CMSC 474, Introduction to Game Theory

Finding Nash Equilibria

Mohammad T. Hajiaghayi

University of Maryland

Finding Mixed-Strategy Nash Equilibria

- In general, it's tricky to compute mixed-strategy Nash equilibria
  - But easier if we can identify the support of the equilibrium strategies
- In 2x2 games, we can do this easily
- We especially use theorem below proved the previous week

Theorem A: *Always there exists a pure best response* s_i *to* $\mathbf{s}_{-i}$

- **Corollary B:** If (s_1, s_2) is a pure Nash equilibrium only among pure strategies, it should be a Nash equilibrium among mixed strategies as well
- Now let (s_1, s_2) be a Nash equilibrium
- If both s_1, s_2 have supports of size one, it should be one of the cells of the normal-form matrix and we are done by Corollary B
- Thus assume at least one of s_1, s_2 has a support of size two.

Finding Mixed-Strategy Nash Equilibria

- Now if the support of one of s_1, s_2, say s_1, is of size one, i.e., it is pure, then s_2 should be pure as well, unless both actions of player 2 have the same payoffs; in this case any mixed strategy of both actions can be Nash equilibrium.
- Thus in the rest we assume both supports have size two.
  - Thus to find s_1 assume agent 1 selects action a_1 with probability p and action a'_1 with probability 1-p.
  - Now since s_2 has a support of size two, its support must include both of agent 2's actions a_2 and a'_2, and they must have the same expected utility
    - Otherwise agent 2's best response would be just one of them and its support has size one.
  - Hence find p such that $u_2(s_1, a_2) = u_2(s_1, a'_2)$, i.e., solve the equation to find p (and thus s_2)
  - Similarly, find s_2 such that $u_1(a_1, s_2) = u_1(a'_1, s_2)$

Finding Mixed-Strategy Nash Equilibria

Example: Battle of the Sexes

Wife \ Husband	Oper a	Football
Opera	2, 1	0, 0
Football	0, 0	1, 2

- We already saw pure Nash equilibria.
- If there's a mixed-strategy equilibrium,
  - both strategies must be mixtures of {Opera, Football}
  - each must be a best response to the other
- Suppose the husband's strategy is s_h = {(p, Opera), (1–p, Football)}
- Expected utilities of the wife's actions:

$$u_w(\text{Opera}, s_h) = 2p; \qquad u_w(\text{Football}, s_h) = 1(1 - p)$$

- If the wife mixes the two actions, they must have the same expected utility
  - Otherwise the best response would be to *always* use the action whose expected utility is higher
  - Thus $2p = 1 - p$, so $p = 1/3$
- So the husband's mixed strategy is s_h = {(1/3, Opera), (2/3, Football)}

Finding Mixed-Strategy Nash Equilibria

Wife \ Husband	Opera	Football
Opera	2, 1	0, 0
Football	0, 0	1, 2

- Similarly, we can show the wife's mixed strategy is
  - s_w = {(2/3, Opera), (1/3, Football)}
- So the mixed-strategy Nash equilibrium is (s_w , s_h), where
  - s_w = {(2/3, Opera), (1/3, Football)}
  - s_h = {(1/3, Opera), (2/3, Football)}
- Questions:
  - Like all mixed-strategy Nash equilibria, (s_w , s_h) is weak
    - Both players have infinitely many other best-response strategies
    - What are they?
  - How do we know that (s_w , s_h) really is a Nash equilibrium?
    - Indeed the proof is by the way that we found Nash equilibria (s_w , s_h)

Finding Mixed-Strategy Nash Equilibria

- s_w = {(2/3, Opera), (1/3, Football)}
- s_h = {(1/3, Opera), (2/3, Football)}

- Wife's expected utility is
  - 2(2/9) + 1(2/9) + 0(5/9) = 2/3
- Husband's expected utility is also 2/3

2/3 • 1/3 = 2/9

2/3 • 2/3 = 4/9

Wife \ Husband	Opera	Football
Opera	2, 1	0, 0
Football	0, 0	1, 2

1/3 • 1/3 = 1/9

1/3 • 2/3 = 2/9

- It's "fair" in the sense that both players have the same expected payoff
- But it's Pareto-dominated by both of the pure-strategy equilibria
  - In each of them, one agent gets 1 and the other gets 2
- Can you think of a fair way of choosing actions that produces a higher expected utility?

Finding Mixed-Strategy Nash Equilibria

Matching Pennies

- Easy to see that in this game, no pure strategy could be part of a Nash equilibrium
  - For each combination of pure strategies, one of the agents can do better by changing his/her strategy

	Heads	Tails
Heads	1, –1	–1, 1
Tails	–1, 1	1, –1

- Thus there isn't a strict Nash equilibrium since it would be pure.

- But again there's a mixed-strategy equilibrium
  - Can be derived the same way as in the Battle of the Sexes
    - Result is (s,s), where s = {(½, Heads), (½, Tails)}

Another Interpretation of Mixed Strategies

- Suppose agent *i* has a deterministic method for picking a strategy, but it depends on factors that aren't part of the game itself
  - If *i* plays a game several times, *i* may pick different strategies
- If the other players don't know how *i* picks a strategy, they'll be uncertain what *i*'s strategy will be
  - Agent *i*'s mixed strategy is **everyone else's assessment** of how likely *i* is to play each pure strategy
- Example:
  - In a series of soccer penalty kicks, the kicker could kick left or right in a deterministic pattern that the goalie thinks is random

Complexity of Finding Nash Equilibria

- We've discussed how to find Nash equilibria in some special cases
  - Step 1: look for pure-strategy equilibria
    - Examine each cell of the matrix
    - If no cell in the same row is better for agent 1, and

      no cell in the same column is better for agent 2

      then the cell is a Nash equilibrium
  - Step 2: look for mixed-strategy equilibria
    - Write agent 2's strategy as $\{(q, b), (1–q, b')\}$;
      look for q such that a and a' have the same expected utility
    - Write agent 1's strategy as $\{(p, a), (1–p, a')\}$;
      look for p such that b and b' have the same expected utility

	b	b'
a	u_1, v_1	u_2, v_2
a'	u_3, v_3	u_4, v_4

2x2 games

- More generally for two-player games with any number of actions

  for each player, if we know support of each, we can find a mixed-Nash equilibrium in polynomial-time by solving linear equations (via linear program).
- What about the general case?

Complexity of Finding Nash Equilibria

- General case: n players, m actions per player, payoff matrix has m^n cells

  (not in the book)

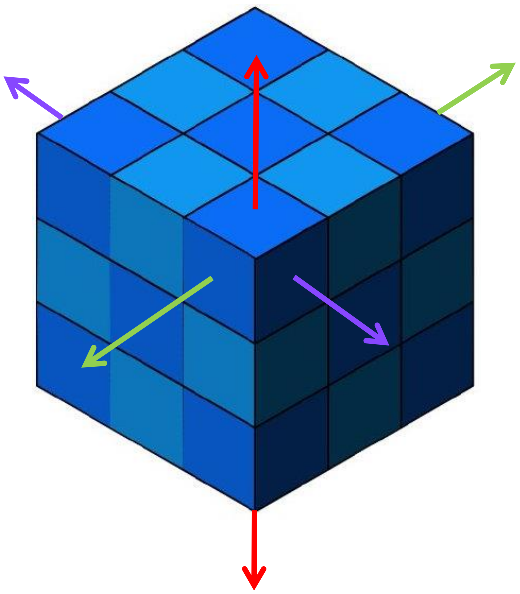

- Brute-force approach:
  - Step 1: Look for pure-strategy equilibria
    - At each cell of the matrix,
      - For each player, can that player do better by choosing a different action?
    - Polynomial time
  - Step 2: Look for mixed-strategy equilibria
    - For every possible combination of supports for $s_1, \ldots, s_n$
      - Solve sets of simultaneous equations
    - Exponentially many combinations of supports
    - Can it be done more quickly?

Complexity of Finding Nash Equilibria

- Two-player games
  - Lemke & Howson (1964): solve a set of simultaneous equations that includes all possible support sets for $s_1, \ldots, s_n$
    - Some of the equations are quadratic => worst-case exponential time
  - Porter, Nudelman, & Shoham (2004)
    - AI methods (constraint programming)
  - Sandholm, Gilpin, & Conitzer (2005)
    - Mixed Integer Programming (MIP) problem
- *n*-player games
  - van der Laan, Talma, & van der Heyden (1987)
  - Govindan, Wilson (2004)
  - Porter, Nudelman, & Shoham (2004)
- Worst-case running time still is exponential in the size of the payoff matrix

Complexity of Finding Nash Equilibria

- There are special cases that can be done in polynomial time in the size of the payoff matrix
  - Finding pure-strategy Nash equilibria
    - Check each square of the payoff matrix
  - Finding Nash equilibria in zero-sum games (see later in this class)
    - Linear programming
- For the general case,
  - It's unknown whether there are polynomial-time algorithms to do it
  - It's unknown whether there are polynomial-time algorithms to compute approximations
  - But we know both questions are PPAD-complete (but not NP-complete) even for two-player games (with some definition of PPAD introduced by Christos Papadimitriou in 1994)
- This is still one of the most important open problems in computational complexity theory

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
- Finding an ε-Nash equilibrium is still PPAD-complete (but not NP-complete) even for two-player games

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*)

  (e.g., 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	$1 + \varepsilon/2$, 1	500, 500