Extensive Form Games

Mihai Manea

MIT

Extensive-Form Games

- N: finite set of players; nature is player $0 \in N$
- tree: order of moves
- payoffs for every player at the terminal nodes
- information partition
- actions available at every information set
- description of how actions lead to progress in the tree
- random moves by nature

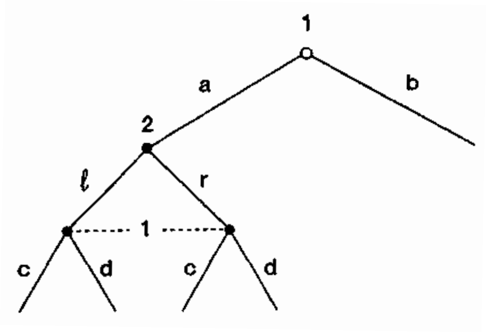

Game Tree

- $(X, >)$: tree
- X: set of nodes
- $x > y$: node x precedes node y
- $\phi \in X$: initial node, $\phi > x, \forall x \in X \setminus \{\phi\}$
- $>$ transitive ($x > y, y > z \Rightarrow x > z$) and asymmetric ($x > y \Rightarrow y \not> x$)
- every node $x \in X \setminus \{\phi\}$ has one immediate predecessor: $\exists x' > x$ s.t. $x'' > x \;\&\; x'' \neq x' \Rightarrow x'' > x'$
- $Z = \{z \mid \nexists x, z > x\}$: set of terminal nodes
- $z \in Z$ determines a unique path of moves through the tree, payoff $u_i(z)$ for player i

Information Partition

- information partition: a partition of $X \setminus Z$
- node x belongs to information set $h(x)$
- player $i(h) \in N$ moves at every node x in information set h
- $i(h)$ knows that he is at some node of h but does not know which one
- same player moves at all $x \in h$, otherwise players might disagree on whose turn it is
- $i(x) := i(h(x))$

Actions

- $A(x)$: set of available actions at $x \in X \setminus Z$ for player $i(x)$
- $A(x) = A(x') =: A(h), \forall x' \in h(x)$ (otherwise $i(h)$ might play an infeasible action)
- each node $x \neq \phi$ associated with the last action taken to reach it
- every immediate successor of x labeled with a different $a \in A(x)$ and vice versa
- move by nature at node x: probability distribution over $A(x)$

Strategies

- $H_i = \{h | i(h) = i\}$
- $S_i = \prod_{h \in H_i} A(h)$: set of pure strategies for player i
- $s_i(h)$: action taken by player i at information set $h \in H_i$ under $s_i \in S_i$
- $S = \prod_{i \in N} S_i$: strategy profiles
- A strategy is a complete contingent plan specifying the action to be taken at each information set.
- Mixed strategies: $\sigma_i \in \Delta(S_i)$
- mixed strategy profile $\sigma \in \prod_{i \in N} \Delta(S_i) \to$ probability distribution $O(\sigma) \in \Delta(Z)$
- $u_i(\sigma) = \mathbb{E}_{O(\sigma)}(u_i(z))$

Strategic Form

- The strategic form representation of the extensive form game is the normal form game defined by (N, S, u)
- A mixed strategy profile is a Nash equilibrium of the extensive form game if it constitutes a Nash equilibrium of its strategic form.

Grenade Threat Game

Player 2 threatens to explode a grenade if player 1 doesn't give him \$1000.

- Player 1 chooses between g and $\neg g$.
- Player 2 observes player 1's choice, then decides whether to explode a grenade that would kill both.

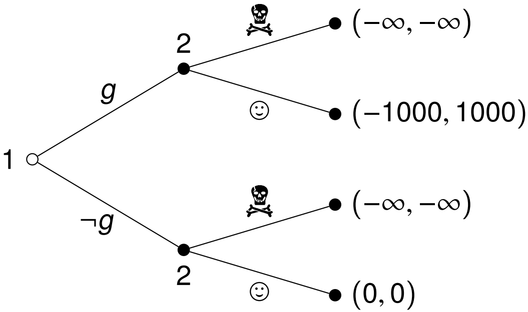

Strategic Form Representation

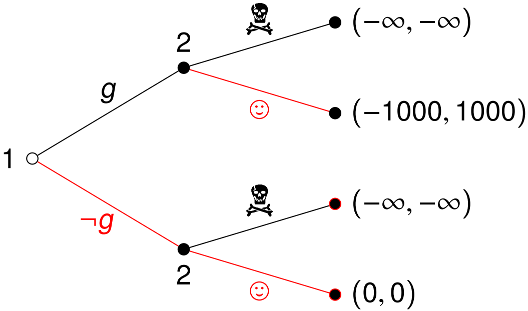

	☠, ☠	☠, ☺	☺, ☠	☺, ☺
g	$-\infty, -\infty$	$-\infty, -\infty$	$-1000, 1000^*$	$-1000, 1000$
$\neg g$	$-\infty, -\infty$	$0, 0^*$	$-\infty, -\infty$	$0, 0^*$

Three pure strategy Nash equilibria. Only $(\neg g, ☺, ☺)$ is subgame perfect.
☠ is not a credible threat.

Behavior Strategies

- $b_i(h) \in \Delta(A(h))$: behavior strategy for player $i(h)$ at information set h
- $b_i(a|h)$: probability of action a at information set h
- behavior strategy $b_i \in \prod_{h \in H_i} \Delta(A(h))$
- independent mixing at each information set
- b_i outcome equivalent to the mixed strategy

$$\sigma_i(s_i) = \prod_{h \in H_i} b_i(s_i(h)|h) \tag{1}$$

- Is every mixed strategy equivalent to a behavior strategy?
- Yes, under perfect recall.

Perfect Recall

No player forgets any information he once had or actions he previously chose.

- If $x'' \in h(x'), x > x'$, and the same player i moves at both x and x' (and thus at x''), then there exists $\hat{x} \in h(x)$ (possibly $\hat{x} = x$) s.t. $\hat{x} > x''$ and the action taken at x along the path to x' is the same as the action taken at $\hat{x}$ along the path to x''.
- x' and x'' distinguished by information i does not have, so he cannot have had it at $h(x)$
- x' and x'' consistent with the same action at $h(x)$ since i must remember his action there
- Equivalently, every node in $h \in H_i$ must be reached via the same sequence of i's actions.

Equivalent Behavior Strategies

- $R_i(h) = \{s_i | h$ is on the path of (s_i, s_{-i}) for some $s_{-i}\}$: set of i's pure strategies that do not preclude reaching information set $h \in H_i$
- Under perfect recall, a mixed strategy σ_i is equivalent to a behavior strategy b_i defined by

$$b_i(a|h) = \frac{\sum_{\{s_i \in R_i(h) | s_i(h) = a\}} \sigma_i(s_i)}{\sum_{s_i \in R_i(h)} \sigma_i(s_i)} \quad (2)$$

when the denominator is positive.

Theorem 1 (Kuhn 1953)

In extensive form games with perfect recall, mixed and behavior strategies are outcome equivalent under the formulae (1) & (2).

Proof

- $h_1, \ldots, h_{\bar{k}}$: player i's information sets preceding h in the tree
- Under perfect recall, reaching any node in h requires i to take the same action a_k at each h_k,

$$R_i(h) = \{s_i | s_i(h_k) = a_k, \forall k = \overline{1, \bar{k}}\}.$$

- Conditional on getting to h, the distribution of continuation play at h is given by the relative probabilities of the actions available at h under the restriction of σ_i to $R_i(h)$,

$$b_i(a|h) = \frac{\sum_{\{s_i | s_i(h_k) = a_k, \forall k = \overline{1, \bar{k}} \ \& \ s_i(h) = a\}} \sigma_i(s_i)}{\sum_{\{s_i | s_i(h_k) = a_k, \forall k = \overline{1, \bar{k}}\}} \sigma_i(s_i)}.$$

Example

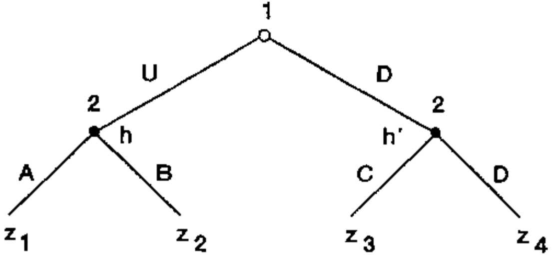

Figure: Different mixed strategies can generate the same behavior strategy.

- $S_2 = \{(A, C), (A, D), (B, C), (B, D)\}$
- Both $\sigma_2 = 1/4(A, C) + 1/4(A, D) + 1/4(B, C) + 1/4(B, D)$ and $\sigma_2 = 1/2(A, C) + 1/2(B, D)$ generate—and are equivalent to—the behavior strategy b_2 with $b_2(A|h) = b_2(B|h) = 1/2$ and $b_2(C|h') = b_2(D|h') = 1/2$.

Example with Imperfect Recall

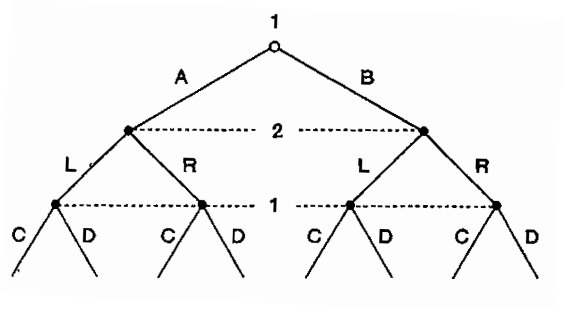

Figure: Player 1 forgets what he did at the initial node.

- $S_1 = \{(A, C), (A, D), (B, C), (B, D)\}$
- $\sigma_1 = 1/2(A, C) + 1/2(B, D) \rightarrow b_1 = (1/2A + 1/2B, 1/2C + 1/2D)$
- b_1 not equivalent to σ_1
- (σ_1, L): prob. 1/2 for paths (A, L, C) and (B, L, D)
- (b_1, L): prob. 1/4 to paths $(A, L, C), (A, L, D), (B, L, C), (B, L, D)$

Imperfect Recall and Correlations

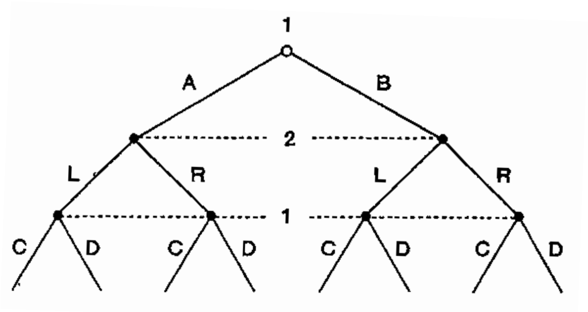

- Since both *A* vs. *B* and *C* vs. *D* are choices made by player 1, the strategy σ_1 under which player 1 makes all his decisions at once allows choices at different information sets to be correlated
- Behavior strategies cannot produce this correlation, because when it comes time to choose between *C* and *D*, player 1 has forgotten whether he chose *A* or *B*.

Absent Minded Driver

Piccione and Rubinstein (1997)

- A drunk driver has to take the third out of five exits on the highway (exit 3 has payoff 1, other exits payoff 0).
- The driver cannot read the signs and forgets how many exits he has already passed.
- At each of the first four exits, he can choose C (continue) or E (exit)... imperfect recall: choose same action.
- C leads to exit 5, while E leads to exit 1.
- Optimal solution involves randomizing: probability p of choosing C maximizes $p^2(1-p)$, so $p = 2/3$.
- "Beliefs" given $p = 2/3$: $(27/65, 18/65, 12/65, 8/65)$
- E has conditional "expected" payoff of $12/65$, C has 0. Optimal strategy: E with probability 1, inconsistent.

Conventions

- Restrict attention to games with perfect recall, so we can use mixed and behavior strategies interchangeably.
- Behavior strategies are more convenient.
- Drop notation *b* for behavior strategies and denote by $\sigma_i(a|h)$ the probability with which player *i* chooses action *a* at information set *h*.

Survivor

THAI 21

- Two players face off in front of 21 flags.
- Players alternate in picking 1, 2, or 3 flags at a time.
- The player who successfully grabs the last flag wins.

Game of luck?

Backward Induction

- An extensive form game has perfect information if all information sets are singletons.
- Can solve games with perfect information using backward induction.
- Finite game $\rightarrow \ \exists$ penultimate nodes (successors are terminal nodes).
- The player moving at each penultimate node chooses an action that maximizes his payoff.
- Players at nodes whose successors are penultimate/terminal choose an optimal action given play at penultimate nodes.
- Work backwards to initial node...

Theorem 2 (Zermelo 1913; Kuhn 1953)

In a finite extensive form game of perfect information, the outcome(s) of backward induction constitutes a pure-strategy Nash equilibrium.

Market Entrance

- Incumbent firm 1 chooses a level of capital K_1 (which is then fixed).
- A potential entrant, firm 2, observes K_1 and chooses its capital K_2.
- The profit for firm $i = 1, 2$ is $K_i(1 - K_1 - K_2)$ (firm i produces output K_i, we use earlier demand function).
- Each firm dislikes capital accumulation by the other.
- A firm's *marginal* value of capital decreases with the other's.
- Capital levels are strategic substitutes.

Stackelberg Competition

- Profit maximization by firm 2 requires

$$K_2 = \frac{1 - K_1}{2}.$$

- Firm 1 anticipates that firm 2 will act optimally, and therefore solves

$$\max_{K_1} \left\{ K_1 \left(1 - K_1 - \frac{1 - K_1}{2} \right) \right\}.$$

- Solution involves $K_1 = 1/2$, $K_2 = 1/4$, $\pi_1 = 1/8$, and $\pi_2 = 1/16$.
- Firm 1 has first mover advantage.
- In contrast, in the simultaneous move game, $K_1 = 1/3$, $K_2 = 1/3$, $\pi_1 = 1/9$, and $\pi_2 = 1/9$.

Centipede Game

- Player 1 has two piles in front of her: one contains 3 coins, the other 1.
- Player 1 can either take the larger pile and give the smaller one to player 2 (T) or push both piles across the table to player 2 (C).
- Every time the piles pass across the table, one coin is added to each.
- Players alternate in choosing whether to take the larger pile (T) or trust opponent with bigger piles (C).
- The game lasts 100 rounds.

What's the backward induction solution?

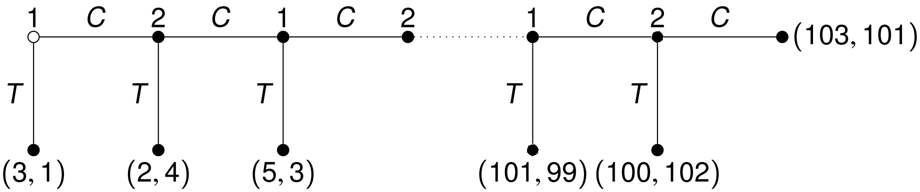

Chess Players and Backward Induction

Palacios-Huerta and Volij (2009)

- chess players and college students behave differently in the centipede game.
- Higher-ranked chess players end the game earlier.
- All Grandmasters in the experiment stopped at the first opportunity.
- Chess players are familiar with backward induction reasoning and need less learning to reach the equilibrium.
- Playing against non-chess-players, even chess players continue in the game longer.
- In long games, common knowledge of the ability to do complicated inductive reasoning becomes important for the prediction.

Subgame Perfection

- Backward induction solution is more than a Nash equilibrium.
- Actions are optimal given others' play—and form an equilibrium—starting at *any* intermediate node: subgame perfection. . . rules out non-credible threats.
- Subgame perfection extends backward induction to imperfect information games.
- Replace "smallest" subgames with a Nash equilibrium and iterate on the reduced tree (if there are multiple Nash equilibria in a subgame, all players expect same play).

Subgames

Subgame: part of a game that can be analyzed separately; strategically and informationally independent. . . information sets not "chopped up."

Definition 1

A **subgame** G of an extensive form game T consists of a single node x and *all* its successors in T, with the property that if $x' \in G$ and $x'' \in h(x')$ then $x'' \in G$. The information sets, actions and payoffs in the subgame are inherited from T.

False Subgames

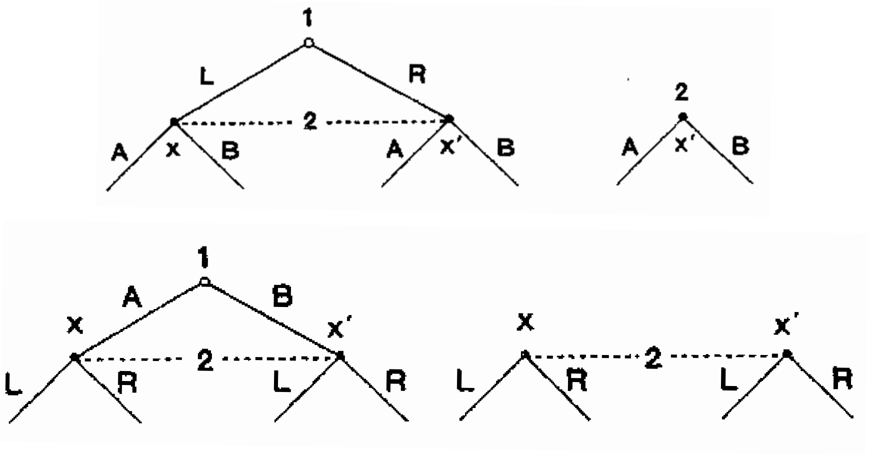

Subgame Perfect Equilibrium

σ: behavior strategy in T

- $\sigma|G$: the strategy profile induced by σ in subgame G of T (start play at the initial node of G, follow actions specified by σ, obtain payoffs from T at terminal nodes)
- Is $\sigma|G$ a Nash equilibrium of G for any subgame G?

Definition 2

A strategy profile σ in an extensive form game T is a **subgame perfect equilibrium** if $\sigma|G$ is a Nash equilibrium of G for every subgame G of T.

- Any game is a subgame of itself $\rightarrow$ a subgame perfect equilibrium is a Nash equilibrium.
- Subgame perfection coincides with backward induction in games of perfect information.

Nuclear Crisis

Russia provokes the US. . .

- The U.S. can choose to escalate (E) or end the game by ignoring the provocation (I).
- If the game escalates, Russia faces a similar choice: to back down (B), but lose face, or escalate (E).
- Escalation leads to nuclear crisis: a simultaneous move game where each nation chooses to either retreat (R) and lose credibility or detonate (☢). Unless both countries retreat, retaliation to the first nuclear strike culminates in nuclear disaster, which is infinitely costly.

The Extensive Form

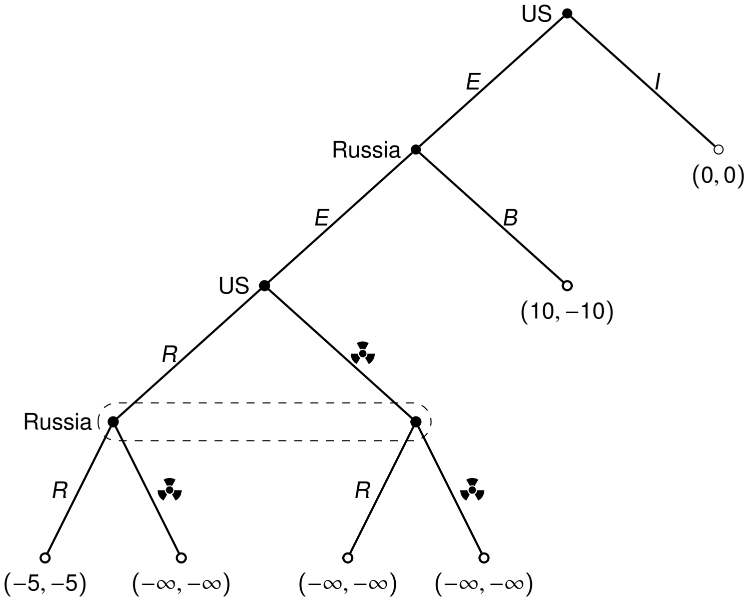

Last Stage

The simultaneous-move game at the last stage has two Nash equilibria.

	R	☢
R	$-5, 5^*$	$-\infty, -\infty$
☢	$-\infty, -\infty$	$-\infty, -\infty^*$

One Subgame Perfect Equilibrium

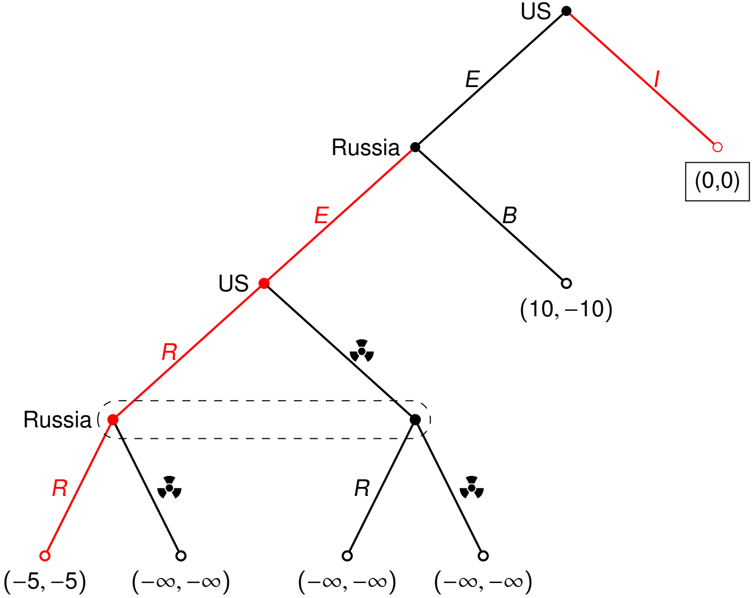

Another Subgame Perfect Equilibrium