This Segment: Techniques for solving games

Lecture 1: Game representations, game-theoretic solution concepts, and complexity

Tuomas Sandholm
Computer Science Department
Carnegie Mellon University

The heart of the problem

- In a 1-agent setting, agent's expected utility maximizing strategy is well-defined
- But in a multiagent system, the outcome may depend on others' strategies also

Terminology

- *Agent = player*
- *Action = move* = choice that agent can make at a point in the game
- *Strategy* s_i = mapping from history (to the extent that the agent i can distinguish) to actions
- *Strategy set* S_i = strategies available to the agent
- *Strategy profile* $(s_1, s_2, ..., s_{|A|})$ = one strategy for each agent
- Agent's utility is determined after each agent (including nature that is used to model uncertainty) has chosen its strategy, and game has been played: $u_i = u_i(s_1, s_2, ..., s_{|A|})$

Game representations

Extensive form

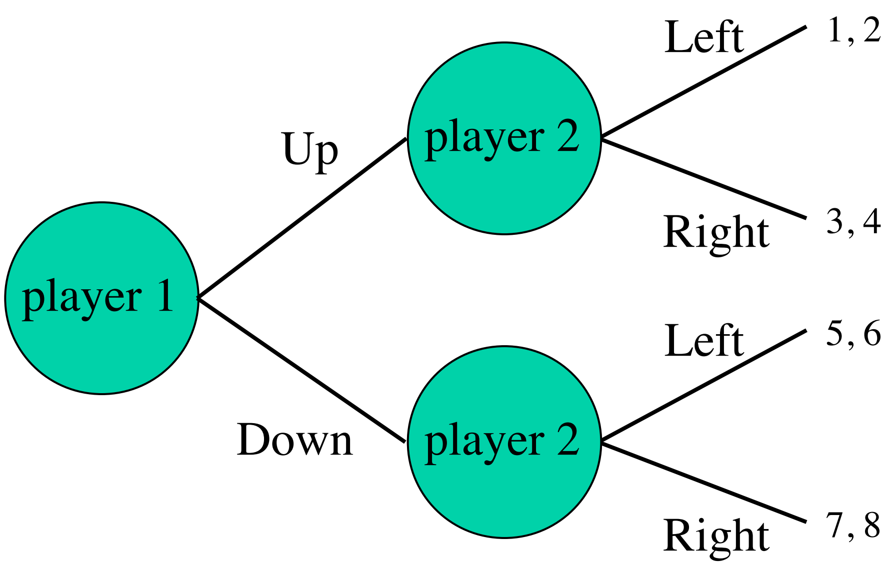

Matrix form
(aka normal form
aka strategic form)

player 2's strategy

player 1's strategy	Left, Left	Left, Right	Right, Left	Right, Right
Up	1, 2	1, 2	3, 4	3, 4
Down	5, 6	7, 8	5, 6	7, 8

Potential combinatorial explosion →

Dominant strategy "equilibrium"

- ***Best response*** s_i^*: for all s_i', $u_i(s_i^*,s_{-i}) \geq u_i(s_i',s_{-i})$
- ***Dominant strategy*** s_i^*: s_i^* is a best response for all s_{-i}
  - Does not always exist
  - Inferior strategies are called "dominated"
- ***Dominant strategy equilibrium*** **is a strategy profile where each agent has picked its dominant strategy**
  - Does not always exist
  - Requires no counterspeculation

	cooperate	defect
cooperate	3, 3	0, 5
defect	5, 0	1, 1

Pareto optimal?

Social welfare maximizing?

Nash equilibrium [Nash50]

- Sometimes an agent's best response depends on others' strategies: a dominant strategy does not exist
- A strategy profile is a ***Nash equilibrium*** if no player has incentive to deviate from his strategy given that others do not deviate: for every agent i, $u_i(s_i^*, s_{-i}) \geq u_i(s_i', s_{-i})$ for all s_i'
  - Dominant strategy equilibria are Nash equilibria but not vice versa
  - Defect-defect is the only Nash eq. in Prisoner's Dilemma
  - Battle of the Sexes game
    - Has no dominant strategy equilibria

		Woman: boxing	Woman: ballet
Man	boxing	2, 1	0, 0
Man	ballet	0, 0	1, 2

Criticisms of Nash equilibrium

- Not unique in all games, e.g. Battle of the Sexes
  - Approaches for addressing this problem
    - Refinements (=strengthenings) of the equilibrium concept
      - Eliminate weakly dominated strategies first
      - Choose the Nash equilibrium with highest welfare
      - Subgame perfection …
    - Focal points
    - Mediation
    - Communication
    - Convention
    - Learning
- Does not exist in all games

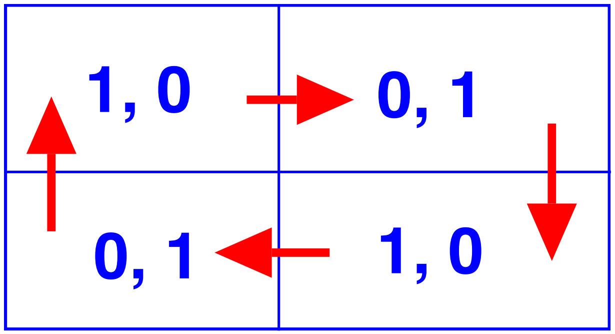

Existence of (pure strategy) Nash equilibria

- Thrm.
  - Any finite game,
  - where each action node is alone in its information set
    - (i.e. at every point in the game, the agent whose turn it is to move knows what moves have been played so far)
  - is dominance solvable by backward induction (at least as long as ties are ruled out)
- Constructive proof: Multi-player minimax search

Rock-scissors-paper game

Sequential moves

Rock-scissors-paper game

Simultaneous moves

Mixed strategy Nash equilibrium

Mixed strategy **= agent's chosen probability distribution over pure strategies from its strategy set**

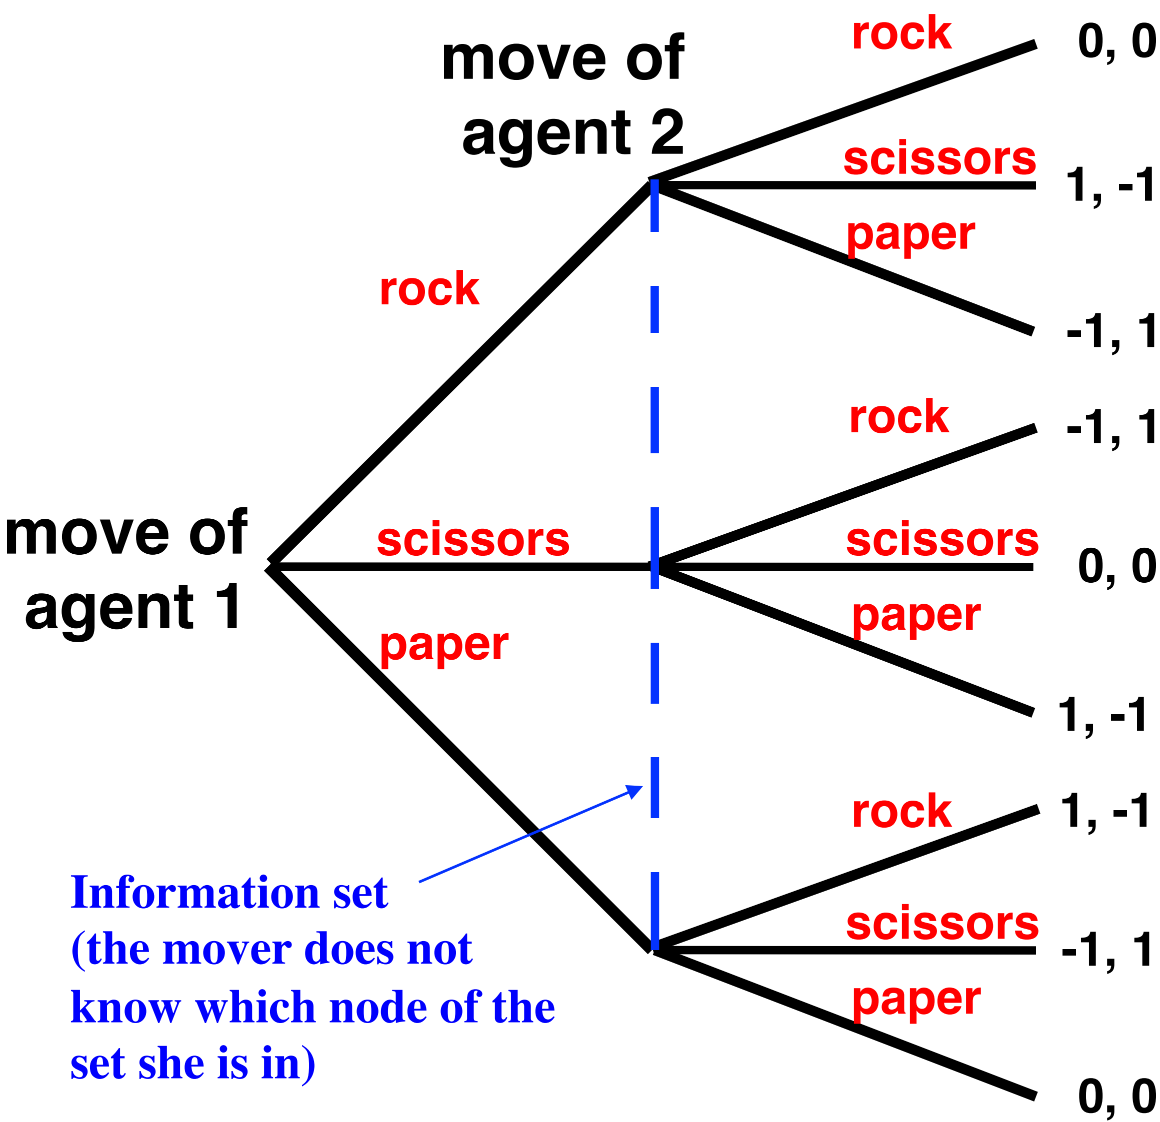

Each agent has a best response strategy and beliefs (consistent with each other)

Symmetric mixed strategy Nash eq: Each player plays each pure strategy with probability 1/3

In mixed strategy equilibrium, each strategy that occurs in the mix of agent i has equal expected utility to i

Existence & complexity of mixed strategy Nash equilibria

- **Every finite player, finite strategy game has at least one Nash equilibrium if we admit mixed strategy equilibria as well as pure [Nash 50]**
  - **(Proof is based on Kakutani's fix point theorem)**

- **May be hard to compute**
  - **Complexity of finding a Nash equilibrium in a normal form game:**
    - **2-player 0-sum games can be solved in polytime with LP**
    - **2-player games are PPAD-complete (even with 0/1 payoffs) [Chen, Deng & Teng JACM-09; Abbott, Kane & Valiant FOCS-05; Daskalakis, Goldberg & Papadimitriou STOC-06],**
      and NP-complete to find an even approximately *good* Nash equilibrium [Conitzer & Sandholm GEB-08]
    - **3-player games are FIXP-complete [Etessami & Yannakakis FOCS-07]**

Ultimatum game

(for distributional bargaining)

Subgame perfect equilibrium [Selten 72] & credible threats

- Proper subgame = subtree (of the game tree) whose root is alone in its information set
- Subgame perfect equilibrium = strategy profile that is in Nash equilibrium in every proper subgame (including the root), whether or not that subgame is reached along the equilibrium path of play
- E.g. Cuban missile crisis

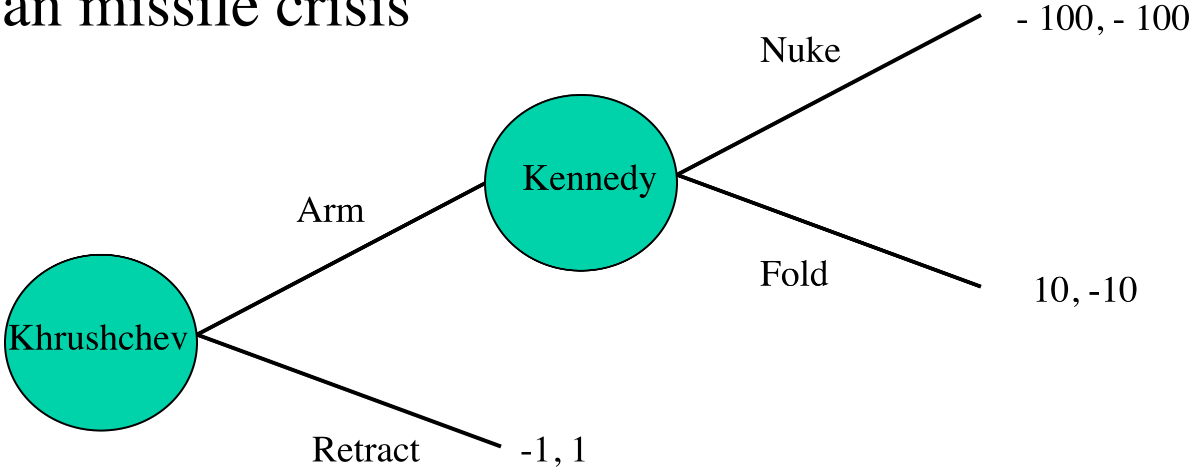

- Pure strategy Nash equilibria: (Arm,Fold), (Retract,Nuke)
- Pure strategy subgame perfect equilibria: (Arm,Fold)
- Conclusion: Kennedy's Nuke threat was not *credible*

Ultimatum game, again

Thoughts on credible threats

- Could use software as a commitment device
  - If one can credibly convince others that one cannot change one's software agent, then revealing the agent's code acts as a credible commitment to one's strategy
  - E.g. nuke in the missile crisis
  - E.g. accept no less than 60% as the second mover in the ultimatum game
- Restricting one's strategy set can increase one's utility
  - This cannot occur in single agent settings
- Social welfare can increase or decrease

Solution concepts

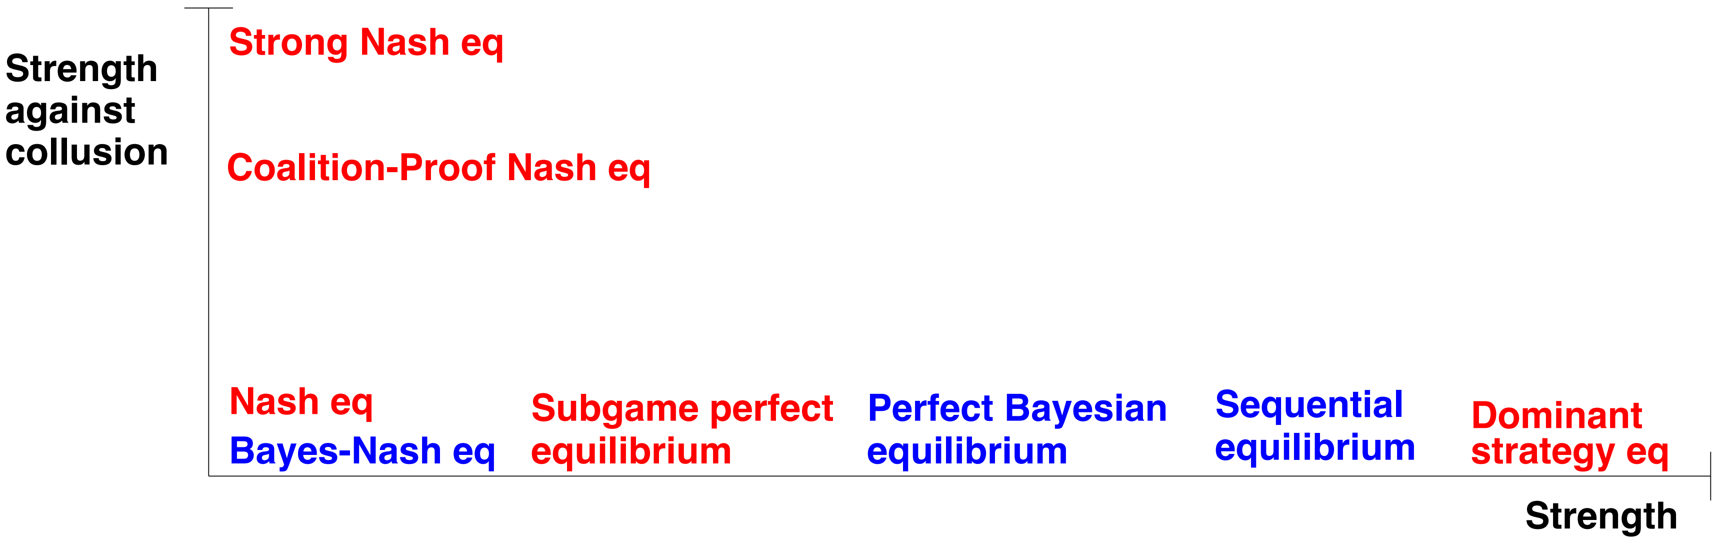

Ex post equilibrium = Nash equilibrium for all priors

There are other equilibrium refinements too (see, e.g., following slides & wikipedia)

Definition of a Bayesian game

[Harsanyi 67/68]

- N is the set of players.
- Ω is the set of the states of nature.
  - For instance, in a card game, it can be any order of the cards.
- A_i is the set of actions for player i. $A = A_1 \times A_2 \times \ldots \times A_n$
- T_i is the type set of player i. For each state of nature, the game will have different types of players (one type per player).
  - For instance, in a car selling game, it will be how much the player values the car
- $C_i \in A_i \times T_i$ defines the available actions for player i of some type in T_i.
- u: $\Omega \times A \rightarrow R$ is the payoff function for player i.
- p_i is the probability distribution over Ω for each player i, i.e., each player has different views of the probability distribution over the states of nature.
  - In the game, they may never know the exact state of nature.

Solution concepts for Bayesian games

More refined

- A (Bayesian) Nash equilibrium is a strategy profile and *beliefs specified for each player about the types of the other players* that maximizes the expected utility for each player given their beliefs about the other players' types and given the strategies played by the other players
- Perfect Bayesian equilibrium (PBE)
  - Players place beliefs on nodes occurring in their information sets
  - A belief system is *consistent* for a given strategy profile if the probability assigned by the system to every node is computed as the probability of that node being reached given the strategy profile, i.e., by Bayes' rule
  - A strategy profile is *sequentially rational* at a particular information set for a particular *belief system* if the expected utility of the player whose information set it is is maximal given the strategies played by the other players
    - A strategy profile is sequentially rational for a particular belief system if it satisfies the above for every information set
  - A PBE is a strategy profile and a belief system such that the strategies are sequentially rational given the belief system and the belief system is *consistent*, wherever possible, given the strategy profile
    - 'wherever possible' clause is necessary: some information sets might be reached with zero probability given the strategy profile; hence Bayes' rule cannot be employed to calculate the probability of nodes in those sets. Such information sets are said to be *off the equilibrium path* and any beliefs can be assigned to them
- Sequential equilibrium [Kreps and Wilson 82]. Refinement of PBE that specifies constraints on beliefs in such zero-probability information sets. Strategies and beliefs must be a limit point of a sequence of totally mixed strategy profiles and associated sensible (in PBE sense) beliefs
- Extensive-form trembling hand ***perfect*** equilibrium [Selten 75]. Require every move at every information set to be taken with non-zero probability. Take limit as tremble probability $\rightarrow 0$
- Extensive-form ***proper*** equilibrium [Myerson 78]. Idea: Costly trembles much less likely. At any information set, for any two actions A and B, if the mover's utility from B is less than from A, then prob(B) $\leq \varepsilon$ prob(A). Take limit as $\varepsilon \rightarrow 0$

Solution concepts for Bayesian games …

- Extensive-form perfect / proper equilibrium can involve playing weakly dominated strategies => argument for other solution concepts:
- *Normal-form* perfect equilibrium
  - Normal- and extensive-form perfect equilibria are incomparable
  - A normal-form perfect equilibrium of an extensive-form game may or may not be sequential (and might not even be subgame perfect)
- Quasi-perfect equilibrium [van Damme 84]
  - Informally, a player takes observed as well as potential future mistakes of his opponents into account but assumes that he himself will not make a mistake in the future, even if he observes that he has done so in the past
  - Incomparable to extensive-form perfect / proper
- *Normal-form* proper equilibrium
  - Always sequential and quasi-perfect
  - For 0-sum games, provides a strategy that maximizes the conditional utility (among minmax strategies), conditioned on the opponent making a mistake. (Mistake is defined as a pure strategy that does not achieve the value of the game against all minmax strategies.)

More refined