Introduction to Game Theory

George Voutsadakis[1]

[1]Mathematics and Computer Science
Lake Superior State University

LSSU Math 500

Subsection 1

Extensive Games With Perfect Information

Introduction to Extensive Games With Perfect Information

- An extensive game is a detailed description of the sequential structure of the decision problems encountered by the players in a strategic situation.
- There is perfect information if each player, when making any decision, is perfectly informed of all the events that have previously occurred.
- We initially restrict attention to games in which:
  - No two players make decisions at the same time;
  - All relevant moves are made by the players (no randomness ever intervenes).
- Later these two restrictions will be removed.

Extensive Games With Perfect Information

Definition (Extensive Game with Perfect Information)

An **extensive game with perfect information** consists of:

- A set N (the set of **players**).
- A set H of sequences (finite or infinite) satisfying:
  - The empty sequence $\emptyset$ is a member of H;
  - If $(a^k)_{k=1,\ldots,K} \in H$ and $L < K$ then $(a^k)_{k=1,\ldots,L} \in H$;
  - If an infinite sequence $(a^k)_{k=1}^{\infty}$ satisfies $(a^k)_{k=1,\ldots,L} \in H$, for every positive integer L, then $(a^k)_{k=1}^{\infty} \in H$.

  Each member of H is a **history**.

  Each component of a history is an **action** taken by a player.

  A history $(a^k)_{k=1,\ldots,K} \in H$ is **terminal** if it is infinite or if there is no a^{K+1} such that $(a^k)_{k=1,\ldots,K+1} \in H$.

  The set of terminal histories is denoted Z.

Extensive Games With Perfect Information (Cont'd)

Definition (Extensive Game with Perfect Information)

- A function $P : (H - Z) \to N$.

  P is the **player function**, $P(h)$ being the player who takes an action after the history h.
- For each player $i \in N$ a preference $\succsim_i$ on Z (the **preference** of i).

- A triple $\langle N, H, P \rangle$, whose components satisfy the first three conditions, is called an **extensive game form with perfect information**.
- If the set H of possible histories is finite, then the game is **finite**.
- If the longest history is finite, then the game has a **finite horizon**.

Interpretation of Extensive Games

- Let h be a history of length k.
- Let (h, a) be the history of length $k + 1$ consisting of h followed by a.
- After any nonterminal history h player $P(h)$ chooses an action from the set $A(h) = \{a : (h, a) \in H\}$.
  - The empty (**initial**) history is the starting point of the game.
  - At this point player $P(\emptyset)$ chooses a member of $A(\emptyset)$.
  - For each possible choice a^0 from this set player $P(a^0)$ subsequently chooses a member of the set $A(a^0)$.
  - This choice determines the next player to move, and so on.
  - A history after which no more choices have to be made is **terminal**.
- We often specify the players' preferences over terminal histories by giving payoff functions that represent the preferences.

Example: Sharing Two Objects

- Two people use the following procedure to share two desirable identical indivisible objects.
  - One of them proposes an allocation;
  - The other either accepts or rejects.
- In the event of rejection, neither person receives either of the objects.
- Each person cares only about the number of objects he obtains.

Example (Formalization)

- An extensive game that models the individuals' predicament is $\langle N, H, P; (\succsim_i)\rangle$ where:
  - $N = \{1, 2\}$;
  - H consists of the ten histories:

$$\begin{array}{l} \emptyset, \\ (2,0),(1,1),(0,2), \\ ((2,0),y),((2,0),n),((1,1),y),((1,1),n),((0,2),y),((0,2),n); \end{array}$$

  - The player function is given by $P(\emptyset) = 1$ and $P(h) = 2$, for every nonterminal history $h \neq \emptyset$;
  - Preferences are determined as follows:

$$\begin{array}{l} ((2,0),y) \succ_1 ((1,1),y) \succ_1 ((0,2),y) \\ \sim_1 ((2,0),n) \sim_1 ((1,1),n) \sim_1 ((0,2),n); \end{array}$$

and

$$\begin{array}{l} ((0,2),y) \succ_2 ((1,1),y) \succ_2 ((2,0),y) \\ \sim_2 ((0,2),n) \sim_2 ((1,1),n) \sim_2 ((2,0),n). \end{array}$$

Example (Tree-Based Representation)

- A convenient representation of this game is

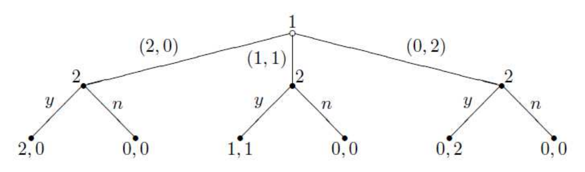

- The top circle represents the initial history $\emptyset$.
- The 1 above indicates that $P(\emptyset) = 1$ (Player 1 makes the first move).
- The three line segments correspond to the three members of $A(\emptyset)$ and are labeled by the names of the actions.
- Each line segment leads to a small disk beside which is the label 2, indicating that Player 2 takes an action after any history of length one.
- The labels beside the line segments that emanate from these disks are the names of Player 2's actions.
- The numbers below the terminal histories are payoffs.

Alternative Tree-Based Formalization

- Based on such a tree

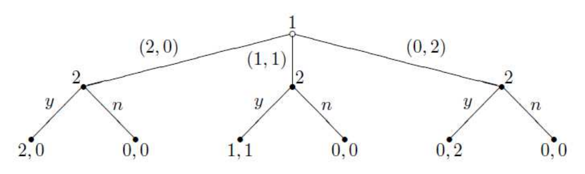

one my give an alternative definition of an extensive game.

- Each node corresponds to a history;
- Any pair of nodes that are connected corresponds to an action;
- The names of the actions are not part of the definition.

Strategies in Extensive Games

- A strategy of a player in an extensive game is a plan that specifies the action chosen by the player, for every history after which it is his turn to move.

Definition (Strategy)

A **strategy** of player $i \in N$ in an extensive game with perfect information $\langle N, H, P, (\succsim_i) \rangle$ is a function that assigns an action in $A(h)$ to each nonterminal history $h \in H - Z$ for which $P(h) = i$.

- The notion of a strategy of a player in a game $\langle N, H, P, (\succsim_i) \rangle$ depends only on the game form $\langle N, H, P \rangle$.

An Example of Strategies

- Consider again the game

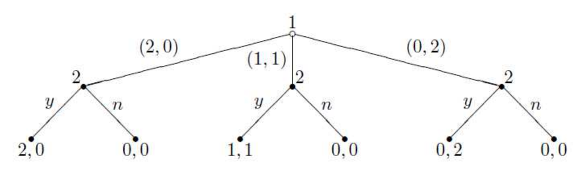

- Player 1 takes an action only after $\emptyset$. So her strategies consist of her possible actions after that history, i.e., $(2, 0)$, $(1, 1)$ and $(0, 2)$.
- Player 2 takes an action after each of the three histories $(2, 0)$, $(1, 1)$ and $(0, 2)$. In each case he has two possible actions. His strategies are triples $a_2b_2c_2$, where a_2, b_2 and c_2 are the actions that he chooses after the histories $(2, 0)$, $(1, 1)$ and $(0, 2)$.
  These are interpreted as contingency plans:
  - If Player 1 chooses $(2, 0)$, then Player 2 will choose a_2;
  - If Player 1 chooses $(1, 1)$, then Player 2 will choose b_2;
  - If Player 1 chooses $(0, 2)$, then Player 2 will choose c_2.

Strategies and Unreachable Histories

- Consider now the game represented by the tree shown in the figure.
- A strategy specifies the action chosen by a player for every history after which it is his turn to move.
- This applies even for histories that, if the strategy is followed, are never reached.

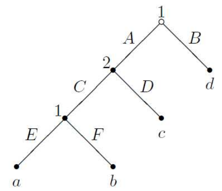

- In this game Player 1 has four strategies AE, AF, BE, and BF.
- Her strategy specifies an action after the history (A, C), even if it specifies that she chooses B at the beginning of the game.
- In this sense, a strategy differs from what we would naturally consider to be a plan of action.
- For some purposes we can regard BE and BF as the same strategy.
- However, in other cases it is important to keep them distinct.

Outcomes and Mixed Strategies

- For each strategy profile $s = (s_i)_{i \in N}$ in $\langle N, H, P, (\succsim_i) \rangle$, we define the **outcome** $O(s)$ of s to be the terminal history that results when each player $i \in N$ follows the precepts of s_i.
- That is, $O(s)$ is the (possibly infinite) history $(a^1, \ldots, a^K) \in Z$, such that, for $0 \leq k < K$, we have

$$s_{P(a^1, \ldots, a^k)}(a^1, \ldots, a^k) = a^{k+1}.$$

- As in a strategic game we can define a **mixed strategy** to be a probability distribution over the set of (pure) strategies.
  - In extensive games with perfect information little is added by considering such strategies.
  - They play a crucial role in the study of extensive games in which the players are not perfectly informed when taking actions.

Nash Equilibria of Extensive Games

- Our first solution concept ignores the sequential structure of the game, treating strategies as choices made before play begins.

Definition (Nash Equilibrium of an Extensive Game)

A **Nash equilibrium of an extensive game with perfect information** $\langle N, H, P, (\succsim_i) \rangle$ is a strategy profile s^* such that, for every player $i \in N$, we have

$$O(s^*_{-i}, s^*_i) \succsim_i O(s^*_{-i}, s_i),$$

for every strategy s_i of player i.

Strategic Form of an Extensive Games

Definition (Strategic Form of an Extensive Game)

The **strategic form of the extensive game with perfect information** $\Gamma = \langle N, H, P, (\succsim_i) \rangle$ is the strategic game $\langle N, (S_i), (\succsim_i) \rangle$ in which, for all $i \in N$:

- S_i is the set of strategies of player i in Γ.
- $\succsim_i'$ is defined, for all $s, s' \in \times_{i \in N} S_i$, by

$$s \succsim_i' s' \quad \text{if and only if} \quad O(s) \succsim_i O(s').$$

- Now we can define a Nash equilibrium of Γ as a Nash equilibrium of the strategic game derived from Γ.

Reduced Strategies

- For Nash equilibria, it suffices to consider only strategies that specify a player's action only after histories that are not inconsistent with the actions that it specifies at earlier points in the game.
- We can define a **reduced strategy** of player i to be a function f_i whose domain is a subset of $\{h \in H : P(h) = i\}$ and satisfies the following conditions:
  (i) It associates with every history h in the domain of f_i an action in $A(h)$;
  (ii) A history h with $P(h) = i$ is in the domain of f_i if and only if all the actions of player i in h are those dictated by f_i.
  That is, if $h = (a^k)$ and $h' = (a^k)_{k=1,\ldots,L}$ is a subsequence of h with $P(h') = i$, then $f_i(h') = a^{L+1}$.

Reduced Strategies and Nash Equilibria

- Each reduced strategy of player i corresponds to a set of strategies of player i.
- For each vector of strategies of the other players, all strategies in the set are **outcome-equivalent**.
- The set of Nash equilibria of an extensive game corresponds to the Nash equilibria of the strategic game in which the set of actions of each player is the set of his reduced strategies.

Example of Reduced Strategies

- As an example of the set of reduced strategies of a player, consider
- Player 1 has three reduced strategies:
  - $f_1(\emptyset) = B$ (with domain $\{\emptyset\}$);
  - $f_1(\emptyset) = A$ and $f_1(A, C) = E$ (with domain $\{\emptyset, (A, C)\}$);
  - $f_1(\emptyset) = A$ and $f_1(A, C) = F$ (with domain $\{\emptyset, (A, C)\}$).

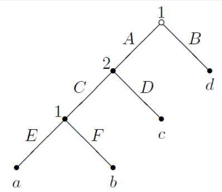

- For some games some of a player's reduced strategies are equivalent.

  Regardless of the strategies of the other players, they generate the same payoffs for all players (though not the same outcome).
- Thus, for some games there is a further redundancy in the definition of a strategy, from the point of view of the players' payoffs.
- E.g., if $a = b$, then player 1's two reduced strategies in which she chooses A at the start of the game are payoff-equivalent.

Reduced Strategic Form of a Game

Definition (Equivalent Strategies and Reduced Strategic Forms)

Let $\Gamma = \langle N, H, P, (\succsim_i)\rangle$ be an extensive game with perfect information and let $\langle N, (S_i), (\succsim'_i)\rangle$ be its strategic form. For any $i \in N$, define the strategies $s_i \in S_i$ and $s'_i \in S_i$ of Player i to be **equivalent** if, for each $s_{-i} \in S_{-i}$, we have

$$(s_{-i}, s_i) \sim'_j (s_{-i}, s'_i), \quad \text{for all } j \in N.$$

The **reduced strategic form of** Γ is the strategic game $\langle N, (S'_i), (\succsim''_i)\rangle$ in which, for each $i \in N$:

- Each set S'_i contains one member of each set of equivalent strategies in S_i;
- $\succsim''_i$ is the preference ordering over $\times_{j \in N} S'_j$ induced by $\succsim'_i$.

Example

- Consider the following game.

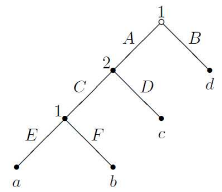

	C	D
AE	a	c
AF	b	c
BE	d	d
BF	d	d

The following is a reduced strategic form of this game.

	C	D
AE	a	c
AF	b	c
B	d	d

Criticism of Nash Equilibria

- The next example illustrates the notion of Nash equilibrium and points to an undesirable feature that equilibria may possess.
- Example: Consider the following game.

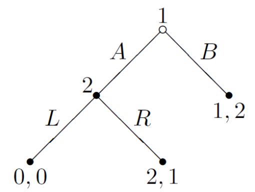

It has two Nash equilibria:

- (A, R), with payoff profile $(2, 1)$;
- (B, L), with payoff profile $(1, 2)$.

For (B, L), we have:

- Given that Player 2 chooses L after the history A, it is optimal for player 1 to choose B at the start of the game.
  If she chooses A, then, given Player 2's choice, she obtains 0 rather than 1;
- Given Player 1's choice of B, it is optimal for Player 2 to choose L (since his choice makes no difference to the outcome).

Interpretation of Equilibrium

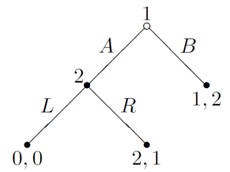

- Our interpretation of a nonterminal history as a point at which a player may reassess his plan of action leads to an argument that the Nash equilibrium (B, L) lacks plausibility.
- The equilibrium (B, L) is sustained by the "threat" of Player 2 to choose L if player 1 chooses A.
- This threat is not credible, since Player 2 has no way of committing himself to this choice.
- Thus, Player 1 can be confident that, if she chooses A, then player 2 will choose R.
- Since she prefers the outcome (A, R) to the Nash equilibrium outcome (B, L), Player 1 has an incentive to deviate from the equilibrium and choose A.

An Additional Example

- Consider again the following game. It has ten Nash equilibria.

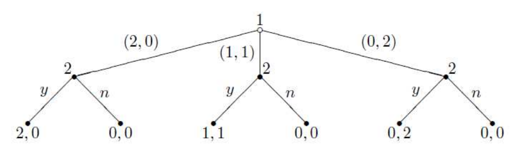

  - The four equilibria

$$((2,0), yyy), ((2,0), yyn), ((2,0); yny), ((2,0), ynn)$$

    result in the division $(2,0)$;
  - The two equilibria

$$((1,1), nyy), ((1,1), nyn)$$

    result in the division $(1,1)$;
  - The equilibrium $((0,2), nny)$ results in the division $(0,2)$;
  - The two equilibria

$$((2,0), nny), ((2,0), nnn)$$

    result in the division $(0,0)$.

An Additional Example (Cont'd)

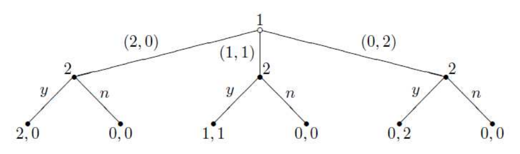

- The only equilibria that do not involve an action of Player 2 that is implausible after some history are

$$((2,0), yyy) \quad \text{and} \quad ((1,1), nyy).$$

- In all other equilibria, Player 2 rejects a proposal that gives him at least one of the objects.
- Like (B, L) in the preceding example, equilibria involving implausible actions are ruled out by the notion of subgame perfect equilibrium.

Subsection 2

Subgame Perfect Equilibrium

Subgames of Extensive Games with Perfect Information

Definition (Subgame)

The **subgame of the extensive game with perfect information** $\Gamma = \langle N, H, P, (\precsim_i) \rangle$ **that follows the history** h is the extensive game $\Gamma(h) = \langle N, H\mid_h, P\mid_h, (\precsim_i\mid_h) \rangle$, where

- $H\mid_h$ is the set of sequences h' of actions for which $(h, h') \in H$;
- $P\mid_h$ is defined by $P\mid_h(h') = P(h, h')$, for each $h' \in H\mid_h$;
- $\precsim_i\mid_h$ is defined by $h' \precsim_i\mid_h h''$ if and only if $(h, h') \precsim_i (h, h'')$.

- In equilibrium, the action prescribed by each player's strategy is optimal, given the other players' strategies, after every history.
- Given a strategy s_i of player i and a history h in the extensive game Γ, denote by $s_i\mid_h$ the strategy that s_i induces in the subgame $\Gamma(h)$, i.e., $s_i\mid_h(h') = s_i(h, h')$, for each $h' \in H\mid_h$.
- Denote by O_h the outcome function of $\Gamma(h)$.

Subgame Perfect Equilibrium

Definition (Subgame Perfect Equilibrium)

A **subgame perfect equilibrium of an extensive game with perfect information** $\Gamma = \langle N, H, P, (\succsim_i) \rangle$ is a strategy profile s^* such that, for every player $i \in N$ and every nonterminal history $h \in H \cap Z$ for which $P(h) = i$, we have $O_h(s^*_{-i}|_h, s^*_i|_h) \succsim_i|_h O_h(s^*_{-i}|_h, s_i)$, for every strategy s_i of player i in the subgame $\Gamma(h)$.

- Equivalently, we can define a subgame perfect equilibrium to be a strategy profile s^* in Γ for which, for any history h, the strategy profile $s^*|_h$ is a Nash equilibrium of the subgame $\Gamma(h)$.
- The notion of subgame perfect equilibrium eliminates Nash equilibria in which the players' threats are not credible.

Examples

- Consider again the game depicted on the left.

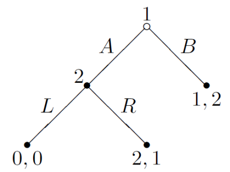

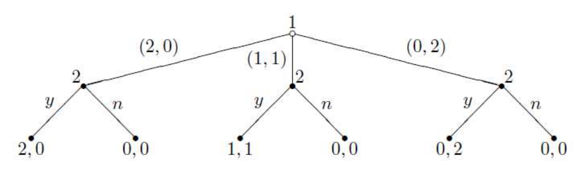

The only subgame perfect equilibrium is (A, R).

- Consider, once more, the game depicted on the right.

The only subgame perfect equilibria are

$$((2,0), yyy) \quad \text{and} \quad ((1,1), nyy).$$

Stackelberg Games

- A **Stackelberg game** is a two-player extensive game with perfect information in which:
  - A "leader" chooses an action from a set A_1;
  - A "follower", informed of the leader's choice, chooses an action from a set A_2.
- The solution usually applied to such games in economics is that of subgame perfect equilibrium.

Solutions of Stackelberg Games

- Some (but not all) subgame perfect equilibria of a Stackelberg game correspond to solutions of the maximization problem

$$\max_{(a_1,a_2)\in A_1\times A_2} u_1(a_1,a_2)$$
$$\text{subject to } a_2 \in \arg\max_{a_2'\in A_2} u_2(a_1,a_2'),$$

  where u_i is a payoff function that represents player i's preferences.
- Under the hypotheses that:
  - The set A_i of actions of each player i is compact;
  - The payoff functions u_i are continuous,

  this maximization problem has a solution.
- There are subgame perfect equilibria of Stackelberg games that do not correspond to a solution of the maximization problem above.

Formalizing a Simplified Criterion

- To verify that a strategy profile s^* is a subgame perfect equilibrium, we must check that, for every player i and every subgame, there is no strategy that leads to an outcome that player i prefers.
- The following result shows that in a game with a finite horizon we can restrict attention, for each player i and each subgame, to alternative strategies that differ from s_i^* in the actions they prescribe after just one history.
- Specifically, a strategy profile is a subgame perfect equilibrium if and only if, for each subgame, the player who makes the first move cannot obtain a better outcome by changing only his initial action.
- To formalize this result, we define, for an extensive game Γ, the **length** of Γ, denoted $\ell(\Gamma)$, to be the length of the longest history in Γ.

The One Deviation Property

Lemma (The One Deviation Property)

Let $\Gamma = \langle N, H, P, (\succsim_i) \rangle$ be a finite horizon extensive game with perfect information. The strategy profile s^* is a subgame perfect equilibrium of Γ if and only if, for every player $i \in N$ and every history $h \in H$ for which $P(h) = i$, we have

$$O_h(s^*_{-i}|_h, s^*_i|_h) \succsim_i|_h O_h(s^*_{-i}|_h, s_i)$$

for every strategy s_i of player i in the subgame $\Gamma(h)$ that differs from $s^*_i|_h$ only in the action it prescribes after the initial history of $\Gamma(h)$.

- A subgame perfect equilibrium s^* of Γ satisfies the condition.
- Now suppose that s^* is not a subgame perfect equilibrium.

  Suppose that player i can deviate profitably in the subgame $\Gamma(h')$.

The One Deviation Property (Cont'd)

- Then there exists a profitably deviant strategy s_i of player i in $\Gamma(h')$, for which $s_i(h) \neq (s_i^* |_{h'})(h)$, for a number of histories h not larger than the length of $\Gamma(h')$.

  Since Γ has a finite horizon, this number is finite.

  From among all the profitable deviations of player i in $\Gamma(h')$, choose a strategy s_i for which the number of histories h, such that $s_i(h) \neq (s_i^* |_{h'})(h)$, is minimal.

  Let h^* be the longest history h of $\Gamma(h')$ for which $s_i(h) \neq (s_i^* |_{h'})(h)$.

  Then the initial history of $\Gamma(h^*)$ is the only history in $\Gamma(h^*)$ at which the action prescribed by s_i differs from that prescribed by $s_i^* |_{h'}$.

  Further, $s_i |_{h^*}$ is a profitable deviation in $\Gamma(h^*)$, since otherwise there would be a profitable deviation in $\Gamma(h')$ that differs from $s_i^* |_{h'}$ after fewer histories than does s_i.

  Thus $s_i |_{h^*}$ is a profitable deviation in $\Gamma(h^*)$ that differs from $s_i^* |_{h^*}$ only in the action that it prescribes after the initial history of $\Gamma(h^*)$.

Introducing Kuhn's Theorem

- We now prove that every finite extensive game with perfect information has a subgame perfect equilibrium.
- The proof is constructive.
  - For each of the longest nonterminal histories in the game:
    - Choose an optimal action for the player whose turn it is to move;
    - Replace each of these histories with a terminal history in which the payoff profile is that which results when the optimal action is chosen;
  - Repeat the procedure, working all the way back to the start of the game.

Kuhn's Theorem

Proposition (Kuhn's Theorem)

Every finite extensive game with perfect information has a subgame perfect equilibrium.

- Consider a finite extensive game with perfect information

$$\Gamma = \langle N, H, P, (\succsim_i) \rangle.$$

Construct a subgame perfect equilibrium of Γ by induction on $\ell(\Gamma(h))$.

At the same time, define a function R that associates a terminal history with every history $h \in H$.

We show that this history is a subgame perfect equilibrium outcome of the subgame $\Gamma(h)$.

Proof of Kuhn's Theorem

- If $\ell(\Gamma(h)) = 0$ (i.e., h is a terminal history of Γ) define $R(h) = h$.

  Now suppose that $R(h)$ is defined for all $h \in H$, with $\ell(\Gamma(h)) \leq k$, for some $k \geq 0$.

  Let h^* be a history for which $\ell(\Gamma(h^*)) = k + 1$ and let $P(h^*) = i$.

  Since $\ell(\Gamma(h^*)) = k + 1$, we have $\ell(\Gamma(h^*, a)) \leq k$, for all $a \in A(h^*)$.

  Define $s_i(h^*)$ to be a $\succsim_i$-maximizer of $R(h^*, a)$, over $a \in A(h^*)$.

  Define $R(h^*) = R(h^*, s_i(h^*))$.

  By induction, we have now defined a strategy profile s in Γ.

  By the One Deviation Property, this strategy profile is a subgame perfect equilibrium of Γ.
- The procedure used in the proof is referred to as **backwards induction**.

Kuhn's Theorem and Chess

- Consider chess under the rule that a game is a draw once a position is repeated three times.
- Under this hypothesis, chess is finite.
- Thus, Kuhn's Theorem implies that it has a subgame perfect equilibrium and, hence, also a Nash equilibrium.
- Chess is strictly competitive.
- So the equilibrium payoff is unique.
- Moreover, any Nash equilibrium strategy of a player guarantees the player his equilibrium payoff.
- Thus, we conclude that one of the following must hold:
  - White has a strategy that guarantees that it wins;
  - Black has a strategy that guarantees that it wins;
  - Each player has a strategy that guarantees that the outcome of the game is either a win for him or a draw.

Subsection 3

Two Extensions of the Definition of a Game

Games with Perfect Information and Chance Moves

- First we extend the model to cover situations in which there is some exogenous uncertainty.
- An **extensive game with perfect information and chance moves** is a tuple

$$\langle N, H, P, f_c, (\succsim_i) \rangle,$$

  where:
  - N is a finite set of **players**;
  - H is a set of **histories**;
  - P is a function from the nonterminal histories in H to $N \cup \{c\}$;
    If $P(h) = c$, then the action after h is determined by chance;
  - For $h \in H$, with $P(h) = c$, $f_c(\cdot \mid h)$ is a probability measure on $A(h)$;
    $f_c(a \mid h)$ is the probability that a occurs after the history h;
    Each $f_c(\cdot \mid h)$ is independent of every other such measure;
  - For each player $i \in N$, $\succsim_i$ is a preference relation on lotteries over the set of terminal histories.

Strategies, Outcomes and Equilibria

- A **strategy** for each player $i \in N$ is defined as before.
- The **outcome** of a strategy profile is a probability distribution over terminal histories.
- The definition of a **subgame perfect equilibrium** is the same.

Games with Perfect Information and Simultaneous Moves

- We model situations in which players move simultaneously after certain histories, each being fully informed of all past events.
- An **extensive game with perfect information and simultaneous moves** is a tuple

$$\langle N, H, P, (\succsim_i) \rangle,$$

where:

  - N, H, and $\succsim_i$, for each $i \in N$ are the same as before;
  - P is a function that assigns to each nonterminal history a set of players;
  - H and P jointly satisfy the condition that, for every nonterminal history h, there is a collection $\{A_i(h)\}$, $i \in P(h)$, of sets for which

$$A(h) = \{a : (h, a) \in H\} = \bigtimes_{i \in P(h)} A_i(h).$$

Histories and Actions

- A **history** in such a game is a sequence of vectors.

  The components of each vector a^k are the actions taken by the players whose turn it is to move after the history $(a^\ell)_{\ell=1}^{k-1}$.
- The set of **actions** among which each player $i \in P(h)$ can choose after the history h is $A_i(h)$.

  The interpretation is that the choices of the players in $P(h)$ are made simultaneously.

Strategies and Equilibria

- A **strategy** of player $i \in N$ in such a game is a function that assigns an action in $A_i(h)$ to every nonterminal history h for which $i \in P(h)$.
- The definition of a **subgame perfect equilibrium** is the same as before, with the exception that $P(h) = i$ is replaced by $i \in P(h)$.
- For an extensive game with perfect information and chance moves:
  - The One Deviation Property still holds;
  - Kuhn's Theorem is still valid.
- On the other hand, for an extensive game with perfect information and simultaneous moves:
  - The One Deviation Property holds;
  - Kuhn's Theorem is no longer valid.

Subsection 4

Two Notable Finite Horizon Games

Games for Subgame Perfect Equilibrium

- We demonstrate some of the strengths and weaknesses of the concept of subgame perfect equilibrium by examining two well known games.
- To describe each of these games, we introduce a variable time that is discrete and starts at period 1.
- This variable is not an addition to the formal model of an extensive game.
- It is merely a device to simplify the description of the games and highlight their structures.

The Chain-Store Game

- A chain-store (player CS) has branches in K cities, $1, \ldots, K$.
- In each city k, there is a single potential competitor, Player k.
- In each period k, Player k decides whether or not to compete with CS.
- If Player k decides to compete, then CS can either fight (F) or cooperate (C).
- CS responds to Player k before Player $k+1$ makes its decision.
- Thus, in period k, the set of possible outcomes is

$$Q = \{\text{Out}, (\text{In}, C), (\text{In}, F)\}.$$

The Chain-Store Game (Individual Rounds)

- If challenged in any given city, the chain-store prefers to cooperate rather than fight.
- However, it obtains the highest payoff if there is no entry.

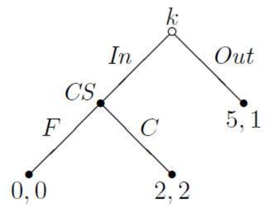

- Each potential competitor is better off staying out than entering and being fought.
- However, it obtains the highest payoff when it enters and the chain-store is cooperative.

The Chain-Store Game (Formalization)

- Two assumptions complete the description of the game.
  - First, at every point in the game, all players know all the actions previously chosen.
    So we may use an extensive game with perfect information:
    - The set of histories is
      $$\left(\bigcup_{k=0}^{K} Q^k\right) \cup \left(\bigcup_{k=0}^{K-1} (Q^k \times \{\mathrm{In}\})\right),$$
      where Q^k is the set of all sequences of k members of Q;
    - The player function is given, for $k = 0, \ldots, K-1$, by
      $$P(h) = \begin{cases} k+1, & \text{if } h \in Q^k \\ \mathrm{CS}, & \text{if } h \in Q^k \times \{\mathrm{In}\} \end{cases}.$$
  - Second, the payoff of the chain-store in the game is the sum of its payoffs in the K cities.

The Chain-Store Game (Equilibria)

- The game has a multitude of Nash equilibria.
  - Every terminal history in which the outcome in any period is either Out or (In, C) is the outcome of a Nash equilibrium.
    Note that, in any equilibrium in which player k chooses Out. the chain-store's strategy specifies that it will fight if player k enters.
  - In contrast, the game has a unique subgame perfect equilibrium.
    In this equilibrium every challenger chooses In and the chain-store always chooses C.
    - In city K, the chain-store must choose C, regardless of the history;
    - So, in city $K - 1$, it must do the same;
    - Continuing the argument, one sees that the chain-store must always choose C.

The Chain-Store Game (Comments)

- For small values of K:
  - The Nash equilibria that are not subgame perfect are intuitively unappealing;
  - The subgame perfect equilibrium is appealing.
- When K is large, the subgame perfect equilibrium loses some of its appeal.

  The strategy of the chain-store in this equilibrium dictates that it cooperate with every entrant, regardless of its past behavior.

Problems with Interpreting Subgame Perfect Equilibria

- Given our interpretation of a strategy, cooperating with every entrant, regardless of its past behavior, means that even a challenger who has observed the chain-store fight with many entrants still believes that the chain-store will cooperate with it.
- Although the chain-store's unique subgame perfect equilibrium strategy does indeed specify that it cooperate with every entrant, it seems more reasonable for a competitor who has observed the chain-store fight repeatedly to believe that its entry will be met with an aggressive response, especially if there are many cities still to be contested.
- If a challenger enters, then it is in the myopic interest of the chain-store to be cooperative.

  Intuition, however, suggests that it may be in its long-term interest to build a reputation for aggressive behavior, in order to deter future entry.

The Centipede Game

- Two players are involved in a process that they alternately have the opportunity to stop.
  - Each prefers the outcome when he stops the process in any period t to that in which the other player does so in period $t + 1$.
  - However, better still is any outcome that can result if the process is not stopped in either of these periods.

  After T periods, where T is even, the process ends.
- For T = 6 the game is

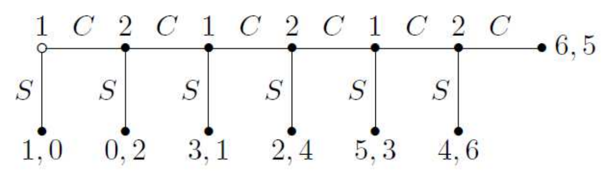

Formal Description of the Centipede Game

- Formally, the set of histories in the game consists of:
  - All sequences $C(t) = (C, \ldots, C)$ of length t, for $0 \leq t \leq T$;
  - All sequences $S(t) = (C, \ldots, C, S)$ consisting of $t-1$ repetitions of C followed by a single S, for $1 \leq t \leq T$.
- The player function is defined by

$$P(C(t)) = \begin{cases} 1, & \text{if } t \text{ is even and } t \leq T-2 \\ 2, & \text{if } t \text{ is odd} \end{cases} .$$

- Preferences are specified by the following clauses:
  - Player $P(C(t))$ prefers $S(t+3)$ to $S(t+1)$ to $S(t+2)$ for $t \leq T-3$;
  - Player 1 prefers $C(T)$ to $S(T-1)$ to $S(T)$;
  - Player 2 prefers $S(T)$ to $C(T)$.

Subgame Perfect Equilibrium

- The game has the unique subgame perfect equilibrium in which each player chooses S in every period. The outcome of this equilibrium is the same as the outcome of every Nash equilibrium.

  First note that there is no equilibrium in which the outcome is $C(T)$.

  Now assume that there is a Nash equilibrium that ends with Player i choosing S in period t (i.e., after the history $C(t-1)$).

  If $t \geq 2$, then Player j can increase his payoff by choosing S in $t-1$.

  Hence, in any equilibrium, Player 1 chooses S in the first period.

  For this to be optimal for Player 1, Player 2 must choose S in period 2.
- The notion of Nash equilibrium imposes no restriction on the players' choices in later periods.

  Any pair of strategies in which Player 1 chooses S in period 1 and Player 2 chooses S in period 2 is a Nash equilibrium.

Interpretation of Subgame Perfect Equilibria

- In the unique subgame perfect equilibrium of this game each player believes that the other player will stop the game at the next opportunity, even after a history in which that player has chosen to continue many times in the past.
- As is the case with the chain-store game, such a belief is not intuitively appealing.
- Unless T is very small it seems unlikely that player 1 would immediately choose S at the start of the game.
- The intuition in the centipede game is slightly different from that in the chain-store game.
- After any long history, both players have repeatedly violated the precepts of rationality enshrined in the notion of subgame perfect equilibrium.

Subsection 5

Iterated Elimination of Weakly Dominated Strategies

Iterated Elimination and Backwards Induction

- We defined the procedure of iterated elimination of weakly dominated actions for a strategic game.
  - It is less appealing than the procedure of iterated elimination of strictly dominated actions (since a weakly dominated action is a best response to some belief).
  - Still, it is a natural method for a player to use to simplify a game.
- In the proof of Kuhn's Theorem:
  - We define the procedure of backwards induction for finite extensive games with perfect information;
  - We show that it yields the set of subgame perfect equilibria of the game.

Iterated Elimination and Subgame Perfect Equilibrium

- The procedures of Iterated Elimination and Backwards Induction are related.
- Let Γ be a finite extensive game with perfect information in which no player is indifferent between any two terminal histories.
- Γ has a unique subgame perfect equilibrium.
- We define a sequence for eliminating weakly dominated actions in the strategic form G of Γ.
- All the surviving action profiles of G generate the unique subgame perfect equilibrium outcome of Γ.

The Elimination Process

- Let h be a history of Γ with $P(h) = i$ and $\ell(\Gamma(h)) = 1$.
- Let $a_i^* \in A(h)$ be the unique action selected by the procedure of backwards induction for the history h.
- Backwards induction eliminates every strategy of Player i that chooses an action different from a_i^* after the history h.
- Among these strategies, those consistent with h are weakly dominated actions in G.
- These weakly dominated actions are eliminated from G at this stage.

The Elimination Process (Cont'd)

- After elimination for each history h with $\ell(\Gamma(h)) = 1$, we perform elimination for histories h with $\ell(\Gamma(h)) = 2$.
- We continue back to the beginning of the game in this way.
- Every strategy of Player i that remains chooses the action after any history that is consistent with Player i's subgame perfect equilibrium strategy.
- The subgame perfect equilibrium remains.
- Moreover, the strategy profiles that remain generate the unique subgame perfect equilibrium outcome.

Remarks I

- Other orders of elimination may remove all subgame perfect equilibria.
- Consider the game

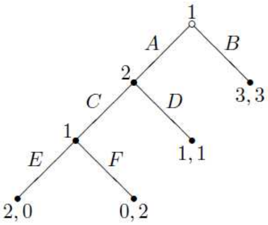

	C	D
AE	$2, 0$	$1, 1$
AF	$0, 2$	$1, 1$
BE	$3, 3$	$3, 3$
BF	$3, 3$	$3, 3$

- The unique subgame perfect equilibrium is (BE, D).
- If, in the strategic form, the weakly dominated action AE is eliminated, then D is weakly dominated in the remaining game.
- If AF is eliminated after D, then neither of the two remaining action profiles (BE, C) and (BF, C) are subgame perfect equilibria of the extensive game.

Remarks II

- If some player is indifferent between two terminal histories, then:
  (i) There may be an order of elimination that eliminates a subgame perfect equilibrium outcome.
  (ii) There may exist no order of elimination for which all surviving strategy profiles generate subgame perfect equilibrium outcomes.

Remarks II (Example)

- Consider the game

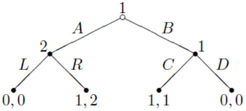

	L	R
AC	0, 0	1, 2
AD	0, 0	1, 2
BC	1, 1	1, 1
BD	0, 0	0, 0

- The strategies AC, AD, and BD of Player 1 are all weakly dominated by BC.
- After they are eliminated, no remaining pair of actions yields the subgame perfect equilibrium outcome (A, R).
- Suppose the payoff $(1, 2)$ is replaced by $(2, 0)$.

  Then, in the modified game, the outcome (A, L), which is not even a Nash equilibrium outcome, survives any order of elimination.

BoS with an Outside Option

- Consider the extensive game with perfect information and simultaneous moves

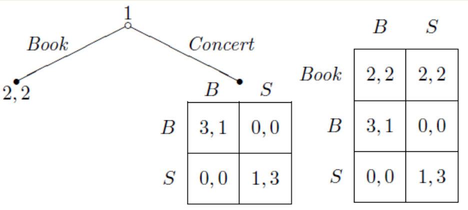

- Player 1 first decides whether to stay at home and read a book or to go to a concert.
    - If she decides to read a book, then the game ends.
    - If she decides to go, she is engaged in the game BoS with player 2. After the history Concert, the players choose actions simultaneously.

BoS with an Outside Option (Preferences)

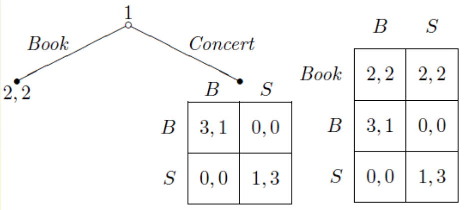

- Each player prefers to hear the music of his favorite composer in the company of the other player rather than either go to a concert alone or stay at home.
- However, each player prefers to stay at home rather than either go out alone or hear the music of his less-preferred composer.

BoS with an Outside Option (Elimination)

- In the reduced strategic form of the game

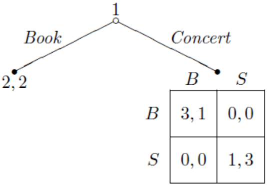

	B	S
Book	2, 2	2, 2
B	3, 1	0, 0
S	0, 0	1, 3

  - S is strictly dominated for player 1 by Book.
  - If it is eliminated, then S is weakly dominated for player 2 by B.
  - Finally, Book is strictly dominated by B for player 1.
  - The outcome that remains is (B, B).
- This sequence of eliminations corresponds to forward induction:
  - If Player 2 has to decide, he knows that Player 1 has not chosen Book.
  - Such a choice makes sense for Player 1 only if she plans to choose B.
  - Thus, Player 2 should choose B also.

Burning Money

- Consider a version of BoS with a twist.

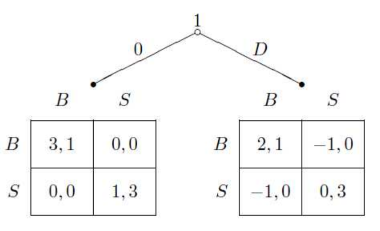

- At the start, Player 1 has two options.
  - Discard a dollar (D);
  - Refrain from doing so (0).

  Her move is observed by Player 2.
- Then, they play BoS with payoffs as in the left-hand table.
- Both players are risk neutral (the two subgames that follow Player 1's initial move are strategically identical).

Burning Money (Elimination)

- The reduced strategic form of the game is also shown

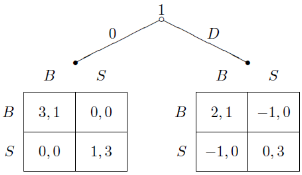

	BB	BS	SB	SS
$0B$	$3,1$	$3,1$	$0,0$	$0,0$
$0S$	$0,0$	$0,0$	$1,3$	$1,3$
DB	$2,1$	$-1,0$	$2,1$	$-1,0$
DS	$-1,0$	$0,3$	$-1,0$	$0,3$

- Weakly dominated actions can be eliminated:
  1. DS is weakly dominated for Player 1 by $0B$;
  2. SS is weakly dominated for Player 2 by SB;
  3. BS is weakly dominated for Player 2 by BB;
  4. $0S$ is strictly dominated for Player 1 by DB;
  5. SB is weakly dominated for Player 2 by BB;
  6. DB is strictly dominated for Player 1 by $0B$.

  The single strategy pair that remains is $(0B, BB)$.
- So, under elimination, this outcome is Player 1's favorite.

Burning Money (Argument)

- An intuitive supporting argument is the following:
  - Player 1 must anticipate that, if she chooses 0, then she will obtain an expected payoff of at least $\frac{3}{4}$, since, for every belief about the behavior of Player 2, she has an action that yields her at least this expected payoff.
  - Thus, if Player 2 observes that Player 1 chooses D, then he must expect that Player 1 will subsequently choose B (since the choice of S cannot possibly yield Player 1 a payoff in excess of $\frac{3}{4}$).
  - Given this, Player 2 should choose B, if Player 1 chooses D.
  - Player 1, knowing this, expects to obtain a payoff of 2 by choosing D.
  - But now Player 2 can rationalize the choice 0 by Player 1 only by believing that Player 1 will choose B (since S can yield Player 1 no more than 1).
  - So the best action of Player 2 after observing 0 is B.
  - This makes 0 the best action for Player 1.