

Game Theory

9. Extensive Games with Imperfect Information

Albert-Ludwigs-Universität Freiburg



Bernhard Nebel and Robert Mattmüller

1 Motivation



- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Summary

Motivation



- **So far:** All state information is completely known by all players
- **Often in practice:** Only partial knowledge (e.g. card games)
- Extensive games **with imperfect information** model such situations using **information sets**, which are sets of histories.
- **Idea:** Decision points are now information sets.
- **Strategies:** **Mixed** (over pure strategies) or **behavioral** (collections of independent mixed decisions for each information set)
- Different from **incomplete** information games, in which there is uncertainty about the utility functions of the other players.

- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Summary

2 Definitions



- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Summary

Definition (Extensive game)

An **extensive game** is a tuple $\Gamma = \langle N, H, P, f_c, (\mathcal{I}_i)_{i \in N}, (u_i)_{i \in N} \rangle$ that consists of:

- A finite non-empty set N of **players**.
 - A set H of (finite or infinite) sequences, called **histories**, such that
 - it contains the empty sequence $\langle \rangle \in H$,
 - H is closed under prefixes: if $\langle a^1, \dots, a^k \rangle \in H$ for some $k \in \mathbb{N} \cup \{\infty\}$, and $l < k$, then also $\langle a^1, \dots, a^l \rangle \in H$, and
 - H is closed under limits: if for some infinite sequence $\langle a^i \rangle_{i=1}^\infty$, we have $\langle a^i \rangle_{i=1}^k \in H$ for all $k \in \mathbb{N}$, then $\langle a^i \rangle_{i=1}^\infty \in H$.
- All infinite histories and all histories $\langle a^i \rangle_{i=1}^k \in H$, for which there is no a^{k+1} such that $\langle a^i \rangle_{i=1}^{k+1} \in H$ are called **terminal histories** Z . Components of a history are called **actions**.

Motivation
Definitions
Recall
Strategies and outcomes
Solution Concepts
Summary


Definition (Extensive game, ctd.)

- A **player function** $P : H \setminus Z \rightarrow N \cup \{c\}$ that determines which player's turn it is to move after a given nonterminal history, c signifying a **chance move**.
 - $f_c(\cdot|h)$ is a **probability distribution** over $A(h)$.
 - \mathcal{I}_i is the **information partition** for player i of $\{h \in H | P(h) = i\}$ with the property that $A(h) = A(h')$ whenever h and h' are in the same member of the partition. Members of the partition $I_i \in \mathcal{I}_i$ are called **information sets**.
 - For each player $i \in N$, a **utility function** (or **payoff function**) $u_i : Z \rightarrow \mathbb{R}$ defined on the set of terminal histories.
- Γ is **finite**, if H is finite; **finite horizon**, if histories are bounded.

Motivation
Definitions
Recall
Strategies and outcomes
Solution Concepts
Summary

Example

Example



After player 1 chooses L , player 2 makes a move (A or B) player 1 cannot observe.

Motivation
Definitions
Recall
Strategies and outcomes
Solution Concepts
Summary

Simultaneous moves

- We have already chance moves, but could we extend the model with **simultaneous moves** as well?
- Actually, we can already **model** them somehow.
- In the example game after the history $\langle L \rangle$, we have essentially a simultaneous move of player 1 and 2:
 - When player 2 moves, he does not know what player 1 will do.
 - After player 2 has made his move, player 1 does not know whether A or B was chosen.
 - Only after both players have acted, they are presented with the outcome.


Motivation
Definitions
Recall
Strategies and outcomes
Solution Concepts
Summary

3 Recall

Recall

Information sets can be arbitrary. However, often we want to assume that agents always remember what they have learned before and which actions they have performed: **Perfect recall**.

Example (Imperfect recall)



- Left: Player 1 forgets that he made a move!
- Right: Player 1 cannot remember what his last move was.

Experience record

Definition (Experience record)

Given a history h of an extensive game, the function $X_i(h)$ is the sequence consisting of information sets that player i encounters in h and the actions that player i takes at them. X_i is called the **experience record** of player i in h .

Example

In our example game, Player 1 encounters two information sets in the history $h = \langle L, A \rangle$, namely $\langle \rangle$ and $\{\langle L, A \rangle, \langle L, B \rangle\}$. In the first information set, he chooses L . So $X_1(h) = \langle \langle \rangle, L, \{\langle L, A \rangle, \langle L, B \rangle\} \rangle$.

Perfect recall

Definition (Perfect Recall)

An extensive game has **perfect recall** if for each player i , we have $X_i(h) = X_i(h')$ whenever the histories h and h' are in the same information set of player i .

Example

In our example game, the only non-singleton information set satisfies the condition, since for $h = \langle L, A \rangle$ and $h' = \langle L, B \rangle$ we have $X_1(h) = X_1(h') = \langle \langle \rangle, L, \{\langle L, A \rangle, \langle L, B \rangle\} \rangle$. For the imperfect recall examples, the actions are different for the two histories ending up in the non-singleton information set.

In most cases, our games will have perfect recall.

4 Strategies and outcomes

Pure strategies

Definition (Pure strategy in an extensive game)

A **pure strategy** of a player i in an extensive game $\Gamma = \langle N, H, P, f_C, (\mathcal{I}_i)_{i \in N}, (u_i)_{i \in N} \rangle$ is a function s_i that assigns an action from $A(I_i)$ to each information set I_i .

Remark: Note that the outcome of a strategy profile s is now a probability distribution (because of the chance moves).

Remark: Because of the chance moves and because of the imperfect information, it probably makes more sense to consider randomized strategies.


Mixed and behavioral strategies

Definition (Mixed and behavioral strategies)

A **mixed strategy** of a player i in an extensive game $\Gamma = \langle N, H, P, f_C, (\mathcal{I}_i)_{i \in N}, (u_i)_{i \in N} \rangle$ is a probability distribution over the set of i 's pure strategies. A **behavioral strategy** of player i is a collection $(\beta_i(I_i))_{I_i \in \mathcal{I}_i}$ of independent probability distributions, where $\beta_i(I_i)$ is a probability distribution over $A(I_i)$. For any history $h \in I_i \in \mathcal{I}_i$ and action $a \in A(h)$, we denote by $\beta_i(h)(a)$ the probability $\beta_i(I_i)(a)$ assigned by $\beta_i(I_i)$ to action a .

Exampel

Example



- Player 1 has four **pure strategies** (two information sets, two actions at each).
- A **mixed strategy** is a probability distribution over those.
- A **behavioral strategy** is a pair of probability distributions, one for $\{\langle \rangle\}$ and one for $\{\langle L, A \rangle, \langle L, B \rangle\}$.

Outcomes



The outcome of a (mixed or behavioral) strategy profile σ is a probability distribution over histories $O(\sigma)$, resulting from following the individual strategies.

- For any history $h = \langle a^1, \dots, a^k \rangle$ define a **pure strategy s_i of i to be consistent with h** if for any subhistory $h' = \langle a^1, \dots, a^{\ell} \rangle$ with $P(h') = i$ we have $s_i(h') = a^{\ell+1}$.
- For any history, let $\pi_i(h)$ be the sum of probabilities of pure strategies s_i from σ_i consistent with h .
- Then for any mixed profile σ , the probability that $O(\sigma)$ assigns to a terminal history h is: $\prod_{i \in NU\{c\}} \pi_i(h)$.
- For any behavioral profile β , the probability that $O(\beta)$ assigns to $h = \langle a^1, \dots, a^k \rangle$ is:

$$\prod_{k=0}^{K-1} \beta_{P(\langle a^1, \dots, a^k \rangle)}(\langle a^1, \dots, a^k \rangle)(a^{k+1})$$

- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Summary

Outcome equivalence



Definition

Two (mixed or behavioral) strategies of a player i are called **outcome-equivalent** if for every partial profile of pure strategies of the other players, the two strategies induce the same outcome.

Question: Can we find outcome-equivalent mixed strategies for behavioral strategies and vice versa?


Partial answer: Sometimes.

- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Summary

Counter-example (1)



Example (Behavioral strategy without a mixed strategy)




- A behavioral strategy assigning non-zero probability to a and b generates outcomes $\langle a, a \rangle$, $\langle a, b \rangle$, and $\langle b \rangle$ with non-zero probability
- Since there are only the pure strategies playing a or b , no mixed strategy can produce $\langle a, b \rangle$.

- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Summary

Counter-example (2)



Example (Mixed strategy without a behavioral strategy)



- Mix the two pure strategies LL and RR equally, resulting in the distribution $(1/2, 0, 0, 1/2)$.
- No behavioral strategy can accomplish this.

- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Summary

Equivalence of behavioral and mixed strategies



If we restrict ourselves to games with perfect recall, however, everything works.

Theorem (Equivalence of mixed and behavioral strategies (Kuhn))

In a game of perfect recall, any mixed strategy of a given agent can be replaced by an equivalent behavioral strategy, and any behavioral strategy can be replaced by an equivalent mixed strategy.

- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Summary

5 Solution Concepts



- Assessments
- Sequential Rationality
- Sequential equilibrium
- Examples

- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Assessments
- Sequential Rationality
- Sequential equilibrium
- Examples
- Summary


Expected utility



Similar to the case of mixed strategies for strategic games, we define the utility for mixed and behavioral strategies as expected utility, summing over all histories:

$$U_i(\sigma) = \sum_{h \in H} u_i(h) \cdot O(\sigma)(h)$$

Example



- Mixed strategy (mixing a and b) σ : $U_1(\sigma) = 0$.
- Behavioral strategy β with $p = 1/2$ for a: $U_1(\beta) = 1/4$.

- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Assessments
- Sequential Rationality
- Sequential equilibrium
- Examples
- Summary

Nash equilibria



Definition (Nash equilibrium in mixed strategies)

A **Nash equilibrium in mixed strategies** is a profile σ^* of mixed strategies with the property that for every player i :

$$U_i(\sigma_{-i}^*, \sigma_i^*) \geq U_i(\sigma_{-i}^*, \sigma_i) \text{ for every mixed strategy } \sigma_i \text{ of } i.$$

Note: Support lemma applies here as well.

Definition (Nash equilibrium in behavioral strategies)

A **Nash equilibrium in behavioral strategies** is a profile β^* of mixed strategies with the property that for every player i :

$$U_i(\beta_{-i}^*, \beta_i^*) \geq U_i(\beta_{-i}^*, \beta_i) \text{ for every behavioral strategy } \beta_i \text{ of } i.$$


Remark: Equivalent, provided we have perfect recall.

- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Assessments
- Sequential Rationality
- Sequential equilibrium
- Examples
- Summary

Eliminating imperfect equilibria



Example




Nash equilibria: (M,L) and (L,R)
 Unreasonable ones: (L,R), because in the info set of player 2, L dominates R

- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Assessments
- Sequential Rationality
- Sequential equilibrium
- Examples
- Summary

How have we got here?



Example



Nash equilibria: (L,R)
 What should player 2 do, when he ends up in his info set?
 Depends on his belief: if probability that M has been played $\geq 1/2$, then R is optimal, otherwise L.

- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Assessments
- Sequential Rationality
- Sequential equilibrium
- Examples
- Summary

Assessments



Let us take the beliefs about what has been played into account when defining an equilibrium.

Definition (Assessment)

An **assessment** in an extensive game is a pair (β, μ) , where β is a profile of behavioral strategies and μ is a function that assigns to every information set a probability distribution on the set of histories in the information set.

$\mu(I)(h)$ is the probability that player $P(I)$ assigns to the history $h \in I$, given I is reached. We have to modify the **outcome** function. Let $h^* = \langle a^1, \dots, a^K \rangle$ be a terminal history. Then:

- $O(\beta, \mu|I)(h^*) = 0$, if there is no subhistory of h^* in I ,
- $O(\beta, \mu|I)(h^*) = \mu(I)(h) \cdot \prod_{k=L}^{K-1} \beta_{P(\langle a^1, \dots, a^k \rangle)}(\langle a^1, \dots, a^k \rangle)(a^{k+1})$, if the subhistory $\langle a^1, \dots, a^L \rangle$ of h^* is in I with $L < K$.

- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Assessments
- Sequential Rationality
- Sequential equilibrium
- Examples
- Summary

Sequential rationality



Similar to O , we extend U_i : $U_i(\beta, \mu|I_i) = O(\beta, \mu|I)(h^*) \cdot u_i(h^*)$.

Definition (Sequential rationality)

Let Γ be an extensive game with perfect recall. The assessment (β, μ) is **sequentially rational** if for every player i and every information $I_i \in \mathcal{I}_i$ we have

$$U_i(\beta, \mu|I_i) \geq U_i((\beta_{-i}, \beta'_i), \mu|I_i) \text{ for every } \beta'_i \text{ of } i.$$

Note: μ could be arbitrary!

- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Assessments
- Sequential Rationality
- Sequential equilibrium
- Examples
- Summary

Consistency with strategies



We would at least require that the beliefs are consistent with the strategies, meaning they should be derived by the strategies.

In our earlier example, player 2's belief should be derived from the behavioral strategy of player 1. E.g., the probability that M has been played should be:

$$\mu(\{\langle M, \langle R \rangle\})(M) = \beta_1(\langle \cdot \rangle)(M) / (\beta_1(\langle \cdot \rangle)(M) + \beta_1(\langle \cdot \rangle)(R)).$$

In other words, we use Bayes' rule to determine μ . However, what to do when the denominator is 0?

- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Assessments
- Sequential Rationality
- Sequential equilibrium
- Examples
- Summary

Structural consistency



By viewing an assessment as a limit of a sequence of **completely mixed** strategy profiles (all strategies are in the support), one can enforce the Bayes' condition also on information set that are not reached by an equilibrium profile.

Definition (Structural consistency)

Let Γ by a finite extensive game with perfect recall. An assessment (β, μ) is **structural consistent** if there is a sequence $((\beta^n, \mu^n))_{n=1}^\infty$ of assessments that converges to (β, μ) in Euclidian space and has the properties that each strategy profile β^n is completely mixed and that each belief system μ_n is derived from β_n using Bayes' rule.

Note: Kreps (1990) wrote: "a lot of bodies are buried in this definition."

- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Assessments
- Sequential Rationality
- Sequential equilibrium
- Examples
- Summary

Sequential equilibrium



Definition (Sequential equilibrium)

An assessment is a **sequential equilibrium** of a finite extensive game with perfect recall if it is sequentially rational and structural consistent.

Note: There is always at least one such equilibrium.


Note: In an extensive game with **perfect** information, (β, μ) is sequential equilibrium iff β is a **subgame-perfect equilibrium**.

- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Assessments
- Sequential Rationality
- Sequential equilibrium
- Examples
- Summary

Example 1



Example




Let (β, μ) be as follows: $\beta_1(L) = 1, \beta_2(R) = 1,$
 $\mu(\{\langle M, \langle R \rangle\})(M) = \alpha$ for $0 \leq \alpha \leq 1.$ (β, μ) is consistent since
 $\beta_1^\epsilon = (1 - \epsilon, \alpha\epsilon, (1 - \alpha)\epsilon), \beta_2^\epsilon = (\epsilon, 1 - \epsilon),$ and
 $\mu^\epsilon(\{\langle M, \langle R \rangle\})(M) = \alpha$ converges to (β, μ) for $\epsilon \rightarrow 0.$ For
 $\alpha \geq 1/2, (\beta, \mu)$ is sequentially rational.

- Motivation
- Definitions
- Recall
- Strategies and outcomes
- Solution Concepts
- Assessments
- Sequential Rationality
- Sequential equilibrium
- Examples
- Summary

Example 2: Selten's horse

Example



Two types of NE (for $I = \{\langle D \rangle, \langle C, d \rangle\}$):


- 1 $\beta_1(\langle \rangle)(D) = 1, 1/3 \leq \beta_2(\langle C \rangle)(c) \leq 1, \beta_3(I)(L) = 1$
- 2 $\beta_1(\langle \rangle)(C) = 1, \beta_2(\langle C \rangle)(c) = 1, 3/4 \leq \beta_3(I)(R) \leq 1.$

Are these also sequential equilibria?

Motivation
Definitions
Recall
Strategies and outcomes
Solution Concepts
Assessments
Sequential Rationality
Sequential equilibrium
Examples
Summary

Selten's horse: Type 1 Nash Equilibrium

Example




$\beta_1(\langle \rangle)(D) = 1, 1/3 \leq \beta_2(\langle C \rangle)(c) \leq 1, \beta_3(I)(L) = 1$ violates sequential rationality for player 2!

Motivation
Definitions
Recall
Strategies and outcomes
Solution Concepts
Assessments
Sequential Rationality
Sequential equilibrium
Examples
Summary

Selten's horse: Type 2 Nash Equilibrium

Example



For each NE, $\beta_1(\langle \rangle)(C) = 1, \beta_2(\langle C \rangle)(c) = 1, 3/4 \leq \beta_3(I)(R) \leq 1$, there exists a sequential equilibrium (β, μ) with $\mu(I)(D) = 1/3$.

For consistency consider: $\beta_1^\varepsilon(\langle \rangle)(C) = 1 - \varepsilon$,
 $\beta_2^\varepsilon(\langle C \rangle)(c) = 2\varepsilon/(1 - \varepsilon), \beta_3^\varepsilon(I)(R) = \beta_3(I)(R) - \varepsilon.$
 Note: $\beta_1^\varepsilon(\langle \rangle)(D) + (\beta_1^\varepsilon(\langle \rangle)(C) \cdot \beta_2^\varepsilon(\langle C \rangle)(d)) = 3\varepsilon.$

Motivation
Definitions
Recall
Strategies and outcomes
Solution Concepts
Assessments
Sequential Rationality
Sequential equilibrium
Examples
Summary

6 Summary

Motivation
Definitions
Recall
Strategies and outcomes
Solution Concepts
Summary

- Extensive games with **imperfect information** can model situations, in which the player know only part of the world.
- Modeled by **information sets**, which are the histories, an agent cannot distinguish.
- **Perfect recall** requires that agents remember know what they have done and learned.
- Without it, a number of results do not hold.
- Strategies can be **mixed** or **behavioral**, which is equivalent in the case of perfect recall.
- Nash equilibria can be defined this way, however, similar to perfect information games, are not always reasonable.
- **Sequential equilibria** are the refinement, which is based on **assessments**.

Motivation

Definitions

Recall

Strategies and outcomes

Solution Concepts

Summary