Extensive Games with Perfect Information

There is perfect information if each player making a move observes all events that have previously occurred.

Start by restricting attention to games without simultaneous moves and without nature (no randomness).

Definition 89.1: An **extensive game with perfect information** consists of the following components:

1. The set of players, N.

2. A set, H, of sequences (histories of actions) satisfying the following properties:

(2a) The empty sequence, $\varnothing$, is an element of H.

(2b) If $(a^k)_{k=1,\ldots,K} \in H$, (where K may be infinity) and $L < K$, then $(a^k)_{k=1,\ldots,L} \in H$.

(2c) If an infinite sequence, $(a^k)_{k=1,\ldots,\infty}$ satisfies

$(a^k)_{k=1,\ldots,L} \in H$ for every positive integer L, then $(a^k)_{k=1,\ldots,\infty} \in H$.

3. A function, P, that assigns to each non-terminal history a member of N. (P is the player function, with $P(h)$ being the player who takes an action after the history, h. A history, $(a^k)_{k=1,\ldots,K}$, is **terminal** if it is infinite or if there is no a^{K+1} such that $(a^k)_{k=1,\ldots,K+1} \in H$.)

4. For each player, $i \in N$, a preference relation $\succeq_i$on the set of terminal histories, Z.

If $< N, H, P >$ satisfies (1)-(3), but preferences are not specified, it is called an extensive game **form** with perfect information. (Example: auction rules are specified, but not preferences over the objects.)

If the set of histories is finite, the game is finite. If the longest history is finite, the game has a finite horizon.

If h is a history of length k, then (h, a) is the history of length $k + 1$ consisting of h followed by the action a.

After any non-terminal history, h, the player $P(h)$ chooses an action from the set $A(h) = \{a : (h, a) \in H\}$.

Note: we could equivalently define an extensive game as a game tree (a connected graph with no cycles) Each node corresponds to a history, and the connection between two nodes corresponds to an action.

A strategy is a plan specifying the action a player takes for every history after which it is his/her turn to move. Note: some simple games have many, many strategies.

Definition 92.1: A **strategy** of player $i \in N$ in an extensive game with perfect information $\langle N, H, P, (\succeq_i) \rangle$ is a function that assigns an action in $A(h)$ to every nonterminal history, $h \in H \backslash Z$ for which we have $P(h) = i$.

Note: The definition of a strategy requires us to specify an action after histories that are impossible to reach, if the strategy is followed. One could argue that a plan does not have to specify such contingencies. One interpretation is that this part of the strategy represents the beliefs that other players have about what the player would do if he/she did not follow the plan.

For each strategy profile, $s = (s_i)_{i \in N}$, we define the **outcome** of s (denoted $O(s)$) to be the terminal node resulting when each player i chooses actions according to s_i.

Definition 93.1: A **Nash equilibrium** of an extensive game with perfect information is a strategy profile s^* such that for every player i we have $O(s^*_{-i}, s^*_i) \succeq_i O(s^*_{-i}, s_i)$ for all s_i.

For every extensive game, there is a corresponding strategic game.

Definition 94.1: The strategic form of the extensive game with perfect information, $\Gamma = \langle N, H, P, (\succeq_i) \rangle$ is the strategic game $G =< N, (S_i), (\succeq_i') >$, where S_i is player i's set of strategies in Γ and set of actions in G, and $\succeq_i'$ is defined by the condition, $s \succeq_i' s'$ if and only if $O(s) \succeq_i O(s')$ for all profiles s and s'.

By definition, for every extensive game Γ and corresponding strategic game G, s^* is a Nash equilibrium of Γ if and only if s^* is a Nash equilibrium of G.

A given strategic game G corresponds to many extensive games, but all of these extensive games have the same set of Nash equilibria. (Not true for other concepts like subgame perfection.)

Subgame Perfection

Some Nash equilibria possess the undesirable property that a player would not credibly follow through with their strategy in all circumstances.

		player 2 L	R
player 1	A	$1, 2$	$1, 2$
	B	$0, 0$	$2, 1$

This game has two Nash equilibria: (A,L) and (B,R).

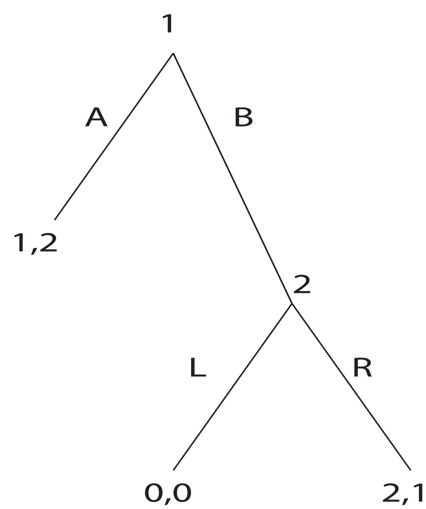

Figure 1:

Will player 2 really play L if called upon to move?

Definition 97.1: The **subgame** of the extensive game with perfect information, $\Gamma = \langle N, H, P, (\succeq_i)\rangle$, that follows the history h is the extensive game,

$$\Gamma(h) = \langle N, H\mid_h, P\mid_h, (\succeq_i\mid_h)\rangle,$$

where $H\mid_h$ is the set of sequences, h', for which we have $(h, h') \in H$; $P\mid_h$ is defined by $P\mid_h(h') = P(h, h')$; and $\succeq_i\mid_h$ is defined by $h' \succeq_i\mid_h h''$ if and only if we have $(h, h') \succeq_i (h, h'')$.

Definition: Given a strategy s_i and a history h for the extensive game Γ, denote by $s_i\mid_h$ the strategy that s_i induces in the subgame $\Gamma(h)$. That is, $s_i\mid_h(h') = s_i(h, h')$ for each $h' \in H\mid_h$. Also, denote by O_h the outcome function of $\Gamma(h)$.

Definition 97.2: A **subgame perfect equilibrium** of an extensive game with perfect information, $\Gamma = \langle N, H, P, (\succeq_i) \rangle$, is a strategy profile s^* such that, for every player $i \in N$ and every non-terminal history $h \in H \setminus Z$ for which $P(h) = i$, we have

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

for every strategy s_i of player i in the subgame $\Gamma(h)$.

The following Lemma is useful for checking whether we have a subgame perfect equilibrium.

Lemma 98.2 (the one-deviation property): Let $\Gamma = \langle N, H, P, (\succeq_i) \rangle$ 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 for which $P(h) = i$, we have

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

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

In each subgame, just check that the player who makes the first move cannot improve by changing his/her initial action.

Proof: Clearly a subgame perfect equilibrium must satisfy the one-deviation property; this is the easy direction.

Suppose s^* satisfies the one deviation property but is not a subgame perfect equilibrium, so some player i can deviate profitably in some subgame $\Gamma(h')$. Then there is a profitable deviation s_i of player i in the subgame $\Gamma(h')$ for which

$$s_i(h) \neq s_i^* \,|_{h'}\, (h) \tag{1}$$

holds for some histories h of $\Gamma(h')$. Since Γ has a finite horizon, so does $\Gamma(h')$. From among all profitable deviations, choose s_i for which the number of histories h for which (1) holds is minimal, and let h^* be the longest such history.

Then the initial history of $\Gamma(h', h^*)$ is the only history h at which we have $s_i(h) \neq s_i^* \,|_{h'}\, (h)$. Since s_i^* cannot be improved by a single deviation, replacing $s_i \,|_{h',h^*}$ with $s_i^* \,|_{h',h^*}$ (but otherwise maintaining s_i) is also a profitable deviation in the subgame $\Gamma(h')$. This contradicts the fact that s_i is chosen so that the number of histories h for which (1) holds is minimal. ■

For infinite horizon games, there may be strategy profiles that cannot be improved by a single deviation at any history, but which are not subgame perfect equilibrium profiles. Thus, the one-deviation property does not hold for infinite horizon games.

Proposition 99.2 (Kuhn's theorem): Every finite extensive game with perfect information has a subgame perfect equilibrium.

Proof is constructive, by **backwards induction**.

Not all finite horizon games with perfect information have a subgame perfect equilibrium, and not all games with a finite number of actions after every history have a subgame perfect equilibrium.

Is there an extensive game with perfect information, for which there exists a NE but no subgame perfect equilibrium? there exists a subgame perfect equilibrium but not a NE?

Extensive Games with Exogenous Uncertainty (chance moves)

Definition: An extensive game with perfect information and chance moves is a tuple, $\langle N, H, P, f_c, (\succeq_i)\rangle$, with new elements

1. P is a function from non-terminal histories to $N \cup \{c\}$. If $P(h) = c$, then chance or nature determines the action taken after the history, h.

2. For each $h \in H$ with $P(h) = c$, we have a probability measure on $A(h)$, denoted by $f_c(\cdot \mid h)$. Each such measure is assumed to be independent of every other such measure.

3. For each player $i \in N$, $\succeq_i$ is a preference relation over lotteries over the set of terminal histories.

Note: The one-deviation property and Kuhn's theorem both hold when chance moves are introduced. (chess vs. backgammon)

Introducing Simultaneous Moves

Definition: An extensive game with perfect information and simultaneous moves is a tuple, $\langle N, H, P, (\succsim_i)\rangle$, where N, H, and $(\succsim_i)$ are as in Definition 89.1, and P is a function that assigns to each non-terminal history h a **set** of players, such that there is a collection of sets, $\{A_i(h)\}_{i\in P(h)}$, for which $A(h) = \{a : (h, a) \in H\} = \times_{i\in P(h)} A_i(h)$.

For a player who chooses an action after the history h, the set of possible actions is $A_i(h)$.

$\times_{i\in P(h)} A_i(h)$ is the set of action profiles that can occur immediately after the history, h. The interpretation is that players in $P(h)$ choose their actions simultaneously.

A history is now a sequence of vectors (of actions), rather than a sequence of actions.

A strategy of player $i \in N$ is a function that assigns an action in $A_i(h)$ for every non-terminal history h for which $i \in P(h)$.

The definition of subgame perfect equilibrium is the same as Definition 97.2, except that $i \in P(h)$ replaces $i = P(h)$.

Notice that this framework allows us easily to extend subgame perfect equilibrium to simultaneous play. If we defined an extensive game in terms of a game tree with information sets, such an extension would be awkward.

We will now go over some examples of games that illustrate and question the logic of backwards induction, and games that illustrate the concept of forward induction.

The Chain Store Game

A chain store has branches in K cities, and in each city, $k = 1, ..., K$, there is a competitor. In period k, the competitor in city k decides whether to enter the market or stay out, and if firm k enters, the chain store must decide whether to fight or cooperate. This is a game of perfect information, with the payoffs in each city given in the figure (next slide). Firm k cares only about the actions taken in its city, but the chain store's payoff is the sum of the payoffs it generates in each city.

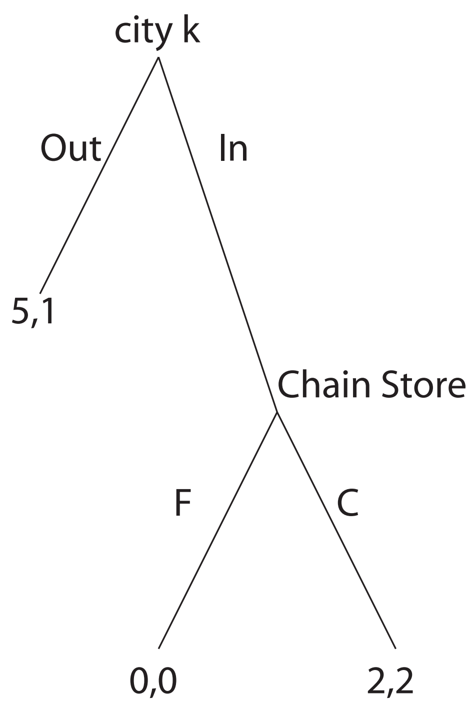

city k
Out
In
5,1
Chain Store
F
C
0,0
2,2

To formalize the game, let $Q = \{out, (in, C), (in, F)\}$ denote the set of outcomes in each city.

The set of histories is

$$\left(\bigcup_{k=0}^{K} Q^k\right) \bigcup \left(\bigcup_{k=0}^{K-1} (Q^k \times \{in\}\right)$$

and the player function is

$$\begin{aligned} P(h) &= k+1 \quad \text{if } \; h \in Q^k \\ P(h) &= \text{chain store} \quad \text{if } h \in Q^k \times \{in\}. \end{aligned}$$

Every terminal history in which the outcome in any period is either out or (in, C) is a **Nash equilibrium** outcome.

There is a unique **subgame perfect equilibrium**, where each competitor chooses in and the chain store always chooses C.

For $K = 1$, subgame perfection eliminates the bad NE.

For large K, isn't it more reasonable to think that the chain store will establish a reputation for being tough? Moreover, if we see the chain store fighting the first 10 competitors, is it reasonable for the next competitor to enter?

The Centipede Game

At each stage, a player can either stop the game, or continue the game, thereby sacrificing one dollar so that the other player can receive more than one dollar.

There is a unique subgame perfect equilibrium, where each player stops the game after every history.

There are several Nash equilibria, but all of them involve both players stopping the game at their first opportunity.

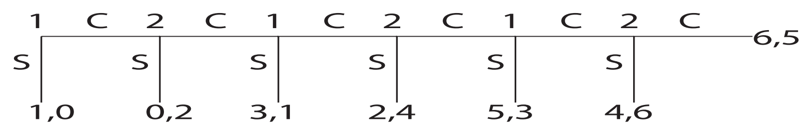

For a very long centipede, with payoffs in the hundreds, will player 1 stop immediately?

Since player 1 starting with C is not consistent with backwards induction, is it reasonable for player 2 to believe that player 1 will use backwards induction in the future?

Forward Induction

Suppose that, as an alternative to going to the concert and playing the battle of the sexes game, player 1 has an opportunity to stay home and read a book.

1. Player 1 chooses "book" or "concert."

2. If player 1 chooses "book" the game is over and payoffs are $(2,2)$. If player 1 chooses "concert" the two players choose either the Bach concert or the Stravinsky concert, according to the battle of the sexes game:

		player 2 B	player 2 S
player 1	B	$3,1$	$0,0$
	S	$0,0$	$1,3$

Because of the book option, $(1, 3)$ cannot be the payoff in any Nash equilibrium. The two subgame perfect equilibria are (book,S,S) and (concert,B,B).

In the concert subgame, is (S,S) a reasonable Nash equilibrium after player 1 decided not to read the book? Why would player 1 choose concert, only to choose an action that cannot yield as high a payoff that she could have guaranteed by choosing book? Playing B is the only justification of choosing concert.

Under forward induction, when player 1 chooses concert, player 2 infers that player 1 intends to play B. Since player 1 knows this, the only remaining equilibrium is (concert,B,B).

Similar example: instead of having the option to read a book, player 1 has the option to burn a dollar. Burning the dollar must signal the intention to play B.

burn \$1

		player 2: B	player 2: S
player 1	B	$2, 1$	$-1, 0$
	S	$-1, 0$	$0, 3$

burn \$0

		player 2: B	player 2: S
player 1	B	$3, 1$	$0, 0$
	S	$0, 0$	$1, 3$

But what is player 2 to infer if player 1 does not burn the money? Player 1 must intend to play B here as well, because she can receive a payoff of 2 by burning the dollar!