CMSC 474, Introduction to Game Theory

9. Dominant Strategies and Correlated Equilibrium

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
  - ε-Nash equilibrium
  - Evolutionarily stable strategies

Dominant Strategies

- Let s_i and s_i' be two strategies for agent i
  - Intuitively, s_i dominates s_i' if agent i does better with s_i than with s_i' for *every* strategy profile $\mathbf{s}_{-i}$ of the remaining agents
- Mathematically, there are three gradations of dominance:
  - s_i **strictly dominates** s_i' if for every $\mathbf{s}_{-i}$,

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

  - s_i **weakly dominates** s_i' if for every $\mathbf{s}_{-i}$,

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

    and for at least one $\mathbf{s}_{-i}$,

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

  - s_i **very weakly dominates** s_i' if for every $\mathbf{s}_{-i}$,

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

Dominant Strategy Equilibria

- A strategy is **strictly** (resp., **weakly**, **very weakly**) **dominant** for an agent if it strictly (weakly, very weakly) dominates any other strategy for that agent

- A strategy profile $(s_1, . . . , s_n)$ in which every s_i is dominant for agent i (strictly, weakly, or very weakly) is a Nash equilibrium

    - Why?

  - Such a strategy profile forms an **equilibrium in strictly** (**weakly**, **very weakly**) **dominant strategies**

Examples

- Example: **the Prisoner's Dilemma**
  - http://www.youtube.com/watch?v=ED9gaAb2BEw

	C	*D*
C	3, 3	0, 5
D	**5**, 0	**1**, 1

- For agent 1, *D* is strictly dominant
  - If agent 2 uses *C, then*
    - Agent 1's payoff is higher with *D* than with *C*
  - If agent 2 uses *D, then*
    - Agent 1's payoff is higher with *D* than with *C*
- Similarly, *D* is strictly dominant for agent 2
- So (*D*,*D*) is a Nash equilibrium in strictly dominant strategies

	C	*D*
C	3, 3	0, 5
D	5, 0	1, 1

- How do strictly dominant strategies relate to strict Nash equilibria?

Example: Matching Pennies

- **Matching Pennies**
  - If agent 2 uses Heads, then
    - For agent 1, Heads is better than Tails
  - If agent 2 uses Tails, then
    - For agent 1, Tails is better than Heads
  - Agent 1 doesn't have a dominant strategy

    => no Nash equilibrium in dominant strategies

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

- **Which Side of the Road**
  - Same kind of argument as above
  - No Nash equilibrium in dominant strategies

	Left	Right
Left	**1**, 1	0, 0
Right	0, 0	**1**, 1

Elimination of Strictly Dominated Strategies

- A strategy s_i is **strictly** (**weakly**, **very weakly**) **dominated** for an agent i if some other strategy s_i' strictly (weakly, very weakly) dominates s_i

	L	R
~~U~~	~~3, 3~~	~~0, 5~~
D	5, 1	1, 0

- A strictly dominated strategy can't be a best response to any move, so we can eliminate it (remove it from the payoff matrix)
  - This gives a **reduced** game
  - Other strategies may now be strictly dominated, even if they weren't dominated before

	L	~~R~~
D	5, 1	~~1, 0~~

- **IESDS** (Iterated Elimination of Strictly Dominated Strategies):
  - Do elimination repeatedly until no more eliminations are possible
  - When no more eliminations are possible, we have the **maximal reduction** of the original game

	L
D	5, 1

IESDS

- If you eliminate a strictly dominated strategy, the reduced game has the same Nash equilibria as the original one
- Thus

  {Nash equilibria of the original game}

  = {Nash equilibria of the maximally reduced game}

	L	*R*
~~*U*~~	~~3, 3~~	~~0, 5~~
D	5, 1	1, 0

- Use this technique to simplify finding Nash equilibria
  - Look for Nash equilibria on the maximally reduced game

	L	~~*R*~~
D	5, 1	~~1, 0~~

- In the example, we ended up with a single cell
  - The single cell *must* be a unique Nash equilibrium in all three of the games

	L
D	5, 1

IESDS

- Even if s_i isn't strictly dominated by a pure strategy, it may be strictly dominated by a mixed strategy
- **Example**: the three games shown at right
  - 1st game:
    - R is strictly dominated by L (and by C)
    - Eliminate it, get 2nd game
  - 2nd game:
    - Neither *U* nor *D* dominates *M*
    - But {(½, *U*), (½, *D*)} strictly dominates *M*
      - This wasn't true before we removed *R*
    - Eliminate it, get 3rd game
  - 3rd game is maximally reduced

	L	*C*	*R*
U	3, 1	0, 1	0, 0
M	1, 1	1, 1	5, 0
D	0, 1	4, 1	0, 0

	L	*C*
U	3, 1	0, 1
M	1, 1	1, 1
D	0, 1	4, 1

	L	*C*
U	3, 1	0, 1
D	0, 1	4, 1

Correlated Equilibrium: Pithy Quote

If there is intelligent life on other planets, in a majority of them, they would have discovered correlated equilibrium before Nash equilibrium.

----Roger Myerson

Correlated Equilibrium: Intuition

- Not every correlated equilibrium is a Nash equilibrium but every Nash equilibrium is a correlated equilibrium
- We have a **traffic light**: a fair randomizing device that tells one of the agents to go and the other to wait.
- Benefits:
  - easier to compute than Nash, e.g., it is polynomial-time computable
  - fairness is achieved
  - the sum of social welfare exceeds that of any Nash equilibrium

Correlated Equilibrium

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

- Recall the mixed-strategy equilibrium for the Battle of the Sexes
  - s_w = {(2/3, Opera), (1/3, Football)}
  - s_h = {(1/3, Opera), (2/3, Football)}
- This is “fair”: each agent is equally likely to get his/her preferred activity
- But 5/9 of the time, they’ll choose different activities => utility 0 for both
  - Thus each agent’s expected utility is only 2/3
  - We’ve required them to make their choices independently
- Coordinate their choices (e.g., flip a coin) => eliminate cases where they choose different activities
  - Each agent’s payoff will always be 1 or 2; expected utility 1.5
- Solution concept: **correlated** equilibrium
  - Generalization of a Nash equilibrium

Correlated Equilibrium

- Let G be an n-agent game
- Let $v_1, \ldots, v_n$ be random variables, one for each agent
  - For each i, let D_i be the domain (the set of possible values) of v_i
- Let π be a joint distribution over $v_1, \ldots, v_n$
  - $\pi(d_1, \ldots, d_n) = \Pr[v_1=d_1, \ldots, v_n=d_n]$
- "Nature" uses π to choose values $\mathbf{d} = (d_1, \ldots, d_n)$ for $\mathbf{v} = (v_1, \ldots, v_i)$
- "Nature" tells each agent i the value of v_i (privately)
  - An agent can condition his/her action on the value of v_i
  - An agent's strategy is a deterministic mapping $\sigma_i : D_i \rightarrow A_i$ (note that we might have $\sigma_i(d_1) = \sigma_i(d_2)$ for d_1 not equal to d_2)
    - As book says mixed strategies wouldn't give any greater generality
  - A strategy profile is $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$
- The games we've been considering before now are a degenerate case in which the random variables $v_1, \ldots, v_n$ are independent

Correlated Equilibrium

- G is an n-player game
- $\mathbf{v} = (v_1, \ldots, v_n)$ are random variables with domains $\mathbf{D} = (D_1, \ldots, D_n)$
  - Joint distribution $\pi(\mathbf{d}) = \pi(d_1, \ldots, d_n) = \Pr[v_1{=}d_1, \ldots, v_n{=}d_n]$
- $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$ is a strategy profile
  - Each strategy σ_i is a mapping from D_i to A_i
- Then the expected utility for agent i is

$$u_i(\boldsymbol{\sigma}) = \textstyle\sum_{\mathbf{d}} \pi(\mathbf{d})\, u_i(\boldsymbol{\sigma}(\mathbf{d})),$$

i.e.,
$$u_i(\sigma_1, \ldots, \sigma_n) = \sum_{d_1, \ldots, d_n} \pi(d_1, \ldots, d_n)\, u_i\big(\sigma_1(d_1), \ldots, \sigma_n(d_n)\big)$$

- $(\boldsymbol{v}, \pi, \boldsymbol{\sigma})$ is a **correlated equilibrium** if for every agent i and strategy σ_i',

$$u_i(\boldsymbol{\sigma}) \ \geq\ u_i(\sigma_i', \boldsymbol{\sigma}_{-i})$$

i.e.,
$$u_i(\sigma_1, \ldots, \sigma_{i-1}, \sigma_i, \sigma_{i+1}, \ldots, \sigma_n) \ \geq\ u_i(\sigma_1, \ldots, \sigma_{i-1}, \sigma_i', \sigma_{i+1}, \ldots, \sigma_n)$$

Correlated Equilibrium

Theorem. For every Nash equilibrium $\mathbf{s} = (s_1, \ldots, s_n)$, there's a corresponding correlated equilibrium $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$

- "Corresponding" means they produce the same distribution on outcomes

Basic idea of the proof: for each i, set up v_i and σ_i to mimic s_i

- $v_1, \ldots, v_n$ independently distributed
- Each v_i has domain A_i and probability distribution s_i
- Each σ_i is the identity function (i.e., do the action that you're told to do)
- When the agents play the strategy profile $\boldsymbol{\sigma}$, the distribution over outcomes is identical to that under $\mathbf{s}$
- No agent i can benefit by deviating from σ_i, so $\boldsymbol{\sigma}$ is a correlated equilibrium

- But not every correlated equilib. is equivalent to a Nash equilib.e.g.,Battle of Sexes
- Intuitively, correlated equilibrium is computable in polynomial time since it has only a single randomization over outcomes, whereas in NE this is constructed as a product of independent probabilities.

Computing CE

$$\sum_{a \in A | a_i \in a} p(a) u_i(a) \geq \sum_{a \in A | a_i' \in a} p(a) u_i(a_i', a_{-i}) \quad \forall i \in N, \forall a_i, a_i' \in A_i$$

$$p(a) \geq 0 \quad \forall a \in A$$

$$\sum_{a \in A} p(a) = 1$$

- variables: $p(a)$; constants: $u_i(a)$
- we could find the social-welfare maximizing CE by adding an objective function

$$\text{maximize:} \quad \sum_{a \in A} p(a) \sum_{i \in N} u_i(a).$$