma_1^* \in \operatorname{argmax}_{\sigma_1 \in \Sigma_1} \min_{\ell \in S_2} u_1(\sigma_1, \ell)$ **iff** assigning $\sigma_1(k) := \sigma_1^*(k)$ and $v := \min_{\ell \in S_2} u_1(\sigma_1^*, \ell)$ gives an optimal solution.

Zero-Sum Two-Player Games – Computing NE

Summary:

- ▶ We have reduced computation of NE to 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.

Strategic-Form Games – Conclusion

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

Strategic-Form Games – Conclusion

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, IESDS preserves the set of Nash equilibria
3. In finite games, rationalizability preserves Nash equilibria

In mixed setting:

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

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

Algorithms

- ▶ 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 and linear programming)
 - ▶ in exponential time using support enumeration
 - ▶ in PPAD using Lemke-Howson

Complexity of Nash Eq. – FNP (Roughly)

Let R be a binary relation on words (over some alphabet) that is polynomial-time computable and polynomially balanced.

I.e., membership to R is decidable in polynomial time, and $(x, y) \in R$ implies $|y| \leq |x|^k$ where k is independent of x, y .

A *search problem* associated with R is this: Given an input x , return a y such that $(x, y) \in R$ if such y exists, and return "NO" otherwise.

Note that the problem of computing NE can be seen as a search problem R where $(x, y) \in R$ means that x is a strategic-form game and y is a Nash equilibrium of polynomial size. (We already know from support enumeration that there is a NE of polynomial size.)

The class of all search problems is called FNP. A class $FP \subseteq FNP$ contains all search problem that can be solved in polynomial time.

A search problem determined by R is *polynomially reducible* to a search problem R' iff there exist polynomially computable functions f, g such that

- ▶ if $(x, y) \in R$ for some y , then $(f(x), y') \in R'$ for some y'
- ▶ if $(f(x), y) \in R'$, then $(x, g(y)) \in R$
- ▶ if $(f(x), y) \notin R'$ for all y , then $(x, y) \notin R$ for all y

Complexity of Nash Eq. – PPAD (Roughly)

The class PPAD is defined by specifying one of its complete problems (w.r.t. the polynomial time reduction) known as *End-Of-The-Line*:

- ▶ **Input:** Two *Boolean circuits (with basis \wedge, \vee, \neg)* S and P , each with m input bits and m output bits, such that $P(0^m) = 0^m \neq S(0^m)$.
- ▶ **Problem:** Find an input $x \in \{0, 1\}^m$ such that $P(S(x)) \neq x$ or $S(P(x)) \neq x \neq 0^m$.

Intuition: *End-Of-The-Line* creates a directed graph $H_{S,P}$ with vertex set $\{0, 1\}^m$ and an edge from x to y whenever both $y = S(x)$ ("successor") and $x = P(y)$ ("predecessor").

All vertices of $H_{S,P}$ have indegree and outdegree at most one. There is at least one source (i.e., x satisfying $P(x) = x$, namely 0^m), so there is at least one sink (i.e., x satisfying $S(x) = x$).

The goal is to find either a source or a sink different from 0^m .

Theorem 50

The problem of computing Nash equilibria is complete for PPAD.

That is, Nash belongs to PPAD and End-Of-The-Line is polynomially reducible to Nash.

Loose Ends – Modes of Dominance

Let $\sigma_i, \sigma'_i \in \Sigma_i$. Then σ'_i is *strictly dominated* by σ_i if $u_i(\sigma_i, \sigma_{-i}) > u_i(\sigma'_i, \sigma_{-i})$ for all $\sigma_{-i} \in \Sigma_{-i}$.

Let $\sigma_i, \sigma'_i \in \Sigma_i$. Then σ'_i is *weakly dominated* by σ_i if $u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i})$ for all $\sigma_{-i} \in \Sigma_{-i}$ and there is $\sigma'_{-i} \in \Sigma_{-i}$ such that $u_i(\sigma_i, \sigma'_{-i}) > u_i(\sigma'_i, \sigma'_{-i})$.

Let $\sigma_i, \sigma'_i \in \Sigma_i$. Then σ'_i is *very weakly dominated* by σ_i if $u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i})$ for all $\sigma_{-i} \in \Sigma_{-i}$.

A strategy is (strictly, weakly, very weakly) dominant in mixed strategies if it (strictly, weakly, very weakly) dominates any other mixed strategy.

Claim 4

Any mixed strategy profile $\sigma \in \Sigma$ such that each σ_i is very weakly dominant in mixed strategies is a mixed Nash equilibrium.

The same claim can be proved in pure strategy setting.

Dynamic Games of Complete Information
Extensive-Form Games
Definition
Sub-Game Perfect Equilibria

Dynamic Games of Perfect Information

(Motivation)

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 how to model situations in which 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 to express sequential nature of games.

We start with perfect information games, where each player always knows 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).

Perfect-Info. Extensive-Form Games (Example)



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.

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 left most terminal node gives 3 to player 1 and 1 to player 2.

Perfect-Information Extensive-Form Games

A *perfect-information extensive-form game* is a tuple

$G = (N, A, H, Z, \chi, \rho, \pi, h_0, u)$ 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* (choice) nodes, Z is a set of *terminal* nodes (assume $Z \cap H = \emptyset$), denote $\mathcal{H} = H \cup Z$,
- ▶ $\chi : H \rightarrow (2^A \setminus \{\emptyset\})$ is the *action function*, which assigns to each choice node a *non-empty* set of *enabled* actions,
- ▶ $\rho : H \rightarrow N$ is the *player function*, which assigns to each non-terminal node a player $i \in N$ who chooses an action there, we define $H_i := \{h \in H \mid \rho(h) = i\}$,
- ▶ $\pi : H \times A \rightarrow \mathcal{H}$ 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 (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 $a_1 = a_2$,
- ▶ $u = (u_1, \dots, u_n)$, where each $u_i : Z \rightarrow \mathbb{R}$ is a *payoff function* for player i in the terminal nodes of Z .

Some Notation

A *path* from $h \in \mathcal{H}$ to $h' \in \mathcal{H}$ is a sequence $h_1 a_2 h_2 a_3 h_3 \cdots h_{k-1} a_k h_k$ where $h_1 = h$, $h_k = h'$ and $\pi(h_{j-1}, a_j) = h_j$ for every $1 < j \leq k$.

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

Assumption: For every $h \in \mathcal{H}$ there is a unique path from h_0 to h and there is no infinite path (i.e., a sequence $h_1 a_2 h_2 a_3 h_3 \cdots$ such that $\pi(h_{j-1}, a_j) = h_j$ for every $j > 1$).

Note that the assumption is satisfied when \mathcal{H} is finite.

Indeed, uniqueness follows immediately from the definition of π . Now let X be the set of all h' from which there is a path to h . If $h_0 \in X$ we are done.

Otherwise, let h' be a node of X with the longest path to h . As $h' \neq h_0$, there is h'' and $a \in \chi(h'')$ such that $h' = \pi(h'', a)$. But then there is a path from h'' to h that is longer than the path from h' , a contradiction.

The above claim implies that every perfect-information extensive-form game can be seen as a game on a *rooted tree* (\mathcal{H}, E, h_0) where

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


Some More Notation

h' is a *child* of h , and h is a *parent* of h' if there is $a \in \chi(h)$ such that $h' = \pi(h, a)$.

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


If h' is reachable from h we say that h' is a descendant of h and h is an ancestor of h' (note that, by definition, h is both a descendant and an ancestor of itself).

Example: Trust Game



- ▶ Two players, both start with 5\$
- ▶ 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 money is tripled by the experimenter and sent to 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))

Example: Trust Game (Cont.)



- ▶ $N = \{1, 2\}$, $A = \{D, T, K, S\}$
- ▶ $H = \{h_0, h_1\}$, $Z = \{z_1, z_2, z_3\}$
- ▶ $\chi(h_0) = \{D, T\}$, $\chi(h_1) = \{K, S\}$
- ▶ $\rho(h_0) = 1$, $\rho(h_1) = 2$
- ▶ $\pi(h_0, D) = z_1$, $\pi(h_0, T) = h_1$, $\pi(h_1, K) = z_2$, $\pi(h_1, S) = z_3$
- ▶ $u_1(z_1) = 5$, $u_1(z_2) = 0$, $u_1(z_3) = 7.5$, $u_2(z_1) = 5$, $u_2(z_2) = 20$, $u_2(z_3) = 12.5$

Stackelberg Competition

Very similar to Cournot duopoly ...

- ▶ 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 .

Except that ...

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

Stackelberg Competition – Extensive-Form Model

An extensive-form game model:

- ▶ $N = \{1, 2\}$
- ▶ $A = [0, \infty)$
- ▶ $H = \{h_0, h_1^{q_1} \mid q_1 \in [0, \infty)\}$
- ▶ $Z = \{z^{q_1, q_2} \mid q_1, q_2 \in [0, \infty)$
- ▶ $\chi(h_0) = [0, \infty), \quad \chi(h_1^{q_1}) = [0, \infty)$
- ▶ $\rho(h_0) = 1, \quad \rho(h_1^{q_1}) = 2$
- ▶ $\pi(h_0, q_1) = h_1^{q_1}, \quad \pi(h_1^{q_1}, q_2) = z^{q_1, q_2}$
- ▶ The payoffs are
 - ▶ $u_1(z^{q_1, q_2}) = q_1(\kappa - q_1 - q_2) - q_1c$
 - ▶ $u_2(z^{q_1, q_2}) = q_2(\kappa - q_1 - q_2) - q_2c$

Example: Chess (a bit simplified)

There are infinitely many representations of chess, this one is different from the one presented at the lecture.

- ▶ $N = \{1, 2\}$
- ▶ Denoting *Boards* the set of all (appropriately encoded) board positions, we define $\mathcal{H} = B \times \{1, 2\}$ where
$$B = \{w \in \text{Boards}^+ \mid \text{no board repeats } \geq 3 \text{ times in } w\}$$
(Here *Boards*⁺ is the set of all non-empty sequences of boards)
- ▶ Z consists of all nodes (wb, i) (here $b \in \text{Boards}$) 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 two occurrences in w
- ▶ $\chi(wb, i)$ is the set of all legal moves of player i in b
- ▶ $\rho(wb, i) = i$
- ▶ π is defined by $\pi((wb, i), a) = (wbb', 2 - i + 1)$ where b' is obtained from b according to the move a
- ▶ $h_0 = (b_0, 1)$ where b_0 is the initial board
- ▶ $u_j(wb, i) \in \{1, 0, -1\}$, here 1 means "win", 0 means "draw", and -1 means "loss" for player j

Pure Strategies

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

Definition 51

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

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 $s \in S$ determines a unique path $w_s = h_0 a_1 h_1 \cdots h_{k-1} a_k h_k$ from h_0 to a terminal node h_k by

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

Denote by $O(s)$ the terminal node reached by w_s .

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

Example: Trust Game



A pure strategy profile (s_1, s_2) where

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

is usually written as TK (BFS & left to right traversal) determines the path $h_0 T h_1 K z_2$

The resulting payoffs: $u_1(s_1, s_2) = 0$ and $u_2(s_1, s_2) = 20$.

Extensive-Form vs Strategic-Form

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

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 \bar{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.

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 strategic-form area make sense in the extensive-form. Also, naive application of algorithms may result in unnecessarily high complexity.

For now, let us consider pure strategies only!

Example: Trust Game




Is any strategy strictly (weakly, very weakly) dominant?

Is any strategy never best response?

Is there a Nash equilibrium in pure strategies ?

Example



Find all pure strategies of both players.


Is any strategy (strictly, weakly, very weakly) dominant?

Is any strategy (strictly, weakly, very weakly) dominated?

Is any strategy never best response?

Are there Nash equilibria in pure strategies ?

Example



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

Find all pure strategies of both players.


Is any strategy (strictly, weakly, very weakly) dominant?

Is any strategy (strictly, weakly, very weakly) dominated?

Is any strategy never best response?

Are there Nash equilibria in pure strategies ?

Criticism of Nash Equilibria



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


Two Nash equilibria in pure strategies: (L, UU') and (R, UK')

Examine (L, UU') :

- ▶ 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' .

Criticism of Nash Equilibria



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

Two Nash equilibria in pure strategies: (L, UU') and (R, UK')

Examine (R, UK') : This equilibrium 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*.

Subgame Perfect Equilibria

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

Definition 52 (Subgame)

A *subgame* G^h of G rooted in $h \in \mathcal{H}$ is the restriction of G to nodes reachable from h in the game tree. More precisely,

$G^h = (N, A, H^h, Z^h, \chi^h, \rho^h, \pi^h, h, u^h)$ where $H^h = H \cap \mathcal{H}^h$, $Z^h = Z \cap \mathcal{H}^h$, χ^h and ρ^h are restrictions of χ and ρ to H^h , resp., (Given a function $f : A \rightarrow B$ and $C \subseteq A$, a restriction of f to C is a function $g : C \rightarrow B$ such that $g(x) = f(x)$ for all $x \in C$.)

- ▶ 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

Definition 53

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

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

Stackelberg Competition – SPE

- ▶ $N = \{1, 2\}$, $A = [0, \infty)$
- ▶ $H = \{h_0, h_1^{q_1} \mid q_1 \in [0, \infty)\}$, $Z = \{z^{q_1, q_2} \mid q_1, q_2 \in [0, \infty)$
- ▶ $\chi(h_0) = [0, \infty)$, $\chi(h_1^{q_1}) = [0, \infty)$, $\rho(h_0) = 1$, $\rho(h_1^{q_1}) = 2$
- ▶ $\pi(h_0, q_1) = h_1^{q_1}$, $\pi(h_1^{q_1}, q_2) = z^{q_1, q_2}$
- ▶ The payoffs are $u_1(z^{q_1, q_2}) = q_1(\kappa - c - q_1 - q_2)$,
 $u_2(z^{q_1, q_2}) = q_2(\kappa - c - q_1 - q_2)$

Denote $\theta = \kappa - c$

Player 1 chooses q_1 , 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 $q_2 = \theta/4$.

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

Note that firm 1 has an advantage as a leader.

Existence of SPE

From this moment on we consider only **finite games!**

Theorem 54

Every finite perfect-information extensive-form game has a SPE in pure strategies.

Proof: By induction on the number of nodes.

Base case: If $|\mathcal{H}| = 1$, the only node is terminal, and the trivial pure strategy profile is SPE.

Induction step: Consider a game with more than one node. Let $K = \{h_1, \dots, h_k\}$ be the set of all children of the root h_0 .

By induction, for every h_ℓ there is a SPE s^{h_ℓ} in G^{h_ℓ} .

For every $i \in N$, define a strategy s_i of player i in G as follows:

- ▶ for $i = \rho(h_0)$ we have $s_i(h_0) \in \operatorname{argmax}_{h_\ell \in K} u_i^{h_\ell}(s^{h_\ell})$
- ▶ for all $i \in N$ and $h \in H$ we have $s_i(h) = s_i^{h_\ell}(h)$ where $h \in H^{h_\ell} \cap H_i$

We claim that $s = (s_1, \dots, s_n)$ is a SPE in pure strategies.

By definition, s is NE in all subgames except (possibly) the G itself.



Existence of SPE (Cont.)

Let $h_\ell = s_{\rho(h_0)}(h_0)$.

Consider a possible deviation of player i .

Let \bar{s} be another pure strategy profile in G obtained from $s = (s_1, \dots, s_n)$ by changing s_i .

First, assume that $i \neq \rho(h_0)$. Then

$$u_i(s) = u_i^{h_\ell}(s^{h_\ell}) \geq u_i^{h_\ell}(\bar{s}^{h_\ell}) = u_i(\bar{s})$$

Here the first equality follows from $h_\ell = s_{\rho(h_0)}(h_0)$ and that s behaves similarly as s^{h_ℓ} in G^{h_ℓ} , the inequality follows from the fact that s^{h_ℓ} is a NE in G^{h_ℓ} , and the second equality follows from $h_\ell = s_{\rho(h_0)}(h_0) = \bar{s}_{\rho(h_0)}(h_0)$.

Second, assume that $i = \rho(h_0)$.

Let $h_r = \bar{s}_i(h_0) = \bar{s}_{\rho(h_0)}(h_0)$.

Then $u_i^{h_\ell}(s^{h_\ell}) \geq u_i^{h_r}(s^{h_r})$ because h_ℓ maximizes the payoff of player $i = \rho(h_0)$ in the children of h_0 .

But then

$$u_i(s) = u_i^{h_\ell}(s^{h_\ell}) \geq u_i^{h_r}(s^{h_r}) \geq u_i^{h_r}(\bar{s}^{h_r}) = u_i(\bar{s})$$

Recall that in the model of chess, the payoffs were from $\{1, 0, -1\}$ and $u_1 = -u_2$ (i.e. it is zero-sum).

By Theorem 54, there is a SPE in pure strategies (s_1^*, s_2^*) .

However, then one of the following holds:

1. White has a winning strategy

If $u_1(s_1^*, s_2^*) = 1$ and thus $u_2(s_1^*, s_2^*) = -1$

2. Black has a winning strategy

If $u_1(s_1^*, s_2^*) = -1$ and thus $u_2(s_1^*, s_2^*) = 1$

3. Both players have strategies to force a draw

If $u_1(s_1^*, s_2^*) = 0$ and thus $u_2(s_1^*, s_2^*) = 0$

Question: Which one is the right answer?

Answer: Nobody knows yet ... the tree is too big!

Even with ~ 200 depth & ~ 5 moves per node: 5^{200} nodes!

Backward Induction

The proof of Theorem 54 gives an efficient procedure for computing SPE for finite perfect-information extensive-form games.

Backward Induction: We inductively "attach" to every node h a SPE s^h in G^h , together with a vector of expected payoffs

$$u(h) = (u_1(h), \dots, u_n(h)).$$

- ▶ **Initially:** Attach to each terminal node $z \in Z$ the empty profile $s^z = (\emptyset, \dots, \emptyset)$ and the payoff vector $u(z) = (u_1(z), \dots, u_n(z))$.
- ▶ **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 SPE s^h where

- ▶ $s_{\rho(h)}^h(h) = h_{\max}$

- ▶ for all $i \in N$ and all $h' \in H_i$ define $s_i^h(h') = \bar{s}_i^h(h')$ where

$$h' \in H^{\bar{h}} \cap H_i \quad (\text{in } G^{\bar{h}}, \text{ each } s_i^h \text{ behaves as } \bar{s}_i^h \text{ i.e. } (s^h)^{\bar{h}} = \bar{s}^{\bar{h}})$$

4. Attach to h the expected payoffs $u_i(h) = u_i(h_{\max})$ for $i \in N$.

Efficient Algorithms for Pure Nash Equilibria

In the step 2. of the backward induction, the algorithm may choose *an arbitrary* $h_{\max} \in \operatorname{argmax}_{h' \in K} u_{\rho(h)}(h')$ and always obtain a SPE.

In order to compute all SPE, the algorithm may systematically search through all possible choices of h_{\max} throughout the induction.

Backward induction is too inefficient (unnecessarily searches through the whole tree).


There are better algorithms, such as α - β -pruning.

For details, extensions etc. see e.g.

- ▶ PB016 Artificial Intelligence I
- ▶ Multi-player alpha-beta pruning, R. Korf, *Artificial Intelligence* 48, pages 99-111, 1991
- ▶ Artificial Intelligence: A Modern Approach (3rd edition), S. Russell and P. Norvig, *Prentice Hall*, 2009

Example

Centipede game:



SPE in pure strategies: (DDD, DD) ... Isn't it weird?

There are serious issues here ...


- ▶ In 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 :-)

Dynamic Games of Complete Information
Extensive-Form Games
Imperfect-Information Games

Extensive-form of Matching Pennies

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

	<i>H</i>	<i>T</i>
<i>H</i>	1, -1	-1, 1
<i>T</i>	-1, 1	1, -1




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.

Reversing the order of players does not help.

We need to extend the formalism to be able to hide some information about previous moves.

Extensive-form of Matching Pennies

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



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.

Imperfect Information Games

An *imperfect-information extensive-form game* is a tuple

$G_{imp} = (G_{perf}, I)$ where

- ▶ $G_{perf} = (N, A, H, Z, \chi, \rho, \pi, h_0, u)$ is a perfect-information extensive-form game (called *the underlying game*),
- ▶ $I = (I_1, \dots, I_n)$ 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)
- ▶ 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 .

Given an information set $I_{i,j}$, we denote by $\chi(I_{i,j})$ the set of all actions enabled in some (and hence all) nodes of $I_{i,j}$.

Imperfect Information Games – Strategies

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

Definition 55


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

Note that each s_i can also be seen as a function $s_i : I_i \rightarrow A$ 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 $S = S_1 \times \dots \times S_n$ the set of all pure strategy profiles.

As in the perfect-information case we have a corresponding strategic-form game $\bar{G}_{imp} = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$.

Matching Pennies



$I_1 = \{I_{1,1}\}$ where $I_{1,1} = \{h_0\}$


$I_2 = \{I_{2,1}\}$ where $I_{2,1} = \{h_1, h_2\}$

Example of pure strategies:

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

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

Weird Example




Note that $I_1 = \{I_{1,1}\}$ where $I_{1,1} = \{h_0, h_3\}$

and that $I_2 = \{I_{2,1}\}$ where $I_{2,1} = \{h_1, h_2\}$

What pure strategies are in this example?

SPE with Imperfect Information



What we designate as subgames to allow the backward induction?

Only subtrees rooted in h_1 , h_2 , and h_0 (together with all subtrees rooted in terminal nodes)

Note that subtrees rooted in h_3 and h_4 cannot be considered as "independent" subgames because their individual solutions cannot be combined to a single best response in the information set $\{h_3, h_4\}$.

SPE with Imperfect Information

Let $G_{imp} = (G_{perf}, I)$ be an imperfect-information extensive-form game where $G_{perf} = (N, A, H, Z, \chi, \rho, \pi, h_0, u)$ 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 , we have that either all nodes of $I(h')$ are reachable from h , or no node of $I(h')$ is reachable from h .

Intuitively, $h \in H_{proper}$ 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 56

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

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 57

A strategy profile $s \in S$ is a subgame perfect equilibrium (SPE) if s^h is a Nash equilibrium in every subgame G_{imp}^h of G_{imp} (here $h \in H_{proper}$).

Backward Induction with Imperfect Info

The backward induction generalizes to imperfect-information extensive-form games along the following lines:

1. As in the perfect-information case, the goal is to label each node $h \in H_{proper} \cup Z$ with a SPE s^h and a vector of payoffs $u(h) = (u_1(h), \dots, u_n(h))$ for individual players according to s^h .
2. Starting with terminal nodes, the labeling proceeds bottom up. Terminal nodes are labeled similarly as in the perfect-inf. case.
3. Consider $h \in H_{proper}$, let K be the set of all $h' \in (H_{proper} \cup Z) \setminus \{h\}$ that are h 's **closest descendants out of $H_{proper} \cup Z$** .

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 $\mathcal{H} \setminus H_{proper}$ (except the first and the last node).

For every $h' \in K$ we have already computed a SPE $s^{h'}$ in $G_{imp}^{h'}$ and the vector of corresponding payoffs $u(h')$.

4. Now consider all nodes of K as terminal nodes where each $h' \in K$ has payoffs $u(h')$. This gives a new game in which we compute an equilibrium \bar{s}^h together with the vector $u(h)$.

The equilibrium s^h is then obtained by "concatenating" \bar{s}^h with all $s^{h'}$, here $h' \in K$, in the subgames $G_{imp}^{h'}$ of G_{imp}^h .

Mutually Assured Destruction

Analysis of Cuban missile crisis of 1962
(as described in *Games for Business and Economics* by R. Gardner)

- ▶ The crisis started with 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".

Question: Could this indeed be a credible threat?


Mutually Assured Destruction (Cont.)

Model as an extensive-form game:

- ▶ First, player 1 (US) chooses to either ignore the incident (I), resulting in maintenance of status quo (payoffs $(0, 0)$), or escalate the situation (E).
- ▶ Following escalation by player 1, player 2 can back down (B), causing it to lose face (payoffs $(10, -10)$), or it can choose to proceed to a nuclear confrontation (N).
- ▶ Upon this choice, the players play a simultaneous-move game in which they can either retreat (R), or choose doomsday (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 $(-100, -100)$.

Find SPE in pure strategies.

Mutually Assured Destruction (Cont.)



Solve $G_{imp}^{h_2}$ (a strategic-form game). Then $G_{imp}^{h_1}$ by solving a game rooted in h_1 with terminal nodes h_2, z_5 (payoffs in h_2 correspond to an equilibrium in $G_{imp}^{h_2}$). Finally solve G_{imp} by solving a game rooted in h_0 with terminal nodes h_1, z_6 (payoffs in h_1 have been computed in the previous step).

Dynamic Games of Complete Information
Repeated Games
Finitely Repeated Games

Example

	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 what follows 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 payoffs and the long-run average payoff

We analyze Nash equilibria and sub-game perfect equilibria.

We stick to pure strategies only!

Finitely Repeated Games

Let $G = (\{1, 2\}, (S_1, S_2), (u_1, u_2))$ be a finite strategic-form game of two players.

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 $s^t = (s_1^t, s_2^t)$.

After T stages, both players collect the average payoff $\sum_{t=1}^T u_i(s^t) / T$.

A *history of length* $0 \leq t \leq T$ is a sequence $h = s^1 \cdots s^t \in S^t$ of t strategy profiles. Denote by $H(t)$ the set of all histories of length t .

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

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

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

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

Given a pure strategy profile τ in $G_{T\text{-rep}}$ such that $w_\tau = s^1 \cdots s^T$, define the payoffs $u_i(\tau) = \sum_{t=1}^T u_i(s^t) / T$.

Example

	C	S
C	-5, -5	0, -20
S	-20, 0	-1, -1

Consider a 3-stage game.

Examples of histories: ϵ , (C, S) , $(C, S)(S, S)$, $(C, S)(S, S)(C, C)$

Here the last one is terminal, obtained using τ_1, τ_2 s.t.:

$$\tau_1(\epsilon) = C, \tau_1((C, S)) = S, \tau_1((C, S)(S, S)) = C$$

$$\tau_2(\epsilon) = S, \tau_2((C, S)) = S, \tau_2((C, S)(S, S)) = C$$

Thus $w_{(\tau_1, \tau_2)} = (C, S)(S, S)(C, C)$

$$u_1(\tau_1, \tau_2) = (0 + (-1) + (-5))/3 = -2$$

$$u_2(\tau_1, \tau_2) = (-20 + (-1) + (-5))/3 = -26/3$$

Finitely Repeated Games in Extensive-Form

Every T -stage game $G_{T\text{-rep}}$ can be defined as an imperfect information extensive-form game.

Define an imperfect-information extensive-form game $G_{\text{imp}}^{\text{rep}} = (G_{\text{perf}}^{\text{rep}}, I)$ such that $G_{\text{perf}}^{\text{rep}} = (\{1, 2\}, A, H, Z, \chi, \rho, \pi, h_0, u)$ where

- ▶ $A = S_1 \cup S_2$
- ▶ $H = (S_1 \times S_2)^{\leq T} \cup (S_1 \times S_2)^{< T} \cdot S_1$
Intuitively, elements of $(S_1 \times S_2)^{\leq k}$ are possible histories;
 $(S_1 \times S_2)^{< k} \cdot S_1$ is used to simulate a simultaneous play of G by 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 $u_i((s_1^1, s_2^1)(s_1^2, s_2^2) \cdots (s_1^T, s_2^T)) = \sum_{t=1}^T u_i(s_1^t, s_2^t) / T$

Finitely Repeated Games in Extensive-Form

The set 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 $s_1 \in S_1$.

Intuitively, in every round, player 1 has a complete information about results of past plays,

player 1 chooses a pure strategy $s_1 \in S_1$,

player 2 is *not* informed about s_1 but still has a complete information about results of all previous rounds,

player 2 chooses a pure strategy $s_2 \in S_2$ and both players are informed about the result.

Finitely Repeated Games – Equilibria

Definition 58

A strategy profile $\tau = (\tau_1, \tau_2)$ 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

$$u_i(\tau_1, \tau_2) \geq u_i(\tau'_i, \tau_{-i})$$

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

$$\tau_i^h(\bar{s}^1 \cdots \bar{s}^t) = \tau_i(s^1 \cdots s^t \bar{s}^1 \cdots \bar{s}^t) \quad \text{for every sequence } \bar{s}^1 \cdots \bar{s}^t$$

(i.e. τ_i^h behaves as τ_i after h)

Definition 59

A strategy profile $\tau = (\tau_1, \tau_2)$ in a T -stage game $G_{T\text{-rep}}$ is a subgame-perfect Nash equilibrium (SPE) if 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 .

SPE with Single NE in G

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

Consider a T -stage game based on Prisoner's dilemma.

For every T , find a SPE.

... there is one, play (C, C) all the time. Is it all?

Theorem 60

Let G be an arbitrary finite strategic-form game. If G has a unique Nash equilibrium, then playing this equilibrium all the time is the unique SPE in the T -stage game based on G .

Proof.

By backward induction, players have to play the NE 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. \square

Further Discussion of Prisoner's Dilemma

	C	S
C	-5, -5	0, -20
S	-20, 0	-1, -1

Are there other NE (that are not SPE) in the repeated Prisoner's dilemma?

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

There are 4 NE: They are the four profiles that lead to $(C, C)(C, C)$, i.e., each player plays either C - CC , or C - CS .

Further Discussion of Prisoner's Dilemma

	C	S
C	-5, -5	0, -20
S	-20, 0	-1, -1

The strategy C strictly dominates S in the Prisoner's dilemma.

Is there a strictly dominant strategy in the 2-stage game based on the Prisoner's dilemma?

If player 2 plays $S-CC$, then the best responses of player 1 are $S-CC$ and $S-SC$.

(The strategy $S-CS$ is usually called "tit-for-tat".)

If player 2 plays $S-SC$, 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.)

SPE with Multiple NE in G

Let $s = (s_1, s_2)$ be a Nash equilibrium in G .

Define a strategy profile $\tau = (\tau_1, \tau_2)$ in $G_{T\text{-rep}}$ where

- ▶ τ_1 chooses s_1 in every stage
- ▶ τ_2 chooses s_2 in every stage

Proposition 3

τ is a SPE 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 proposition may be generalized by allowing players to play different equilibria in particular stages

I.e., consider a sequence of NE s^1, s^2, \dots, s^T in G and assume that in stage ℓ player i plays s_i^ℓ

Does this cover all possible SPE in finitely repeated games?

SPE with Multiple NE in 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

NE in the above game G : (F, f) and (R, r)

Consider 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 1. In stage 2 plays r if (M, m) was played in the first stage, and plays f otherwise.

Is this SPE?

Note that here the players **do not** play a NE in the first step.

The idea is that both players agree to play a Pareto optimal profile. If both comply, then a favorable NE is played in the second stage. If one of them betrays then a "punishing" NE is played.

Dynamic Games of Complete Information
Repeated Games
Infinitely Repeated Games

Infinitely Repeated Games

Let $G = (\{1, 2\}, (S_1, S_2), (u_1, u_2))$ be a 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 $s^t = (s_1^t, s_2^t)$.

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

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

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

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

Every pure strategy profile $\tau = (\tau_1, \tau_2)$ in G_{irep} induces a sequence of pure strategy profiles $w_\tau = s^1 s^2 \cdots$ in G so that $s_i^t = \tau_i(s^1 \cdots s^{t-1})$.

(Here for $t = 0$ we have that $s^1 \cdots s^{t-1} = \epsilon$.)

Infinitely Repeated Games & Discounted Payoff

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

Given $0 < \delta < 1$, we define a *δ -discounted payoff* by

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

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.

Infinitely Repeated Games & Discounted Payoff

Definition 61

A strategy profile $\tau = (\tau_1, \tau_2)$ is a Nash equilibrium in G_{irep}^δ if for both $i \in \{1, 2\}$ and for every τ'_i we have that

$$u_i^\delta(\tau_i, \tau_{-i}) \geq u_i^\delta(\tau'_i, \tau_{-i})$$

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

$$\tau_i^h(\bar{s}^1 \dots \bar{s}^t) = \tau_i(s^1 \dots s^t \bar{s}^1 \dots \bar{s}^t) \quad \text{for every sequence } \bar{s}^1 \dots \bar{s}^t$$

(i.e. τ_i^h behaves as τ_i after h)

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

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

Example

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

	C	S
C	-5, -5	0, -20
S	-20, 0	-1, -1

What are the Nash equilibria and SPE in G_{irep}^δ for a given δ ?

Consider a pure strategy profile (τ_1, τ_2) where $\tau_i(s^1 \dots s^T) = C$ for all $T \geq 1$ and $i \in \{1, 2\}$. Is it a NE? A SPE?

Consider a "grim trigger" profile (τ_1, τ_2) where

$$\tau_i(s^1 \dots s^T) = \begin{cases} S & T = 0 \\ S & s^\ell = (S, S) \text{ for all } 1 \leq \ell \leq T \\ C & \text{otherwise} \end{cases}$$

Is it a NE? Is it a SPE?

One-Shot Deviation Principle

A pure strategy profile $\tau = (\tau_1, \tau_2)$ in G_{irep} satisfies **one-shot deviation property in G_{irep}^δ** if for every $i \in \{1, 2\}$ and every $\bar{\tau}_i$, differing from τ_i just on a single history h , we have $u_i^\delta(\bar{\tau}_1^h, \tau_2^h) \leq u_i^\delta(\tau_1^h, \tau_2^h)$.

Theorem 62

Let $G = (\{1, 2\}, (S_1, S_2), (u_1, u_2))$ be a two-player strategic-form game such that both u_1 and u_2 are bounded on $S = S_1 \times S_2$. Let $0 < \delta < 1$. A pure strategy profile $\tau = (\tau_1, \tau_2)$ in G_{irep} is a SPE in G_{irep}^δ **iff** it satisfies the one-shot deviation property in G_{irep}^δ .

Before proving Theorem 62, let us note the following:

- ▶ The one shot deviation property is concerned with all strategies $\bar{\tau}_i$ that differ from τ_i on a single history. This means that we have to consider **all** histories h , even those that *can not* be visited using τ_i with any opponent.
- ▶ The one-shot deviation property immediately implies the following: If $\bar{\tau}_i$ *does not* differ from τ_i on any history of the form $h' = hh''$ where $h'' \neq \varepsilon$ (i.e., on any history obtained by prolonging h), then $u_i^\delta(\bar{\tau}_1^h, \tau_2^h) \leq u_i^\delta(\tau_1^h, \tau_2^h)$.

Indeed, note that τ_i^h differs from $\bar{\tau}_i^h$ only on h .

One-Shot Deviation Principle

Proof. \Rightarrow : Trivial.

\Leftarrow : Assume that τ satisfies the one-shot deviation property but is not a SPE. That is, a deviation may increase payoff of one of the players in a subgame. Assume, w.l.o.g., that player 1 gains by deviation to a strategy $\bar{\tau}_1$ in a subgame starting with a h , i.e.,

$$u_1^\delta(\bar{\tau}_1^h, \tau_2^h) > u_1^\delta(\tau_1^h, \tau_2^h) \quad (2)$$

Since $\delta < 1$ and u_i are bounded on S , we may safely choose $\bar{\tau}_1$ so that $\bar{\tau}_1(h') = \tau_1(h')$ for all sufficiently long histories h' .

Indeed, since u_i is bounded on pure strategies of G , the sum $\sum_{t=\ell}^{\infty} \delta^t \cdot u_i(s^{t+1})$ goes to 0 as ℓ goes to ∞ ; hence the strict inequality (2) remains valid even if $\bar{\tau}_1$ is arbitrarily modified in a very distant future.

One-Shot Deviation Principle

Let h' be a history of *maximum length* such that h is a prefix of h' and $\bar{\tau}_1(h') \neq \tau_1(h')$. (Note that then $\bar{\tau}_1(h'h'') = \tau_1(h'h'')$ for all $h'' \neq \varepsilon$.)

Let $\bar{\tau}_{11}$ be a strategy of player 1 obtained from $\bar{\tau}_1$ by changing $\bar{\tau}_1(h')$ to $\tau_1(h')$. Now note that the **one-shot deviation property** implies, that

$$u_1^\delta(\bar{\tau}_{11}^{h'}, \tau_2^{h'}) = u_1^\delta(\tau_1^{h'}, \tau_2^{h'}) \geq u_1^\delta(\bar{\tau}_1^{h'}, \tau_2^{h'})$$

and thus $u_1^\delta(\bar{\tau}_{11}^h, \tau_2^h) \geq u_1^\delta(\bar{\tau}_1^h, \tau_2^h) > u_1^\delta(\tau_1^h, \tau_2^h)$. Note that $\bar{\tau}_{11}^h$ has a strictly smaller number of deviations from τ_1^h than $\bar{\tau}_1^h$.

Repeating the same argument with $\bar{\tau}_{11}$ in place of $\bar{\tau}_1$ we obtain $\bar{\tau}_{12}$ such that $u_1^\delta(\bar{\tau}_{12}^h, \tau_2^h) \geq u_1^\delta(\bar{\tau}_{11}^h, \tau_2^h) > u_1^\delta(\tau_1^h, \tau_2^h)$. Here $\bar{\tau}_{12}^h$ has even less deviations from τ_1^h than $\bar{\tau}_{11}^h$.

Then repeating with $\bar{\tau}_{12}$ in place of $\bar{\tau}_1$ we obtain $\bar{\tau}_{13}$ such that $u_1^\delta(\bar{\tau}_{13}^h, \tau_2^h) \geq u_1^\delta(\bar{\tau}_{12}^h, \tau_2^h) > u_1^\delta(\tau_1^h, \tau_2^h)$, etc., still decreasing the number of deviations from τ_1^h .

Eventually, as $\bar{\tau}_1^h$ has only finitely many deviations from τ_1^h , we get $\bar{\tau}_{1k}^h = \tau_1^h$ for some k and thus $u_1^\delta(\tau_1^h, \tau_2^h) = u_1^\delta(\bar{\tau}_{1k}^h, \tau_2^h) > u_1^\delta(\tau_1^h, \tau_2^h)$, a contradiction. □

Example

Consider the infinitely repeated game based on Prisoner's dilemma:

	C	S
C	-5, -5	0, -20
S	-20, 0	-1, -1

The grim trigger profile (τ_1, τ_2) where

$$\tau_i(s^1 \cdots s^T) = \begin{cases} S & T = 0 \\ S & s^\ell = (S, S) \text{ for all } 1 \leq \ell \leq T \\ C & \text{otherwise} \end{cases}$$

is a SPE.

A Simple Version of Folk Theorem

Let $G = (\{1, 2\}, (S_1, S_2), (u_1, u_2))$ be a 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 s^* be a Nash equilibrium in G .

Let s be a strategy profile in G satisfying $u_i(s) > u_i(s^*)$ for all $i \in N$.

Consider the following *grim trigger for s using s^** strategy profile $\tau = (\tau_1, \tau_2)$ in G_{irep} where

$$\tau_i(s^1 \cdots s^T) = \begin{cases} s_i & T = 0 \\ s_i & s^\ell = s \text{ for all } 1 \leq \ell \leq T \\ s_i^* & \text{otherwise} \end{cases}$$

Then for

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

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

Proof: Consider a possible one-shot deviation $\bar{\tau}_1$ of player 1, i.e., there is exactly one h such that $\bar{\tau}_1(h) \neq \tau_1(h)$. We distinguish two cases depending on h .

Proof of Simple Folk Theorem (Cont.)

Case 1: $h \neq s \cdots s$. Then there is a deviation from s in h and thus according to (τ_1^h, τ_2^h) both players play s^* forever :

$$u_1^\delta(\tau_1^h, \tau_2^h) = (1 - \delta) \sum_{k=0}^{\infty} \delta^k u_1(s^*) = u_1(s^*) (1 - \delta) \sum_{k=0}^{\infty} \delta^k = u_1(s^*)$$

Now $(\bar{\tau}_1^h, \tau_2^h)$ gives a sequence $w_{(\bar{\tau}_1^h, \tau_2^h)} = (s'_1, s_2^*) s^* s^* \cdots$ where s'_1 is a strategy of player 1 to which he deviates after h .

Here player 2 plays s_2^* all the time after h because one of the players has already deviated in h .

We obtain

$$\begin{aligned} u_1(\bar{\tau}_1^h, \tau_2^h) &= (1 - \delta) \left(u_1(s'_1, s_2^*) + \sum_{k=1}^{\infty} \delta^k u_1(s^*) \right) \\ &\leq (1 - \delta) \left(u_1(s_1^*, s_2^*) + \sum_{k=1}^{\infty} \delta^k u_1(s^*) \right) \\ &= u_1(s^*) \end{aligned}$$

So this deviation cannot be beneficial no matter what δ is.

Proof of Simple Folk Theorem (Cont.)

Case 2: $h = s \cdots s$. Clearly, $u_1(\tau_1^h, \tau_2^h) = u_1(s)$.

Now $(\bar{\tau}_1^h, \tau_2^h)$ gives a sequence $w_{(\bar{\tau}_1^h, \tau_2^h)} = (s'_1, s_2)s^*s^* \cdots$ where s'_1 is a strategy of player 1 to which he deviates after h .

As opposed to the previous case, here player 2 first plays s_2 (since the deviation of player 1 to s'_1 is the first deviation in the history) and then both players react by playing s^* forever.

If $u_1(s'_1, s_2) < u_1(s)$, then

$$\begin{aligned} u_1^\delta(\bar{\tau}_1^h, \tau_2^h) &= (1 - \delta) \left(u_1(s'_1, s_2) + \sum_{k=1}^{\infty} \delta^k u_1(s^*) \right) \\ &< (1 - \delta) \left(u_1(s_1, s_2) + \sum_{k=1}^{\infty} \delta^k u_1(s^*) \right) \\ &< (1 - \delta) \left(u_1(s) + \sum_{k=1}^{\infty} \delta^k u_1(s) \right) = u_1(s) = u_1^\delta(\tau_1^h, \tau_2^h) \end{aligned}$$

and thus this deviation is also not beneficial no matter what δ is.

Proof of Simple Folk Theorem (Cont.)

Finally, if $u_1(s'_1, s_2) \geq u_1(s)$, then

$$\begin{aligned}u_1^\delta(\bar{\tau}_1^h, \tau_2^h) &= (1 - \delta) \left(u_1(s'_1, s_2) + \sum_{k=1}^{\infty} \delta^k u_1(s^*) \right) \\&= (1 - \delta) u_1(s'_1, s_2) + (1 - \delta) u_1(s^*) \cdot \delta \sum_{k=0}^{\infty} \delta^k \\&= u_1(s'_1, s_2) - \delta \cdot u_1(s'_1, s_2) + \delta \cdot u_1(s^*)\end{aligned}$$

Thus

$$u_1^\delta(\bar{\tau}_1^h, \tau_2^h) \leq u_1^\delta(\tau_1^h, \tau_2^h) = u_1(s) \text{ iff}$$

$$u_1(s'_1, s_2) - \delta \cdot u_1(s'_1, s_2) + \delta \cdot u_1(s^*) \leq u_1(s) \text{ iff}$$

$$u_1(s'_1, s_2) - u_1(s) \leq \delta \cdot (u_1(s'_1, s_2) - u_1(s^*)) \text{ iff}$$

$$\delta \geq \frac{u_1(s'_1, s_2) - u_1(s)}{u_1(s'_1, s_2) - u_1(s^*)}$$

Proof of Simple Folk Theorem (Cont.)

Thus (τ_1, τ_2) satisfies the one-shot deviation property in G_{irep}^δ w.r.t. player 1 if

$$\delta \geq \frac{u_1(s'_1, s_2) - u_1(s)}{u_1(s'_1, s_2) - u_1(s^*)} \text{ for all } s'_1 \in S_1 \text{ satisfying } u_1(s'_1, s_2) \geq u_1(s)$$

Note that the right-hand-side expression is maximized when $u_1(s'_1, s_2)$ is maximized and thus we get

$$\delta \geq \frac{\max_{s'_1 \in S_1} u_1(s'_1, s_2) - u_1(s)}{\max_{s'_1 \in S_1} u_1(s'_1, s_2) - u_1(s^*)}$$

Proving the same for player 2 and putting the results together, we obtain that (τ_1, τ_2) satisfies the one-shot deviation property in G_{irep}^δ if

$$\delta \geq \max_{i \in \{1, 2\}} \frac{\max_{s'_i \in S_i} u_i(s'_i, s_{-i}) - u_i(s)}{\max_{s'_i \in S_i} u_i(s'_i, s_{-i}) - u_i(s^*)} \quad (3)$$

Thus by Theorem 62, (τ_1, τ_2) is a SPE in G_{irep}^δ if δ satisfies ineq. (3).

Simple Folk Theorem – Example

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

	m	f	r
M	4,4	-1,5	3,0
F	5,-1	1,1	0,0
R	0,3	0,0	2,2

NE in G : (F, f)

Consider the grim trigger for (M, m) using (F, f) , i.e., the profile (τ_1, τ_2) in G_{irep} where

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

This is a SPE in G_{irep}^δ for all $\delta \geq \frac{1}{4}$. Also, $u_i(\tau_1, \tau_2) = 4$ for $i \in \{1, 2\}$.

Are there other SPE? Yes, a grim trigger for (R, r) using (F, f) . This is a SPE in G_{irep}^δ for $\delta \geq \frac{1}{2}$.

Tacit Collusion

Consider the Cournot duopoly game model $G = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$

- ▶ $N = \{1, 2\}$
- ▶ $S_i = [0, \kappa]$
- ▶ $u_1(q_1, q_2) = q_1(\kappa - q_1 - q_2) - q_1 c_1 = (\kappa - c_1)q_1 - q_1^2 - q_1 q_2$
 $u_2(q_1, q_2) = q_2(\kappa - q_2 - q_1) - q_2 c_2 = (\kappa - c_2)q_2 - q_2^2 - q_2 q_1$

Assume for simplicity 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 NE $(\theta/3, \theta/3)$.

However, such contracts are forbidden in many countries (including US).

Is it 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 $\theta^2/8$?

Tacit Collusion

Consider the Cournot duopoly game model $G = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$

- ▶ $N = \{1, 2\}$
- ▶ $S_i = [0, \infty)$
- ▶ $u_1(q_1, q_2) = q_1(\kappa - q_1 - q_2) - q_1 c_1 = (\kappa - c_1)q_1 - q_1^2 - q_1 q_2$
 $u_2(q_1, q_2) = q_2(\kappa - q_2 - q_1) - q_2 c_2 = (\kappa - c_2)q_2 - q_2^2 - q_2 q_1$

Assume for simplicity that $c_1 = c_2 = c$ and denote $\theta = \kappa - c$.

Consider the grim trigger profile for $(\theta/4, \theta/4)$ using $(\theta/3, \theta/3)$:
Player i will

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

Assuming that $\kappa = 100$ and $c = 10$ (which gives $\theta = 90$), this is a SPE G_{irep}^δ for $\delta \geq 0.5294 \dots$. It results in $(\theta/4, \theta/4)(\theta/4, \theta/4) \dots$ with the discounted payoffs $\theta^2/8$.

Dynamic Games of Complete Information
Repeated Games
Infinitely Repeated Games
Long-Run Average Payoff and Folk Theorems

Infinitely Repeated Games & Average Payoff

In what follows we assume that all payoffs in the game G are positive and that S is finite!

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

Definition 63

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

$$u_i^{avg}(\tau) = \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T u_i(s^t)$$

(Here \limsup is necessary because τ_i may cause non-existence of the limit.)

The long-run average payoff $u_i^{avg}(\tau)$ is *well-defined* if the limit

$$u_i^{avg}(\tau) = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T u_i(s^t) \text{ exists.}$$

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

Infinitely Repeated Games & Average Payoff

Definition 64

A strategy profile τ is a Nash equilibrium if $u_i^{avg}(\tau)$ is well-defined for all $i \in N$, and for every i and every τ'_i we have that

$$u_i^{avg}(\tau_i, \tau_{-i}) \geq u_i^{avg}(\tau'_i, \tau_{-i})$$

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

Moreover, $\tau = (\tau_1, \tau_2)$ is a SPE in G_{irep}^{avg} if for every history h we have that (τ_1^h, τ_2^h) is a Nash equilibrium.

Example

Consider the infinitely repeated game based on Prisoner's dilemma:

	C	S
C	-5, -5	0, -20
S	-20, 0	-1, -1

The grim trigger profile (τ_1, τ_2) where

$$\tau_i(s^1 \cdots s^T) = \begin{cases} S & T = 0 \\ S & s^\ell = (S, S) \text{ for all } 1 \leq \ell \leq T \\ C & \text{otherwise} \end{cases}$$

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

The intuition behind the grim trigger works as 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 .

Example

Consider the infinitely repeated game based on Prisoner's dilemma:

	C	S
C	-5, -5	0, -20
S	-20, 0	-1, -1

However, other payoffs can be supported by NE. Consider e.g. a strategy profile (τ_1, τ_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 (τ_1, τ_2) is also SPE.

Apparently, $u_1^{avg}(\tau_1, \tau_2) = \frac{9}{10} \cdot (-1) + (-20)/10 = -29/10$ and $u_1^{avg}(\tau_1, \tau_2) = \frac{9}{10}(-1) = -9/10$.

Player 2 gets better payoff than from the Pareto optimal profile (S, S)!

Outline of the 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 called Folk Theorems.

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

In what follows we prove several versions of Folk Theorem concerning achievable payoffs for repeated games.

Ordered by increasing technical and conceptual difficulty, we consider the following variants:

- ▶ Long-run average payoffs & SPE
- ▶ Discounted payoffs & SPE
- ▶ Long-run average payoffs & Nash equilibria

Folk Theorems – Feasible Payoffs

Definition 65

We say that a vector of payoffs $v = (v_1, v_2) \in \mathbb{R}^2$ 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 β_s , here $s \in S$, satisfying $\beta_s \geq 0$ and $\sum_{s \in S} \beta_s = 1$ such that for both $i \in \{1, 2\}$ holds

$$v_i = \sum_{s \in S} \beta_s \cdot u_i(s)$$

We assume that there is $m \in \mathbb{N}$ such that each β_s can be written in the form $\beta_s = \gamma_s/m$.

The following theorems can be extended to a notion of feasible payoffs using *arbitrary, possibly irrational*, coefficients β_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.

Folk Theorems – Long-Run Average & SPE

Theorem 66

Let s^* be a pure strategy Nash equilibrium in G and let $v = (v_1, v_2)$ be a **feasible** vector of payoffs satisfying $v_i \geq u_i(s^*)$ for both $i \in \{1, 2\}$.

Then there is a strategy profile $\tau = (\tau_1, \tau_2)$ in G_{irep} such that

- ▶ τ is a SPE in G_{irep}^{avg}
- ▶ $u_i^{avg}(\tau) = v_i$ for $i \in \{1, 2\}$

Proof: Consider a strategy profile $\tau = (\tau_1, \tau_2)$ in G_{irep} which gives the following behavior:

1. Unless one of the players deviates, the players play **cyclically** all profiles $s \in S$ so that each s is always played for γ_s rounds.
2. Whenever one of the players deviates, then, from that moment on, each player i plays s_i^* .

It is easy to see that $u_i^{avg}(\tau) = v_i$.

We verify that τ is SPE.

Folk Theorems – Long-Run Average & SPE

Fix a history h , we show that $\tau^h = (\tau_1^h, \tau_2^h)$ is a NE in G_{irep}^{avg} .

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

$$w_{\tau^h} = s^* s^* \dots$$

and thus $u_i^{avg}(\tau^h) = u_i(s^*)$.

- ▶ Now if a player i deviates to $\bar{\tau}_i^h$ from τ_i^h in G_{irep}^{avg} , then

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

where s_i^1, s_i^2, \dots are strategies of S_i and s'_{-i} is a strat. of S_{-i} . However, then $u_i^{avg}(\bar{\tau}_i^h, \tau_{-i}^h) \leq u_i(s^*) \leq v_i$ since s^* is a Nash equilibrium and thus $u_i(s_i^k, s^*_{-i}) \leq u_i(s^*)$ for all $k \geq 1$.

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



Folk Theorems – Discounted Payoffs & SPE

Theorem 67

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

- ▶ τ is a SPE in $G_{irep}^{\delta'}$ for every $\delta' \in [\delta, 1)$ and
- ▶ $\lim_{\delta' \rightarrow 1} u_i^{\delta'}(\tau) = v_i$.

Proof: The following claim allows us to reduce the discounted payoff to the long-run-average.

Claim 5

Let τ be a well-defined strategy profile. Then

$$\lim_{\delta \rightarrow 1^-} u_i^\delta(\tau) = u_i^{avg}(\tau)$$

Now to prove Theorem 67, consider the strategy profile $\tau = (\tau_1, \tau_2)$ in G_{irep} from the proof of Theorem 66.

We check the one-shot deviation property in G_{irep}^δ for δ close to 1.

Folk Theorems – Discounted Payoffs & SPE

Fix a history h and consider $\tau^h = (\tau_1^h, \tau_2^h)$.

- ▶ If h does not contain any deviation from 1., then both players follow 1., and $u_i^\delta(\tau^h)$ is close to $u_i^{avg}(\tau^h) = v_i$ for δ close to 1.
- ▶ If h contains any deviation from 1., then $w_{\tau^h} = s^* s^* \dots$ and $u_i^\delta(\tau^h) = u_i(s^*)$.
- ▶ Now assume, w.l.o.g., that player 1 deviates exactly after h , which gives a strategy $\bar{\tau}_1^h$ differing from τ_1^h only on h . Thus $w_{(\bar{\tau}_1^h, \tau_2^h)} = (s'_1, s'_2) s^* s^* \dots$ where s'_1 is a strategy of S_1 and s'_2 is either the next step in the cyclic behavior described by 1. (if h follows 1.), or equal to s_2^* (h does not follow 1.)

Note that for δ close to 1, we have that $u_i^\delta(\bar{\tau}_i^h, \tau_{-i}^h)$ is close to $u_i^{avg}(\bar{\tau}_i^h, \tau_{-i}^h) = u_i(s^*)$.

- ▶ If h follows 1., then $u_1^\delta(\tau^h)$ is close to v_1 which is greater than $u_1(s^*)$ to which $u_1^\delta(\bar{\tau}_1^h, \tau_2^h)$ is close.
- ▶ If h does not follow 1., then $s'_2 = s_2^*$ (players punish due to a deviation in h), and thus $u_1^\delta(\bar{\tau}_1^h, \tau_2^h) \leq u_1(s^*) = u_1^\delta(\tau^h)$. \square

Folk Theorems – Individually Rational Payoffs

Definition 68

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

$$v_i \geq \min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_i, s_{-i})$$

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

Example:

	m	f	r
M	4, 4	-1, 5	3, 0
F	5, -1	1, 1	0, 0
R	0, 3	0, 0	2, 2

Here any (v_1, v_2) such that $v_1 \geq 2$ and $v_2 \geq 1$ is individually rational.

Folk Theorems – Long-Run Average & NE

Theorem 69

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

- ▶ τ is a Nash equilibrium in G_{irep}^{avg}
- ▶ $u_i^{avg}(\tau) = v_i$ for $i \in \{1, 2\}$

Proof: It suffices to use a slightly modified strategy profile $\tau = (\tau_1, \tau_2)$ in G_{irep} from Theorem 66:

- ▶ Unless one of the players deviates, the players play **cyclically** all profiles $s \in S$ so that each s is always played for γ_s rounds.
- ▶ Whenever a player i deviates, the opponent $-i$ plays a strategy $s_{-i}^{min} \in \operatorname{argmin}_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_i, s_{-i})$.

It is easy to see that $u_i^{avg}(\tau) = v_i$.

If a player i deviates, then his long-run average payoff cannot be higher than $\min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_i, s_{-i}) \leq v_i$, so τ is a NE. □

Folk Theorems – Long-Run Average & NE

Theorem 70

If a strategy profile $\tau = (\tau_1, \tau_2)$ is a NE in G_{irep}^{avg} , then $(u_1^{avg}(\tau), u_2^{avg}(\tau))$ is individually rational.

Proof: Suppose that $(u_1^{avg}(\tau), u_2^{avg}(\tau))$ is not individually rational. W.l.o.g. assume that $u_1^{avg}(\tau) < \min_{s_2 \in S_2} \max_{s_1 \in S_1} u_1(s_1, s_2)$.

Now let us consider a new strategy $\bar{\tau}_1$ such that for an arbitrary 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

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

So clearly $u_1^{avg}(\bar{\tau}_1, \tau_2) > u_1^{avg}(\tau)$ which contradicts the fact that (τ_1, τ_2) is a NE. \square

Note that if irrational convex combinations are allowed in the definition of feasibility, then vectors of payoffs for Nash equilibria in G_{irep}^{avg} are exactly feasible and individually rational vectors of payoffs. Indeed, the coefficients β_s in the definition of feasibility are exactly frequencies with which the individual profiles of S are played in the NE.

Folk Theorems – Summary

- ▶ We have proved that "any reasonable" (i.e. feasible and individually rational) vector of payoffs can be justified as payoffs for a Nash equilibrium in G_{irep}^{avg} (where the future has "an infinite weight").
- ▶ Concerning SPE, we have proved that any feasible vector of payoffs dominating a Nash equilibrium in G can be justified as payoffs for SPE in G_{irep}^{avg} .

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

- ▶ For discounted payoffs, we have proved that an arbitrary feasible vector of payoffs strictly dominating a Nash equilibrium in G can be approximated using payoffs for SPE in G_{irep}^{δ} as δ goes to 1. Even this result can be extended to feasible and strictly individually rational payoffs.

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

Summary of Extensive-Form Games

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

- ▶ with perfect information
- ▶ with imperfect information
- ▶ with chance nodes (and both perfect and imperfect information)

We have considered pure, mixed and behavioral strategies.

We have considered Nash equilibria (NE) and subgame perfect equilibria (SPE) in pure and behavioral strategies.

Summary of Extensive-Form Games (Cont.)

For perfect information we have shown that

- ▶ mixed and behavioral strategies are equivalent
- ▶ there is a pure strategy SPE in both pure as well as behavioral strategies
- ▶ SPE can be computed using backward induction in polynomial time

For imperfect information we have shown that

- ▶ mixed and behavioral strategies are not equivalent in general (but they are equivalent for games with perfect recall)
- ▶ 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 (for details see the lecture).

Chance nodes do not interfere with any of the above results.

Summary of Extensive-Form Games (Cont.)

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 existence of pure strategy NE and SPE with respect to existence of NE in the original strategic-form game.

For infinitely repeated games we considered both

- ▶ **discounted payoff**: We have proved that
 - ▶ one-shot deviation property is equivalent to SPE
 - ▶ "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 (Simple Folk Theorem)
- ▶ **long-run average payoff**: We have proved that all feasible and individually rational vectors of payoffs can be achieved by Nash equilibria (a variant of grim trigger)

Games of INcomplete Information
Bayesian Games
Auctions

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 both the resource allocation and payments of the agents.

Auctions: Taxonomy

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 homogeneous 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 they should pay (pay-as-bid vs uniform pricing)
- ▶ *Combinatorial auctions*: Here bidders compete for a set of distinct goods. Each player has a valuation function which assigns values to *subsets* of the set (some goods are useful only in groups etc.) Who wins and what he pays?

(We mostly concentrate on the single-item auctions.)

Single Unit Auctions

There are many single-item auctions, 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, 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-auction*:
 - ▶ *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 a reverse auction it is the *k-th minimum*.) The most prominent special cases are *The First-Price Auction* and *The Second-Price Auction*.

Single Unit Auctions (Cont.)



Observe that

- ▶ the English auction is essentially equivalent to the second price auction if the increments in every round are very small.
There exists a "continuous" version, called 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.

Now the question is, which type of auction is better?

Objectives

The goal of the bidders is clear: To get the item at as low price as possible (i.e., they maximize the difference between their private value and the price they pay)

We consider self-interested non-communicating bidders that are rational and intelligent.

There are at least two goals that may be pursued by the auctioneer (in various settings):

- ▶ Revenue maximization

This may lead to auctions that do not always sell the item to the highest bid

- ▶ Incentive compatibility: We want the bidders to spontaneously bid their true value of the item

This means, that such an auction cannot be strategically manipulated by lying.

Auctions vs Games

Consider *single-item sealed-bid auctions* as strategic form games:

$G = (N, (B_i)_{i \in N}, (u_i)_{i \in N})$ where

- ▶ The set of players N is the set of bidders
- ▶ $B_i = [0, \infty)$ where each $b_i \in B_i$ 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 $b = (b_1, \dots, b_n)$:

$$\text{First Price: } u_i(b) = \begin{cases} v_i - b_i & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{otherwise} \end{cases}$$

$$\text{Second Price: } u_i(b) = \begin{cases} v_i - \max_{j \neq i} b_j & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{otherwise} \end{cases}$$

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

Can we use (possibly imperfect information) extensive-form games?

Incomplete Information Games

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 *action profiles*
 $a = (a_1, \dots, a_n)$.
- ▶ Each T_i is a set of *possible types* of player i ,
Denote by $T = \prod_{i=1}^n T_i$ the set of all *type profiles* $t = (t_1, \dots, t_n)$.
- ▶ u_i is a type-dependent payoff function

$$u_i : A_1 \times \dots \times A_n \times T_i \rightarrow \mathbb{R}$$

Given a profile of actions $(a_1, \dots, a_n) \in A$ and a type $t_i \in T_i$, we write $u_i(a_1, \dots, a_n; t_i)$ to denote the corresponding payoff.

A *pure strategy* of player i is a function $s_i : T_i \rightarrow A_i$. 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 $\prod_{i=1}^n S_i$.

Dominant Strategies

- ▶ A pure strategy s_i *very weakly dominates* s'_i if for every $t_i \in T_i$ the following holds: For all $a_{-i} \in A_{-i}$ we have

$$u_i(s_i(t_i), a_{-i}; t_i) \geq u_i(s'_i(t_i), a_{-i}; t_i)$$

A pure strategy s_i *weakly dominates* s'_i if for every $t_i \in T_i$ the following holds: For all $a_{-i} \in A_{-i}$ we have

$$u_i(s_i(t_i), a_{-i}; t_i) \geq u_i(s'_i(t_i), a_{-i}; t_i)$$

and the inequality is strict for at least one a_{-i}

(Such a_{-i} may be different for different t_i .)

- ▶ A pure strategy s_i *strictly dominates* s'_i if for every $t_i \in T_i$ the following holds: For all $a_{-i} \in A_{-i}$ we have

$$u_i(s_i(t_i), a_{-i}; t_i) > u_i(s'_i(t_i), a_{-i}; t_i)$$

Definition 71

s_i is (*very weakly, weakly, strictly*) *dominant* if it (very weakly, weakly, strictly, resp.) dominates all other pure strategies.

Nash Equilibrium

In order to generalize Nash equilibria to incomplete information games, we use the following notation: Given a pure strategy profile $(s_1, \dots, s_n) \in S$ and a type profile $(t_1, \dots, t_n) \in T$, for every player i write

$$s_{-i}(t_{-i}) = (s_1(t_1), \dots, s_{i-1}(t_{i-1}), s_{i+1}(t_{i+1}), \dots, s_n(t_n))$$

Definition 72

A strategy profile $s = (s_1, \dots, s_n) \in S$ is an *ex-post-Nash equilibrium* if for *every* t_1, \dots, t_n we have that $(s_1(t_1), \dots, s_n(t_n))$ is a Nash equilibrium in the strategic-form game defined by the t_i 's.

Formally, $s = (s_1, \dots, s_n) \in S$ is an *ex-post-Nash equilibrium* if for all $i \in N$ and all t_1, \dots, t_n and all $a_i \in A_i$:

$$u_i(s_1(t_1), \dots, s_n(t_n); t_i) \geq u_i(a_i, s_{-i}(t_{-i}); t_i)$$

Example: Single-Item Sealed-Bid Auctions

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 $b = (b_1, \dots, b_n)$ (i.e. bids) we define

$$\text{First Price:} \quad u_i(b; v_i) = \begin{cases} v_i - b_i & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{otherwise.} \end{cases}$$

$$\text{Second Price:} \quad u_i(b; v_i) = \begin{cases} v_i - \max_{j \neq i} b_j & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{otherwise.} \end{cases}$$

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

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

Second-Price Auction

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

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

Theorem 73

Assume the Second-Price Auction. Then for every player i we have that v_i is a weakly dominant strategy. Also, v is the unique ex-post-Nash equilibrium.

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

$$u_i(v_i, b_{-i}; v_i) \geq u_i(b_i, b_{-i}; v_i)$$

with the strict inequality for at least one b_{-i} .

Intuitively, assume that player i bids b_i against b_{-i} and compare his payoff with the payoff he obtains by playing v_i against b_{-i} .

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

Second-Price Auction (Cont.)

Case $b_i < v_i$: We distinguish three sub-cases depending on b_{-i} .

A. If $b_i > \max_{j \neq i} b_j$, then

$$u_i(b_i, b_{-i}; v_i) = v_i - \max_{j \neq i} b_j = u_i(v_i, b_{-i}; v_i)$$

Intuitively, player i wins and pays the price $\max_{j \neq i} b_j < b_i$. However, then bidding v_i , player i wins and pays $\max_{j \neq i} b_j$ as well.

B. If there is $k \neq i$ such that $b_k > \max_{j \neq k} b_j$, then

$$u_i(b_i, b_{-i}; v_i) = 0 \leq u_i(v_i, b_{-i}; v_i)$$

Moreover, if $b_i < b_k < v_i$, then we get the strict inequality

$$u_i(b_i, b_{-i}; v_i) = 0 < v_i - b_k = u_i(v_i, b_{-i}; v_i)$$

Intuitively, if another player k wins, then player i gets 0 and increasing b_i to v_i 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 .

C. If there are $k \neq \ell$ such that $b_k = b_\ell = \max_j b_j$, then

$$u_i(b_i, b_{-i}; v_i) = 0 \leq u_i(v_i, b_{-i}; v_i)$$

Intuitively, there is a tie in (b_i, b_{-i}) and hence all players get 0.

Second-Price Auction (Cont.)

Case $b_i > v_i$: We distinguish four sub-cases depending on b_{-i} .

A. If $b_i > \max_{j \neq i} b_j > v_i$, then

$$u_i(b_i, b_{-i}; v_i) = v_i - \max_{j \neq i} b_j < 0 = u_i(v_i, b_{-i}; v_i)$$

So in this case the inequality is strict.

B. If $b_i > v_i \geq \max_{j \neq i} b_j$, then

$$u_i(b_i, b_{-i}; v_i) = v_i - \max_{j \neq i} b_j = u_i(v_i, b_{-i}; v_i)$$

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 .

C. If there is $k \neq i$ such that $b_k > \max_{j \neq k} b_j > v_i$, then

$$u_i(b_i, b_{-i}; v_i) = 0 = u_i(v_i, b_{-i}; v_i)$$

D. If there are $k \neq k'$ such that $b_k = b_{k'} = \max_j b_j > v_i$, then

$$u_i(b_i, b_{-i}; v_i) = 0 = u_i(v_i, b_{-i}; v_i)$$

First-Price Auction

Consider 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.

Question: Are there any dominant strategies?

Answer: No, to obtain a contradiction, assume that s_i is a very weakly dominant strategy.

Intuitively, if player i wins against some bids of his opponents, then his bid is strictly higher than bids of all his opponents. Thus he may slightly decrement his bid and still win with a better payoff.

Formally, assume that all opponents bid 0, i.e., $b_j = 0$ for all $j \neq i$, and consider $v_i > 0$.

If $s_i(v_i) > 0$, then

$$u_i(s_i(v_i), b_{-i}; v_i) = v_i - s_i(v_i) < v_i - s_i(v_i)/2 = u_i(s_i(v_i)/2, b_{-i}; v_i)$$

If $s_i(v_i) = 0$, then

$$u_i(s_i(v_i), b_{-i}; v_i) = 0 < v_i/2 = u_i(v_i/2, b_{-i}; v_i)$$

Hence, s_i cannot be weakly dominant.

First-Price Auction (Cont.)

Question: Is there a pure strategy Nash equilibrium?

Answer: No, assume that (s_1, \dots, s_n) is a Nash equilibrium.

If there are v_1, \dots, v_n such that some player i wins, i.e., his bid $s_i(v_i)$ satisfies $s_i(v_i) > \max_{j \neq i} s_j(v_j)$, then

$$\begin{aligned} u_i(s_i(v_i), s_{-i}(v_{-i}); v_i) &= v_i - s_i(v_i) \\ &< v_i - (s_i(v_i) - \varepsilon) = u_i(s_i(v_i) - \varepsilon, s_{-i}(v_{-i}); v_i) \end{aligned}$$

for $\varepsilon > 0$ small enough to satisfy $s_i(v_i) - \varepsilon > \max_{j \neq i} s_j(v_j)$
(i.e., player i may help himself by decreasing the bid a bit)

Assume that for no v_1, \dots, v_n there is a winner (this itself is a bit weird). Consider $0 < v_1 < \dots < v_n$. Since there is no winner, there are two players i, j such that $i < j$ satisfying

$$s_j(v_j) = s_i(v_i) \geq \max_{\ell} s_{\ell}(v_{\ell})$$

But then, due to our assumption, $s_j(v_j) = s_i(v_i) \leq v_i < v_j$ and thus

$$u_j(s_j(v_j), s_{-j}(v_{-j}); v_j) = 0 < v_j - (s_j(v_j) + \varepsilon) = u_j(s_j(v_j) + \varepsilon, s_{-j}(v_{-j}); v_j)$$

for $\varepsilon > 0$ small enough to satisfy $s_j(v_j) + \varepsilon < v_j$.

(i.e., player j can help himself by increasing his bid a bit)

Second Price Auction:

- ▶ There is an ex-post Nash equilibrium in weakly dominant strategies
- ▶ It is incentive compatible (players are self-motivated to bid their private values)

First Price Auction:

- ▶ There are neither dominant strategies, nor ex-post Nash equilibria

Question: Can we modify the model in such a way that First Price Auction has a solution?

Answer: Yes, give the players at least some information about private values of other players.

Bayesian Games

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 a distribution on types, i.e.,

- ▶ $N = \{1, \dots, n\}$ is a set of players,
- ▶ A_i is a set of *actions* available to player i ,
- ▶ T_i is a set of *possible types* of player i ,

Recall that $T = \prod_{i=1}^n T_i$ is the set of type profiles, and that $A = \prod_{i=1}^n A_i$ is the set of action profiles.

- ▶ u_i is a type-dependent payoff function

$$u_i : A_1 \times \dots \times A_n \times T_i \rightarrow \mathbb{R}$$

- ▶ P is a *(joint) probability distribution over T* called *common prior*.

Formally, P is a probability measure over an appropriate measurable space on T . However, I will not go into measure theory and consider only two special cases: finite T (in which case $P : T \rightarrow [0, 1]$ so that $\sum_{t \in T} P(t) = 1$) and $T_i = \mathbb{R}$ for all i (in which case I assume that P is determined by a (joint) density function p on \mathbb{R}^n).

Bayesian Games: Strategies & Payoffs

A play proceeds as follows:

- ▶ First, a type profile $(t_1, \dots, t_n) \in T$ is randomly chosen according to P .
- ▶ Then each player i learns his type t_i .
(It is a common knowledge that every player knows his own type but not the types of other players.)
- ▶ Each player i chooses his action based on t_i .
- ▶ Each player receives his payoff $u_i(a_1, \dots, a_n; t_i)$.

A *pure strategy* for player i is a function $s_i : T_i \rightarrow A_i$.

As before, we use S to denote the set of pure strategy profiles.

Properties

- ▶ We assume that u_i depends only on t_i and not on t_{-i} . This is called **private values** model and can be used to model auctions. This model can be extended to **common values** by using $u_i(a_1, \dots, a_n; t_1, \dots, t_n)$.
- ▶ We assume the *common prior* P . This means that all players have *the same* beliefs about the type profile. This assumption is rather strong. More general models allow each player to have
 - ▶ his own individual beliefs about types
 - ▶ ... his own beliefs about beliefs about types
 - ▶ beliefs about beliefs about beliefs about types
 - ▶
 - ▶ (we get an infinite hierarchy)

There is a generic result of Harsanyi saying that the 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 his beliefs.

(This does not mean that common prior suffices.)

Example: Battle of Sexes

Assume that player 1 may suspect that player 2 is angry with him/her (the choice is yours) but cannot be sure.

In other words, there are two types of player 2 giving two different games.

Formally we have a Bayesian Game

$G = (N, (A_i)_{i \in N}, (T_i)_{i \in N}, (u_i)_{i \in N}, P)$ where

- ▶ $N = \{1, 2\}$
- ▶ $A_1 = A_2 = \{F, O\}$
- ▶ $T_1 = \{t_1\}$ and $T_2 = \{t_2^1, t_2^2\}$
- ▶ The payoffs are given by

		t_2^1		t_2^2	
		F	O	F	O
t_1	F	2, 1	0, 0	2, 0	0, 2
	O	0, 0	1, 2	0, 1	1, 0

- ▶ $P(t_2^1) = P(t_2^2) = \frac{1}{2}$

Example: Single-Item Sealed-Bid Auctions

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 $N = \{1, \dots, n\}$ is the set of bidders
- ▶ $B_i = [0, \infty)$ where each action $b_i \in B_i$ corresponds to the bid
- ▶ $V_i = \mathbb{R}$ where each type v_i corresponds to the private value
- ▶ Let $v_i \in V_i$ be the type of player i (i.e. his private value), then given an action profile $b = (b_1, \dots, b_n)$ (i.e. bids) we define

$$\text{First Price:} \quad u_i(b; v_i) = \begin{cases} v_i - b_i & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{otherwise.} \end{cases}$$

$$\text{Second Price:} \quad u_i(b; v_i) = \begin{cases} v_i - \max_{j \neq i} b_j & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{otherwise.} \end{cases}$$

- ▶ P is a probability distribution of the private values such that $P(v \in [0, \infty)^n) = 1$. For example, we may (and will) assume that each v_i is chosen independently and uniformly from $[0, v_{\max}]$ where v_{\max} is a given number. Then P is uniform on $[0, v_{\max}]^n$.

Finite-Type Bayesian Games: Payoffs

For now, let us assume that each player has only finitely many types, i.e., T is finite.

Given a type profile $t = (t_1, \dots, t_n)$, we denote by $P(t_{-i} | t_i)$ the *conditional probability* that the opponents of player i have the type profile t_{-i} conditioned on player i having t_i , i.e.,

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

Intuitively, $P(t_{-i} | t_i)$ is the maximum information player i may squeeze out of P about possible types of other players once he learns his own type t_i .

Given a pure strategy profile $s = (s_1, \dots, s_n)$ and a type $t_i \in T_i$ of player i the *expected payoff* for player i is

$$u_i(s; t_i) = \sum_{t_{-i} \in T_{-i}} P(t_{-i} | t_i) \cdot u_i(s_1(t_1), \dots, s_n(t_n); t_i)$$

(this is the conditional expectation of u_i assuming the type t_i of player i)

Example: Battle of Sexes

		t_2^1			t_2^2	
		F	O		F	O
$t_1 :$	F	2,1	0,0	F	2,0	0,2
	O	0,0	1,2		O	0,1

$$P(t_2^1) = P(t_2^2) = \frac{1}{2}$$

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

- ▶ $u_1(s_1, s_2; t_1) = \frac{1}{2} \cdot 2 + \frac{1}{2} \cdot 0 = 1$
- ▶ $u_2(s_1, s_2; t_2^1) = 1$ and $u_2(s_1, s_2; t_2^2) = 2$

Infinite-Type Bayesian Games: Payoffs

Now assume that for each player i we have $T_i = \mathbb{R}$ and thus that $T = \mathbb{R}^n$. The concrete type is randomly chosen according to P , denote by $\mathbf{t} = (\mathbf{t}_1, \dots, \mathbf{t}_n)$ the corresponding random vector with distribution P (each \mathbf{t}_i is a random variable giving a type of player i).

Assume that the type \mathbf{t} is absolutely continuous which means that there is a (joint) density function p such that for all rectangles $R = [a_1, b_1] \times \dots \times [a_n, b_n]$

$$P[\mathbf{t} \in R] = \int_{a_1}^{b_1} \dots \int_{a_n}^{b_n} p(t_1, \dots, t_n) dt_n \dots dt_1$$

Let p_i be the marginal density function of \mathbf{t}_i , i.e.,

$$p_i(t_i) = \int_{T_{-i}} p(\mathbf{t}_i, \mathbf{t}_{-i}) d\mathbf{t}_{-i}$$

The conditional density of $\mathbf{t}_{-i} = (\mathbf{t}_1, \dots, \mathbf{t}_{i-1}, \mathbf{t}_{i+1}, \dots, \mathbf{t}_n)$ conditioned on $\mathbf{t}_i = t_i$ where $p_i(t_i) > 0$ is

$$p(\mathbf{t}_{-i} | t_i) = p(\mathbf{t}) / p_i(t_i)$$

(Here $t = (t_1, \dots, t_n)$ is a type profile.)

Infinite-Type Bayesian Games: Payoffs

Given a pure strategy profile $s = (s_1, \dots, s_n)$ and a type $t_i \in T_i$ of player i , the *expected payoff* for player i is

$$u_i(s; t_i) = \int_{T_{-i}} u_i(s_1(t_1), \dots, s_n(t_n); t_i) p(t_{-i} | t_i) dt_{-i}$$

Example: 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 $v = (v_1/2, \dots, v_n/2)$ (i.e., each player i plays $v_i/2$). What is $u_i(v; v_i)$?

$$\begin{aligned} u_i(v; v_i) &= P(\text{player } i \text{ wins}) \cdot v_i/2 + P(\text{player } i \text{ loses}) \cdot 0 \\ &= P(\text{all players except } i \text{ bid less than } v_i/2) \cdot v_i/2 \\ &= \left(\frac{v_i}{2v_{\max}} \right)^{n-1} \cdot v_i/2 \\ &= \frac{v_i^n}{2^n v_{\max}^{n-1}} \end{aligned}$$

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 $\frac{1}{2}$.

- ▶ risk neutral person has no preference
- ▶ risk averse person prefers the first alternative
- ▶ risk seeking person prefers the second one

Dominance and Nash Equilibria

A pure strategy s_i *weakly dominates* s'_i if for every $t_i \in T_i$ the following holds: For all $s_{-i} \in S_{-i}$ we have

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

and the inequality is strict for at least one s_{-i} .

The other modes of dominance are defined analogously. Dominant strategies are defined as usual.

Definition 74

A pure strategy profile $s = (s_1, \dots, s_n) \in S$ in the Bayesian game is a *pure strategy Bayesian Nash equilibrium* 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

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

Example: Battle of Sexes

		t_2^1		t_2^2	
		F	O	F	O
t_1	F	2, 1	0, 0	2, 0	0, 2
	O	0, 0	1, 2	0, 1	1, 0

$$P(t_2^1) = P(t_2^2) = \frac{1}{2}$$

Use the following notation: $(X, (Y, Z))$ means that player 1 plays $X \in \{F, O\}$, and player 2 plays $Y \in \{F, O\}$ if his/her type is t_2^1 and $Z \in \{F, O\}$ otherwise.

Are there pure strategy Bayesian Nash equilibria?

$(F, (F, O))$ is a Bayesian NE.

Even though O is preferred by player 2, the outcome (O, O) cannot occur with a positive probability in any BNE.

- ▶ 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 BNE:
 - ▶ The expected payoff of player 1 at $(O, (O, F))$ is $\frac{1}{2}$
 - ▶ The expected payoff of player 1 at $(F, (O, F))$ is 1

Second Price Auction

Consider the second-price sealed-bid auction as a Bayesian game where the types of players are chosen according to an arbitrary distribution.

Proposition 4

In a second-price sealed-bid auction, with any probability distribution P , the truth revealing profile of bids, i.e., $v = (v_1, \dots, v_n)$, is a weakly dominant strategy profile.

Proof.

The exact same proof as for the strict incomplete information games. Indeed, we do not need to assume that the players have a common prior for this! □

First Price Auction

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?

Proposition 5

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 $s = (s_1, \dots, s_n)$ where $s_i(v_i) = \frac{n-1}{n} v_i$ for every player i and every value v_i . Then s is a Bayesian Nash equilibrium.

Proof. We show that $s_i(v_i) = \frac{n-1}{n} v_i$ is the best response to s_{-i} for all i . Let us fix i and consider a pure strategy s'_i of player i .

Fix v_i and define $b_i = s'_i(v_i)$. We show (see the greenboard) that $b_i = \frac{n-1}{n} v_i$ maximizes $u_i(b_i, s_{-i}; v_i)$. This holds for all v_i , and thus $s'_i = s_i$ is the best response to s_{-i} . □

First Price Auction (Cont.)

More generally, assume only that the private values v_i are identically and independently distributed on $[v_{\min}, v_{\max}]$ (this is called **independent private values** model). Let $F(x)$ be the cumulative distribution function of the private value (for each player).

Let us restrict to *strictly increasing strategies*.

Note that this restriction is quite reasonable, intuitively it means, that the higher the private value, the higher is the bid.

Then one may show that there is a symmetric Bayesian Nash equilibrium (s_1, \dots, s_n) where each s_i is defined by

$$s_i(v_i) = v_i - \frac{\int_{v_{\min}}^{v_{\max}} [F(v_i)]^{n-1} dx}{[F(v_i)]^{n-1}}$$

That is, in particular, the bid is always smaller than the private value.

Expected Revenue

Consider the first and second price sealed-bid auctions. For simplicity, assume that the type of each player is chosen independently and uniformly from $[0, 1]$.

What is the expected revenue of the auctioneer from these two auctions when the players play the corresponding Bayesian NE?

- ▶ In the first-price auction, players bid $\frac{n-1}{n} v_i$. Thus the probability distribution of the revenue is

$$F(x) = P(\max_j \frac{n-1}{n} v_j \leq x) = P(\max_j v_j \leq \frac{nx}{n-1}) = \left(\frac{nx}{n-1}\right)^n$$

It is straightforward to show that then 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 is

$$F(x) = P(\max_j v_j \leq x) + \sum_{i=1}^n P(v_i > x \text{ and for all } j \neq i, v_j \leq x)$$

Amazingly, this also gives the expectation $\frac{n-1}{n+1}$.

Revenue Equivalence (Cont.)

The result from the previous slide is a special case of a rather general **revenue equivalence theorem**, first proved by Vickrey (1961) and then generalized by Myerson (1981).

Both Vickrey and Myerson were awarded Nobel Prize in economics for their contribution to the auction theory.

Theorem 75 (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 of $v_i \notin [v_{\min}, v_{\max}]$ is zero). Then any **efficient** auction mechanism in which any player with value v_{\min} has an expected payoff zero yields the same expected revenue.*

Here efficient means that the auction has a symmetric and increasing Bayesian Nash equilibrium and always allocates the item to the player with the highest bid.