CMSC 474, Introduction to Game Theory

10. Epsilon-Nash Equilibria

Mohammad T. Hajiaghayi

University of Maryland

Outline

- Chapter 2 discussed two solution concepts:
  - Pareto optimality and Nash equilibrium
- Chapter 3 discusses several more:
  - Maxmin and Minmax
  - Dominant strategies
  - Correlated equilibrium
  - Trembling-hand perfect equilibrium (complicated defintion)
  - ε-Nash equilibrium
  - Evolutionarily stable strategies

ε-Nash Equilibrium

- Reflects the idea that agents might not change strategies if the gain would be very small
- Let $\varepsilon > 0$. A strategy profile $\mathbf{s} = (s_1, \ldots, s_n)$ is an **ε-Nash equilibrium** if for every agent i and for every strategy $s_i' \neq s_i$,

  $$u_i(s_i, \mathbf{s}_{-i}) \geq u_i(s_i', \mathbf{s}_{-i}) - \varepsilon$$

- ε-Nash equilibria exist for every $\varepsilon > 0$
  - Every Nash equilibrium is an ε-Nash equilibrium, and is surrounded by a region of ε-Nash equilibria
- This concept can be computationally useful
  - Algorithms to identify ε-Nash equilibria need consider only a finite set of mixed-strategy profiles (not the whole continuous space)
  - Because of finite precision, computers generally find only ε-Nash equilibria, where ε is roughly the machine precision

Problems with ε-Nash Equilibrium

- For every Nash equilibrium, there are ε-Nash equilibria that approximate it, but the converse isn't true
  - There are ε-Nash equilibria that aren't close to any Nash equilibrium
- Example: the game at right has just one Nash equilibrium: (*D*, *R*)
  - Use IESDS to show it's the only one:
    - For agent 1, *D* dominates *U*, so remove *U*
    - Then for agent 2, *R* dominates *L*
- (*D*, *R*) is also an ε-Nash equilibrium
- But there's another ε-Nash equilibrium: (*U*, *L*)
  - Neither agent can gain more than ε by deviating
  - But its payoffs aren't within ε of the Nash equilibrium

	L	*R*
U	1, 1	0, 0
D	1 + ε/2, 1	500, 500

Problems with ε-Nash Equilibrium

- Some ε-Nash equilibria are very unlikely to arise
- Same example as before
  - Agent 1 might not care about a gain of ε/2, but might reason as follows:
    - Agent 2 may expect agent 1 to to play *D*, since *D* dominates *U*
    - So agent 2 is likely to play *R*
    - If agent 2 plays *R*, agent 1 does *much* better by playing *D* rather than *U*

	L	*R*
U	1, 1	0, 0
D	1 + ε/2, 1	500, 500

- In general, ε-approximation is much messier in games than in optimization problems

Summary

- Maxmin and minmax strategies, and the Minimax Theorem
  - Matching Pennies, Two-Finger Morra
- dominant strategies
  - Prisoner’s Dilemma, Which Side of the Road, Matching Pennies
  - Iterated elimination of dominated strategies (IESDS)
- rationalizability
  - the *p*-Beauty Contest
- correlated equilibrium
  - Battle of the Sexes
- epsilon-Nash equilibria
- evolutionarily stable strategies
  - Body-Size game, Hawk-Dove game