Lecture 15: Extensive Games and Subgame Perfect Equilibrium
Readings: Osborne and Rubinstein Ch 6.1-6.5

Extensive Games with Perfect Information

An *extensive game* is an explicit description of the sequential structure of the decision problems encountered by the players in a strategic situation. The model allows us to study solutions in which each player is assumed to consider his plan of action not only at the beginning of the game but also at any point of time at which he has to make a decision. By contrast, in strategic game we could talk about a plan covering unlimited contingencies, but the timing structure is "lost" and the model does not allow us to talk about a player reconsidering his strategy after some events in the game have unfolded.

A game is with *perfect information* if each player, when making any decision, is perfectly informed of all the events that have previously occurred.

Definition An extensive game with perfect information is :
▲A set N .
▲A set H of sequences (finite or infinite) that satisfies the following three properties.
$\varnothing \in H$.
If $(a^k)_{k=1,\ldots,K} \in H$ (where K may be infinite) 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 $(a^k)_{k=1}^{\infty} \in H$.
(A member of H is a *history*; each component of a history is an action taken by some 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.
▲A function P that assigns a player to every nonterminal history (P is the *player function)*
▲$\forall i \in N$ a preference relation $\succsim_i$ on Z

We interpret an extensive game as follows: After any nonterminal history h player $P(h)$ chooses an action from the set $A(h) = \{a\colon (h,a) \in H\}$. The empty (initial) history is the starting point of the game. $P(\varnothing)$ chooses a member of $A(\varnothing)$. For each $a^0 \in A(\varnothing)$ 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 until a terminal history is reached.

Example Two people want to share two apples. One of them proposes an allocation, which the other then either accepts or rejects. In the event of rejection, neither person receives an apple. It is assumed that each person cares only about the number of apples he obtains.

An extensive game that models the scenario:
▲$N = \{1,2\}$;
▲$H = \{\varnothing, (2,0), (1,1), (0,2), ((2,0),y), ((2,0),n), ((1,1),y), ((1,1),n), ((0,2),y), ((0,2),n)\}$
▲$P(\varnothing) = 1$ and $P(h) = 2$ for every nonterminal history $h \neq \varnothing$.
▲$((2,0),y) \succ_1 ((1,1),y) \succ_1 ((0,2),y) \sim_1 ((2,0),n) \sim_1 ((1,1),n) \sim_1 ((0,2),n)$ and $((0,2),y) \succ_2 ((1,1),y) \succ_2 ((2,0),y) \sim_2 ((0,2),n) \sim_2 ((1,1),n) \sim_2 ((2,0),n)$.

A Discussion of representation of this game as a tree.

Strategies

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

To illustrate the notion of a strategy consider the game

1–*A*–2–*C*–1–*E*–*a*
|*B* |*D* |*F*
d *c* *b*

Assume that $b \succ_1 a \succ_1 d \succ_1 c$ and $a \succ_2 c \succ_2 b \succ_2 d$

A strategy specifies the action chosen by a player for *every* history after which it is his turn to move, *even for histories that, if the strategy is followed, are not reached*. Player 1 has *four* strategies *AE*, *AF*, *BE*, and *BF*. Thus, his strategy specifies an action after the history (A, C) even if it specifies that he 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 each strategy profile $s = (s_i)_{i \in N}$ 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 $s_{P(a^1,\ldots,a^k)}(a^1, \ldots, a^k) = a^{k+1}$.

Extension: Exogenous Uncertainty

A simple extension of the model is to *extensive game with perfect information and chance moves:* Here some of the moves are controlled by nature. The model is $\langle N, H, P, f_c, (\succsim_i)\rangle$ where, as before, N is a finite set of players, H is a set of histories, and $P : H \to N \cup \{c\}$. ($P(h) = c$ means that *chance* determines the action after h.) In addition we require that for each $h \in H$ with $P(h) = c$, $f_c(\cdot|h)$ is a probability measure on $A(h)$. The preferences are taken to be defined on the set of lotteries with prizes being the set of terminal histories.

Nash Equilibrium

Definition A Nash equilibrium of $\langle N, H, P, (\succsim_i)\rangle$ is a strategy profile s^* such that for every player $i \in N$ we have $O(s^*_{-i}, s^*_i) \succsim_i O(s^*_{-i}, s_i)\ \ \forall s_i$.

Thus, a Nash equilibrium of an extensive game Γ is actually a Nash equilibrium of the following strategic game.

Definition The strategic form of $\Gamma = \langle N, H, P, (\succsim_i)\rangle$ is the strategic game $\langle N, (S_i), (\succsim'_i)\rangle$ in which for each player $i \in N$

▲S_i is the set of strategies of player i in Γ

▲ $\succsim'_i$ is defined by $s \succsim'_i r$ if and only if $O(s) \succsim_i O(r)$.

Example

The Nash equilibria of the above game are the pair of strategies

$((2,0), yyy)$, $((2,0), yyn)$, $((2,0), yny)$, $((2,0), ynn)$ results in the division $(2,0)$;
$((1,1), nyy)$, $((1,1), nyn)$ result in the division $(1,1)$,
$((0,2), nny)$, results in the division $(0,2)$ and

$((2,0), nny)$, $((2,0), nnn)$. result no agreement

Note that if Nash equilibrium were the only solution we would be interested in for extensive games, we could make do with a definition of a strategy so that it specifies a player's action only after histories that are not inconsistent with the actions that it specifies at earlier points in the game. This is so because the outcome $O(s)$ of the strategy profile s is not affected by the actions that the strategy s_i of any player i specifies after contingencies that are inconsistent with s_i.

The Centipede Game

Two players are involved in a process that they alternately have the opportunity to stop. Each prefers the outcome "he stops the process in t" to that in which the other player does so in $t+1$. 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.

Formally,

▲H consists of all sequences $C(t) = (C, \ldots, C)$ of length t, for $0 \le t \le T$, and all sequences $S(t) = (C, \ldots, C, S)$ consisting of $t-1$ repetitions of C followed by a single S, for $1 \le t \le T$.
▲$P(C(t)) = 1$ if t is even and $t \le T-2$ and $P(C(t)) = 2$ if t is odd.
▲ $P(C(t))$ prefers $S(t+3)$ to $S(t+1)$ to $S(t+2)$ for $t \le T-3$, player 1 prefers $C(T)$ to $S(T-1)$ to $S(T)$, and player 2 prefers $S(T)$ to $C(T)$.

The only outcome of any Nash equilibrium is $S(1)$. To see this, note that there is no equilibrium in which the outcome is $C(T)$. Now assume that there is a NE that ends with player i choosing S in period t (i.e. after the history $C(t-1)$). If $t \ge 2$ then player j can profit by choosing S in period $t-1$. Hence in any equilibrium player 1 chooses S in the first period. In order 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 NE.

Subgame Perfect Equilibrium

Definition A *subgame* of $\Gamma = \langle N, H, P, (\succsim_i) \rangle$ that follows h is the extensive game $\Gamma(h) = \langle N, H|_h, P|_h, (\succsim_{i,h}) \rangle$, where $H|_h$ is the set of sequences h' of actions for which $(h,h') \in H$, $P|_h$ is defined by $P|_h(h') = P(h,h')$ for each $h' \in H|_h$, and $\succsim_{i,h}$ is defined by $h' \succsim_{i,h} h''$ iff $(h,h') \succsim_i (h,h'')$.

Given a strategy s_i of player i and a history h in the extensive game Γ, denote by $s_i|_h$ the strategy that s_i induces in the subgame $\Gamma(h)$ (i.e. $s_i|_h(h') = s_i(h,h')$ for each $h' \in H|_h$); denote by O_h the outcome function of $\Gamma(h)$.

Definition 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 $h \in H \setminus Z$ 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 SPE 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. For example., in the "the split of two apples" game there are two subgame perfect equilibria $((2,0), yyy)$ and $((1,1), nyy)$) that are not equivalent in terms of either player's preferences.

The centipede game has a unique SPE; in this equilibrium each player chooses S in every period. 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.

Kuhn's Theorem (Backward Induction)

To verify that a strategy profile s^* is a subgame perfect equilibrium we have to 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* the initial history. A strategy profile is an SPE iff for each subgame the player who makes the first move cannot obtain a better outcome by changing only his initial action. Denote by $\ell(\Gamma)$ the length of the longest history in Γ.

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 SPE of Γ iff $\forall i \in N$ and $\forall h \in H$ 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 $\Gamma(h)$.

Proof If s^* is an SPE of Γ then it satisfies the condition. Now suppose that s^* is not a SPE; suppose that player i can deviate profitably in the subgame $\Gamma(h')$. Then there exists a profitable 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 $\Gamma(h', h^*)$ 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 $\Gamma(h', h^*)$, 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 $s_i|_{h',h^*}$ is a profitable deviation in $\Gamma(h', h^*)$ that differs from $s_i^*|_{h',h^*}$ only in the action that it prescribes after the initial history of $\Gamma(h', h^*)$.

Proposition Every finite extensive game with perfect information has a SPE.

Proof Let $\Gamma = \langle N, H, P, (\succsim_i)\rangle$ be a finite extensive game with perfect information. We construct an SPE by induction on $\ell(\Gamma(h))$; at the same time we define a function R that associates a terminal history with every history $h \in H$ and show that this history is a SPE outcome of the subgame $\Gamma(h)$.

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 $a \in A(h^*)$, and define $R(h^*) = R(h^*, s_i(h^*))$. By induction we have now defined a strategy profile s in Γ; by Lemma this strategy profile is a SPE of Γ.

The procedure used in this proof is referred to as *backwards induction*. In addition to being a means for proving the proposition, the procedure is an algorithm for calculating the set of SPE of a finite game. Part of the appeal of the notion of SPE derives from the fact that the algorithm describes what appears to be a natural way for players to analyze such a game so long as the horizon is relatively short.

Conclusion: Under the rule that Chess ends with a draw once a position is repeated three times, chess is finite, so that Kuhn's theorem implies that it has a SPE and hence also a NE. Because chess is strictly competitive, the equilibrium payoff is unique and thus any NE strategy of a player guarantees the player his equilibrium payoff. Thus either White has a strategy that guarantees that it wins, or Black has a strategy that guarantees that it wins, or each player has a strategy that guarantees that the outcome of the game is either a win for it or a draw. Quite a remarkable conclusion!

The Interpretation of a Strategy

The definition of a strategy does not correspond to a plan of action since it requires a player to specify his actions after histories that are impossible if he carries out his plan.

One interpretation for the components of a player's strategy corresponding to histories that are not possible if the strategy is followed is that they are the *beliefs* of the other players about what the player will do in the event he does not follow his plan. For example, in the Centipede game, the plan of player 1 for the third period, if he plans to stop right away is one belief of player 2 about what player 1 would do had he not followed his strategy.

Note that when we interpret a strategy as a belief it becomes problematic to speak of the "choice of a strategy". Usually we do not think about beliefs as an outcome of choice. Note also that in any equilibrium of a game with more than two players there is an implicit assumption that all the players other than any given player i hold the *same* beliefs about player i's behavior, not only if he follows his plan of action but also if he deviates from this plan. This interpretation requires carefulness in imposing constraints on the strategies (when we do it) since one is then making assumptions not only about the players' plans of action, but also about their beliefs regarding other players' intentions.

The Chain-Store Game

A chain-store (player CS) has branches in K cities, numbered $1,\ldots,K$. In each city k there is a single potential competitor, player k. In period k player k decides whether or not to compete with CS. If player k decides to compete then the chain-store can either fight (F) or cooperate (C). The chain-store responds to player k's decision before player $k+1$ makes its decision. Thus in each period k the set of possible outcomes is $Q = \{Out, (In, C), (In, F)\}$. If challenged in any given city the chain-store prefers to cooperate (payoff 2) rather than fight (payoff 0) , but obtains the highest payoff if there is no entry (payoff 5). Each potential competitor is better off staying out than entering and being fought, but obtains the highest payoff when it enters and the chain-store is cooperative.

Two assumptions complete the description of the game. First, at every point in the game all players know all the actions previously chosen. This allows us to model the situation as an extensive game with perfect information, in which the set of histories is $(\cup_{k=0}^{K} Q^k) \cup (\cup_{k=0}^{K-1} (Q^k \times \{In\}))$, where Q^k is the set of all sequences of k members of Q, and the player function is given by $P(h) = k+1$ if $h \in Q^k$ and $P(h) = CS$ if $h \in Q^k \times \{In\}$, for $k = 0,\ldots,K-1$. Second, the payoff of the chain-store in the game is the sum of its payoffs in the

K cities and each player k cares only about the outcome in his city.

The game has a multitude of Nash equilibria: In particular every terminal history in which the outcome in any period is either *Out* or (In, C) is the outcome of a NE.
In contrast, the game has a unique SPE; 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 that in city $K - 1$ it must do the same; continuing the argument one sees that the chain-store must always choose C.)

For small values of K the NE that are not SPE are intuitively unappealing while the SPE is appealing. However, 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. Given our interpretation of a strategy as a belief this 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. 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, but intuition suggests that it may be in its long-term interest to build a reputation for aggressive behavior, in order to deter future entry. The current model does not allow such build up of reputation.

Problem set 15

1. **(Exercise)** Given an extensive game define a *reduced strategy* of player i in that game to be a function f_i whose domain is a subset of $\{h \in H: P(h) = i\}$ and has the following properties:

(i) it associates with every h in the domain of f_i an action in $A(h)$ and

(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 $f_i(h') = a^{L+1}$).

▲ Express formally the sense in which the set of Nash equilibria of an extensive game with perfect information 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.

2. **(Exercise)** A Stackelberg game is a two-player extensive game with perfect information in which a "leader" chooses an action from a set A_1 and a "follower", informed of the leader's choice, chooses an action from a set A_2. Some SPE of a Stackelberg game correspond to solutions of the maximization problem

$\max_{(a_1,a_2)\in A_1\times A_2} u_1(a_1,a_2)$ *subject to* $a_2 \in \arg\max_{a_2'\in A_2} u_2(a_1,a_2')$, where u_i is i's payoff function.

Show that the solution for the problem induces an SPE and give an example of a SPE of a Stackelberg game that does not correspond to a solution of the maximization problem above.

3. **(Exercise)** Give an example of an infinite horizon game for which the one deviation property does not hold.

4 **(Exercise)** Show that the requirement in Kuhn's theorem that the game be finite cannot be replaced by the requirement that it have a finite horizon, nor by the requirement that after any history each player have finitely many possible actions.

5. **(Exercise)** Say that a finite extensive game with perfect information satisfies the *no indifference condition* if $z \sim_j z' \ \forall j \in N$ whenever $z \sim_i z'$ for some $i \in N$, where z and z' are terminal histories. Show, using induction on the length of subgames, that every player is indifferent among all subgame perfect equilibrium outcomes of such a game. Show also that if s and s' are subgame perfect equilibria then so is s'', where for each player i the strategy s_i'' is equal to either s_i or s_i' (that is, the equilibria of the game are *interchangeable*).

6. **(Exercise)** Armies 1 and 2 are fighting over an island initially held by a battalion of army 2. Army 1 has K battalions and army 2 has L. Whenever the island is occupied by one army the opposing army can launch an attack. The outcome of the attack is that the occupying battalion and one of the attacking battalions are destroyed; the attacking army wins and, so long as it has battalions left, occupies the island with one battalion. The commander of each army is interested in maximizing the number of surviving battalions but also regards the occupation of the island as worth more than one battalion but less than two. (If, after an attack, neither army has any battalions left, then the payoff of each commander is 0.)

Analyze this situation as an extensive game and, using the notion of SPE, predict the winner as a function of K and L.

7. **(More difficult Exercise)** Under the official rules of Chess, a game ends when a position is repeated three times **and** the player who has to move declares a "draw". Thus, Chess is actually not a finite game. Prove that Chess (with this additional detail) still has a "value".

8. (**Exercise)** Show that both the one deviation property and Kuhn's theorem hold for an extensive game with perfect information and chance moves (Read section 6.3.1 in OR).

9. (**Exercise)** Consider the following two-player game. First player 1 can choose either *Stop* or *Continue*. If she chooses *Stop* then the game ends with the pair of payoffs $(1, 1)$. If she chooses *Continue* then the players simultaneously announce nonnegative integers and each player's payoff is the product of the numbers. Formulate this situation as an extensive game with simultaneous moves (read section 6.3.2 in OR) and find its subgame perfect equilibria.