Part I

Foundations

Chapter 1

Representing a Game

Typically we have two ways to represent a game:

- Extensive form (game tree)
- Normal form (payoff matrix)

1.1 Game Tree

Example (Chess).

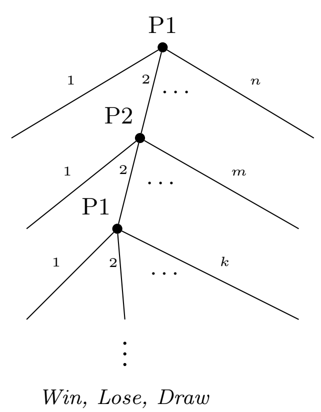

Nodes

Players

Choices at every node

Terminal nodes (outcome)

Information sets

Definition 1.1: Game with Perfect / Imperfect Information

- A game with perfect information is a game where a player at node n knows all the moves that lead to n.
- A game with imperfect information is a game where a player at node n only knows some moves that lead to n.

Example (Matching Pennies).

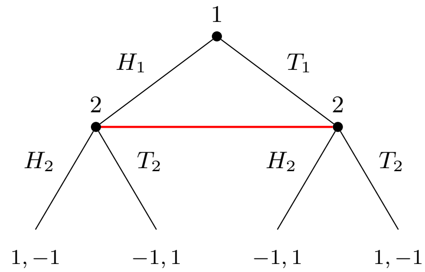

1.2 Information Sets

All the nodes of a game tree are partitioned into information sets.

Rules for Information Sets

- Each node in an information set must belong to the same player.
- Each node in an information set must have the same choices.

Example.

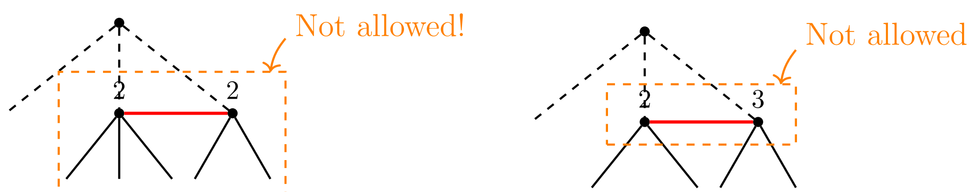

The first tree is invalid because the two nodes in the information set have different choices. The second tree is invalid because the two nodes in the information set belong to different players.

Example (Coin Toss Revisited).

Cointoss Not Revealed to Player 1

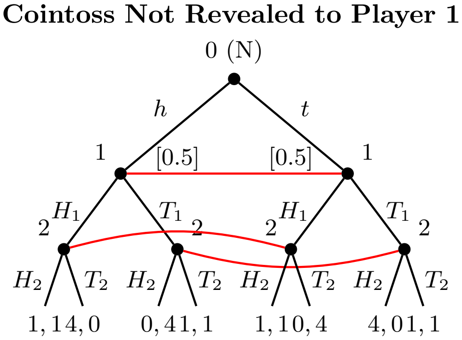

Cointoss Revealed to Player 1

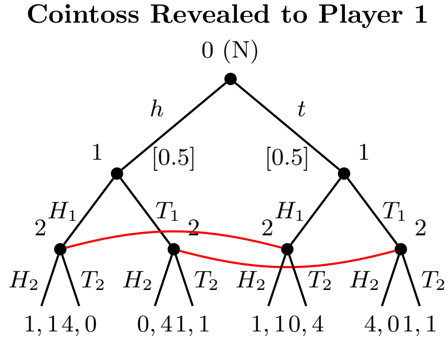

Assumption 1.2

Each path from the root to a terminal node intersects each information set at most once.

Example (Absent-Minded Driver).

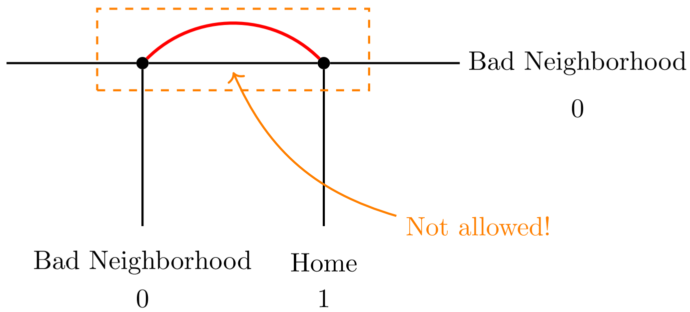

Remark ("No Good Timing").

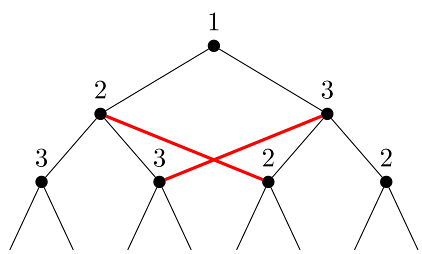

In this game, there is no good notion of "timing" or "stages". So the takeaway is that, the timing of players' moves is not important. What matters is the information structure, i.e., who knows what.

1.3 Strategies

Before introducing the normal form, we need to formalize what a *strategy* is in an extensive-form game. The notion is rich: a strategy specifies what a player plans to do not just on the equilibrium path, but at every information set she might encounter.

Definition 1.3: Strategy

The strategy for player i, denoted by s_i, is a complete plan of which choices to make at every one of her information sets.

Example.

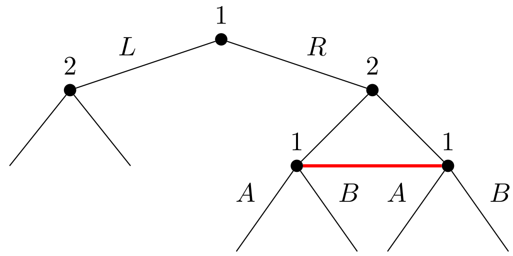

In this game, Player 1's strategies are LA, LB, RA, and RB, where LA means that Player 1 chooses L at the first information set belonging to her, and chooses A at the second. The strategies at all information sets will lead to a unique outcome.

Intuition: A player's strategy is a book with a page for each information set assigned to that player, specifying the choice there.

Definition 1.4: Mixed Strategy

A mixed (or, randomized) strategy for player i, denoted by σ_i, is a probability distribution over S_i, i.e., $\sigma_i \in \Delta(S_i)$.

Pure strategy is a special case of mixed strategy, assigning probability one to a specific action.

1.4 Normal Form

Definition 1.5: Game

A game G of n players consists of the set of strategies S_i and utility functions $u_i : S_1 \times S_2 \times \cdots \times S_n \to \mathbb{R}$ for $i = 1, 2, \cdots, n$:

$$G = (S_1, S_2, \cdots, S_n; u_1, u_2, \cdots, u_n) \equiv (S_i, u_i)_{i=1}^n .$$

Definition 1.6: Belief

A belief of player i, denoted by $\mu_i \in \Delta(S_{-i})$, is a probability distribution over $S_1 \times \cdots \times S_{i-1} \times S_{i+1} \times \cdots \times S_n \equiv S_{-i}$.

Remark.

A belief is a probability distribution over S_{-i}, i.e., $\mu_i \in \Delta_{-i}$. But typically,

$$\Delta(S_i \times S_j) \neq \Delta(S_i) \times \Delta(S_j).$$

Example.

Consider a game of 3 players where Player 2 has strategies s_2 and s_2' to choose from, and Player 3 has s_3 and s_3'. From the Player 1's perspective, she can theoretically form the following two "beliefs":

	s_3	s_3'
s_2	0	1/4
s_2'	1/4	1/2

	s_3	s_3'
s_2	1/9	2/9
s_2'	4/9	4/9

Note that the first belief belongs to $\Delta(S_2 \times S_3)$ but not $\Delta(S_2) \times \Delta(S_3)$, while the second belief belongs to both $\Delta(S_2 \times S_3)$ and $\Delta(S_2) \times \Delta(S_3)$.

Consider the game where Player 1 has strategy space $S_1 = \{s_1, s_1', s_1''\}$, and Player 2 has strategy space $S_2 = \{s_2, s_2'\}$. The normal form of this game is mapping from the strategy profiles to the payoffs of the players. In this example, it can be illustrated by a 3×2 matrix:

	s_2	s_2'
s_1	$\times$	$\times$
s_1'	$\times$	$\times$
s_1''	$\times$	$\times$

Remark (Relationship Between Game Trees and Normal Forms).

- Every tree leads to a unique normal form.
- However, a normal form can lead to multiple game trees that are "equivalent".

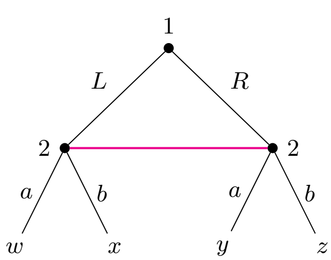

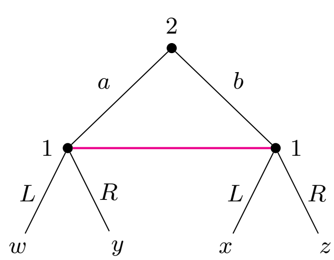

	a	b
l	w	x
r	y	z

- Many game trees can lead to the same normal form.

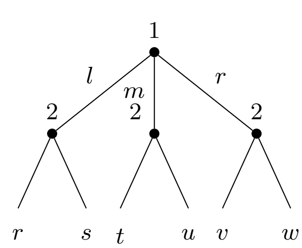

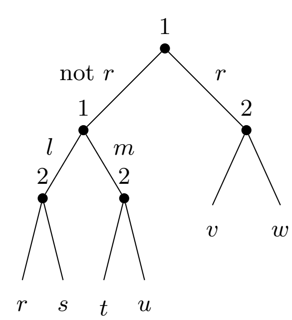

Open Question: Do we lose some important information of the game by looking only at *the normal form?*

1.5 Dominated Strategies and Best Responses

Rather than asking what is an optimal strategy, we approach it in reverse: what are the suboptimal strategies that a rational player would never choose, so that we can rule them out in the analysis of equilibrium?

Definition 1.7: Strictly Dominated Strategy

A pure strategy $s_i \in S_i$ is **strictly dominated** if there exists a (possibly mixed) strategy $\sigma_i \in \Delta(S_i)$ such that

$$u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i}), \quad \forall s_{-i} \in S_{-i},$$

where, by linearity, $u_i(\sigma_i, s_{-i}) = \sum_{s_i' \in S_i} \sigma_i(s_i')\, u_i(s_i', s_{-i})$.

Remark (Three Equivalent Forms).

The same concept can be stated in three equivalent ways. The third form, in terms of beliefs, is the one we will use most often when arguing about rationalizability.

(i) *Pure-strategy form.* There exists $s_i' \in S_i$ with $u_i(s_i', s_{-i}) > u_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}$.

(ii) *Mixed-strategy form.* The form stated in the definition above: there exists $\sigma_i \in \Delta(S_i)$ doing the dominating.

(iii) *Belief form.* There exists $\sigma_i \in \Delta(S_i)$ such that $u_i(\sigma_i, \mu_i) > u_i(s_i, \mu_i)$ for every belief $\mu_i \in \Delta(S_{-i})$.

Form (i) is strictly weaker than (ii)—there are games (e.g. matching pennies with a third strategy) where a pure strategy is not dominated by any other pure strategy but is dominated by a mixture. Forms (ii) and (iii) are equivalent: linearity of u_i in μ_i converts a profile-by-profile inequality into a belief-by-belief one.

Definition 1.8: Best Response

A pure strategy $s_i \in S_i$ is a **best response** to the belief $\mu_i \in \Delta(S_{-i})$ if

$$u_i(s_i, \mu_i) \geq u_i(s_i', \mu_i), \quad \forall s_i' \in S_i,$$

where $u_i(s_i, \mu_i) = \sum_{s_{-i} \in S_{-i}} u_i(s_i, s_{-i})\, \mu_i(s_{-i})$.

Definition 1.9: Never a Best Response (NBR)

A pure strategy s_i is **never a best response** if for every belief $\mu_i \in \Delta(S_{-i})$ there exists $s_i' \in S_i$ with $u_i(s_i', \mu_i) > u_i(s_i, \mu_i)$.

Theorem 1.10: Dominance $\iff$ Never a Best Response

In a finite game, s_i is strictly dominated if and only if s_i is never a best response.

The equivalence is the bridge between the eliminative "dominance" viewpoint and the constructive "rationalizable" viewpoint developed in Chapter **??**. The proof uses a separating-hyperplane argument and is deferred to that chapter.

Definition 1.11: Weakly Dominated Strategy

A pure strategy s_i is **weakly dominated** if there exists $\sigma_i \in \Delta(S_i)$ such that

$$u_i(\sigma_i, s_{-i}) \geq u_i(s_i, s_{-i}) \ \forall s_{-i}, \qquad \text{with strict inequality for some } s_{-i}^* \in S_{-i}.$$

1.6 Iterated Elimination of Dominated Strategies

The previous theorem states that dominated strategy is a belief-free notion. This allows us to develop the concept of iterated elimination of dominated strategies (IESDS), which is a powerful tool to analyze games without having to worry about players' beliefs.

The strategies that survive the IESDS are called *rationalizable strategies*, which are the strategies that can be rationalized by some beliefs about other players' strategies. Rationalizability is a weaker notion than Nash equilibrium, as it does not require mutual consistency of beliefs. However, it is still a useful concept to understand the strategic behavior of players in a game.

Notably, IESDS's outcome is independent of the order of elimination, which is a desirable property. This is because the set of dominated strategies is well-defined and does not depend on the sequence in which they are removed. As a result, we can confidently apply IESDS to simplify games and analyze the strategic interactions without worrying about the order of elimination. This is *not* the case with *IEWDS* (iterated elimination of weakly dominated strategies), where the order of elimination matters.

Example (IEWDS and Order of Elimination).

Consider the following 2-player game and IEWDS.

	L	R
U	$1,1$	$0,0$
M	$1,1$	$2,1$
D	$0,0$	$2,1$

- If Player 1 eliminates U and D, then Player 2 can eliminate nothing, and the process ends. The surviving strategies are M for Player 1, and L and R for Player 2.

- If Player 1 eliminates only U, then Player 2 can eliminate L, and the process ends. The surviving strategies are M and D for Player 1, and R for Player 2.
- If Player 1 eliminates only D, then Player 2 can eliminate R, and the process ends. The surviving strategies are U and M for Player 1, and L for Player 2.

From this example, we can clearly see the "outcome" of IEWDS depends on the order of elimination.

Remark (Chapter Summary).

This chapter introduced two complementary representations of a game and the language we will use throughout the book. The *extensive form* (game tree, with information sets capturing what each player knows when she moves) is the natural object for sequential-move games and games of imperfect information. The *normal form* (a tuple of strategy sets and payoff functions) is the natural object for simultaneous-move analysis and the home of equilibrium concepts; every extensive-form game admits a normal-form representation by enumerating strategies. A *strategy* sits between the two: it is defined on the extensive form (a complete plan for every information set) but its payoff implications live in the normal form. The chapter closes with the first nontrivial solution-concept idea—*dominated strategies and iterated elimination*—which already shows that rationality alone, without any equilibrium fixed-point reasoning, can pin down behavior in some games. Chapter **??** introduces the central equilibrium concept, Nash equilibrium, that the rest of the book builds on.

Part II

Bargaining

Part III

Auctions and Mechanism Design

Part IV

Matching

Part V

Information and Dynamic Games

Part VI

Problem Sets and Solutions

Part VII

Exams and Solutions