CMSC 474, Introduction to Game Theory

Dominant Strategies & Price of Anarchy

Mohammad T. Hajiaghayi

University of Maryland

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 ($s_1, \ldots, s_n$) 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

- 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
~~U~~	~~3, 3~~	~~0, 5~~
D	5, 1	1, 0

	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

The Price of Anarchy (PoA)

- In the Chocolate Game, recall that
  - (T3,T3) is the action profile that provides the best outcome for everyone

	T3	*T1*
T3	3, 3	0, 4
T1	4, 0	1, 1

  - If we assume each payer acts to maximize his/her utility without regard to the other, we get (T1,T1)
  - By choosing (T3,T3), each player could have gotten 3 times as much

	T3	*T1*
T3	3, 3	0, 4
T1	4, 0	1, 1

- Let's generalize "best outcome for everyone"

The Price of Anarchy

- ***Social welfare function***: a function $w(\mathbf{s})$ that measures the players' welfare, given a strategy profile $\mathbf{s}$, e.g.,
  - Utilitarian function: $w(\mathbf{s})$ = average expected utility
  - Egalitarian function: $w(\mathbf{s})$ = minimum expected utility

- ***Social optimum***: benevolent dictator chooses $\mathbf{s}^*$ that optimizes w
  - $\mathbf{s}^* = \arg\max_{\mathbf{s}} w(\mathbf{s})$

- ***Anarchy***: no dictator; every player selfishly tries to optimize his/her own expected utility, disregarding the welfare of the other players
  - Get a strategy profile $\mathbf{s}$ (e.g., a Nash equilibrium)
  - In general, $w(\mathbf{s}) \leq w(\mathbf{s}^*)$

Price of Anarchy (PoA) $= \max_{\mathbf{s}\text{ is Nash equilibrium}} w(\mathbf{s}^*) / w(\mathbf{s})$

- PoA is the most popular measure of inefficiency of equilibria.
- We are generally interested in PoA which is closer to 1, i.e., all equilibria are good approximations of an optimal solution.

The Price of Anarchy

- Example: the Chocolate Game
  - Utilitarian welfare function: $w(\mathbf{s})$ = average expected utility

	T 3	*T1*
T3	3, 3	0, 4
T1	4, 0	1, 1

- Social optimum: $\mathbf{s}^* = (T3,T3)$
  - $w(\mathbf{s}^*) = 3$

	T3	*T1*
T3	3, 3	0, 4
T1	4, 0	1, 1

- Anarchy: $\mathbf{s} = (T1,T1)$
  - $w(\mathbf{s}) = 1$
- Price of anarchy

  $= w(\mathbf{s}^*) / w(\mathbf{s}) = 3/1 = 3$

- What would the answer be if we used the egalitarian welfare function?

The Price of Anarchy

- Sometimes instead of ***maximizing*** a welfare function *w*, we want to ***minimize*** a cost function *c* (e.g. in Prisoner's Dilemma)
  - Utilitarian function: $c(\mathbf{s})$ = avg. expected cost
  - Egalitarian function: $c(\mathbf{s})$ = max. expected cost

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

- Need to adjust the definitions
  - ***Social optimum***: $\mathbf{s}^* = \arg\min_{\mathbf{s}} c(\mathbf{s})$
  - ***Anarchy***: every player selfishly tries to minimize his/her own cost, disregarding the costs of the other players
    - Get a strategy profile **s** (e.g., a Nash equilibrium)
    - In general, $c(\mathbf{s}) \geq c(\mathbf{s}^*)$
  - ***Price of Anarchy (PoA)*** = $\max_{\mathbf{s} \text{ is Nash equilibrium}} c(\mathbf{s}) / c(\mathbf{s}^*)$
    - i.e., the reciprocal of what we had before
    - E.g. in Prisoner's dilemma PoA= 3

Braess's Paradox in Road Networks

- Suppose 1,000 drivers wish to travel from *S* (start) to *D* (destination)
  - ➢ Two possible paths:
    - $S \rightarrow A \rightarrow D$ and $S \rightarrow B \rightarrow D$
  - ➢ The road from *S* to *A* is long: t = 50 minutes
    - But it's also very wide:
      t = 50 no matter how many cars
  - ➢ Same for road from *B* to *D*
  - ➢ Road from *A* to *D* is shorter but is narrow
    - Time = (number of cars)/25

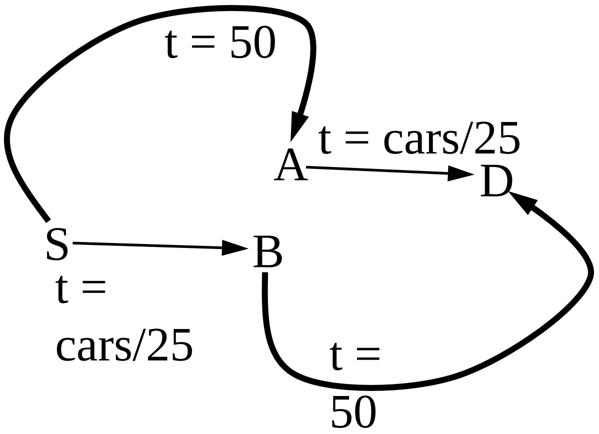

- Nash equilibrium:
  - ➢ 500 cars go through A, 500 cars through B
  - ➢ Everyone's time is 50 + 500/25 = 70 minutes
  - ➢ If a single driver changes to the other route then there are 501 cars on that route, so his/her time goes up

Braess's Paradox (cont'd)

- Add a *very* short and wide road from B to A:
  - 0 minutes to traverse, no matter how many cars
- Nash equilibrium:
  - All 1000 cars go S→B→A→D
  - Time for S→B is 1000/25 = 40 minutes
  - Total time is 80 minutes

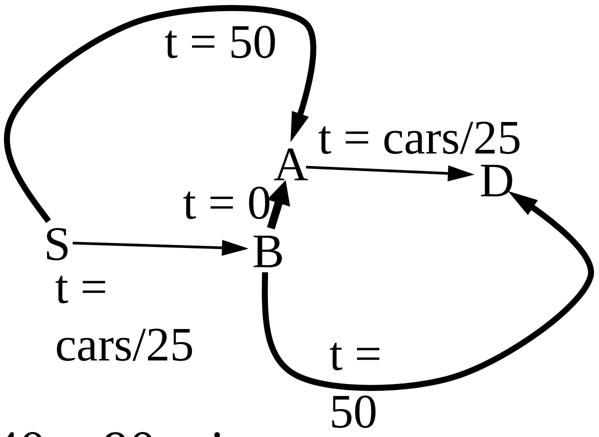

- To see that this is an equilibrium:
  - If driver goes *S*→*A*→*D*, his/her cost is 50 + 40 = 90 minutes
  - If driver goes *S*→*B*→*D*, his/her cost is 40 + 50 = 90 minutes
  - Both are dominated by S→B→A→D
- To see that it's the ***only*** Nash equilibrium:
  - For every traffic pattern, S→B→A→D dominates *S*→*A*→*D* and *S*→*B*→*D*
  - Choose any traffic pattern, and compute the times a driver would get on all three routes

The Price of Anarchy

- Example: Braess's Paradox
  - Utilitarian cost function: $c(\mathbf{s})$ = average expected cost

- Social optimum:
  - $\mathbf{s}^*$ = [500 go S→A→D; 500 go S→B→D]
  - $c(\mathbf{s}^*) = 70$
- Anarchy: $\mathbf{s}$ = [1000 drivers go S→B→A→D]
  - $c(\mathbf{s}) = 80$
- Price of anarchy = $c(\mathbf{s}) / c(\mathbf{s}^*) = 8/7$

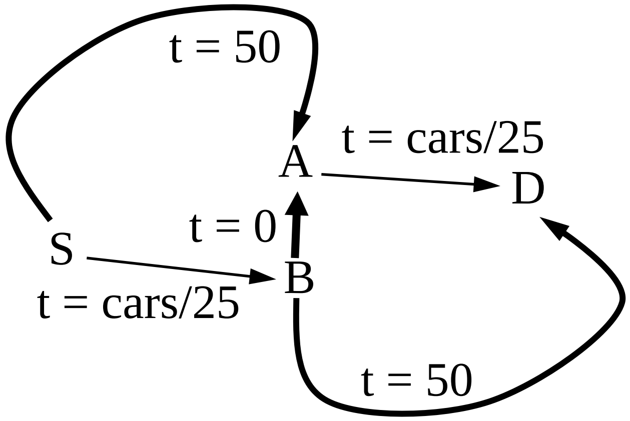

- What would the answer be if we used the egalitarian cost function?
- Note that when we talk about Price of Anarchy for Nash equilibria in general, we consider the **worst case** Nash equilibrium

Discussion

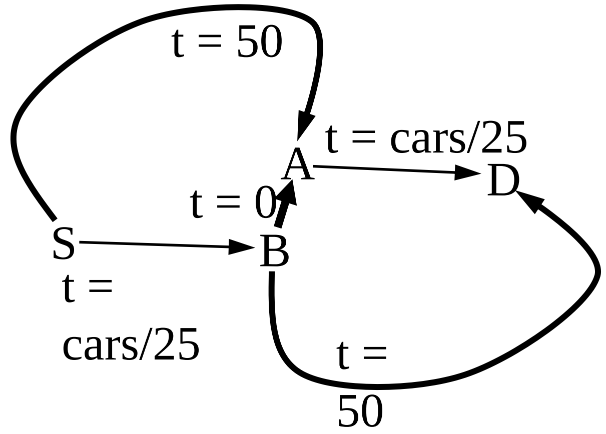

- In the example, adding the extra road increased the travel time from 70 minutes to 80 minutes
  - This suggests that carelessly adding road capacity can actually be hurtful
- But are the assumptions realistic?
- For A→B, *t* = 0 regardless of how many cars
  - Road length = 0? Then S→A and S→B must go to the same location, so how can their travel times be so different?
- For S→A, *t* = 50 regardless of how many cars
  - is it a 1000-lane road?
- For 1000 cars, does "*t* = cars/25" really mean 40 minutes per car?
  - The cars can't all start at the same time
  - If they go one at a time, could have 40 minutes total but 1/25 minute/car
- So can this really happen in practice?

Braess's Paradox in Practice

- 1969, Stuttgart, Germany – when a new road to city the center was opened, traffic got worse; and it didn't improve until the road was closed
- 1990, Earth day, New York – closing 42nd street improved traffic flow
- 1999, Seoul, South Korea – closing a tunnel improved traffic flow
- 2003, Seoul, South Korea – traffic flow was improved by closing a 6-lane motorway and replacing it with a 5-mile-long park
- 2010, New York – closing parts of Broadway has improved traffic flow
- Braess's paradox can also occur in other kinds of networks such as queuing networks or communication networks;
  - In principle, it can occur in Internet traffic though I don't have enough evidence to know how much of a problem it is
- Sources
  - http://www.umassmag.com/transportationandenergy.htm
  - http://www.cs.caltech.edu/~adamw/courses/241/lectures/brayes-j.pdf
  - http://www.guardian.co.uk/environment/2006/nov/01/society.travelsenvironmentalimpact
  - http://www.scientificamerican.com/article.cfm?id=removing-roads-and-traffic-lights
  - http://www.lionhrtpub.com/orms/orms-6-00/nagurney.html