CSC200: Lecture 11

- Today
  - Extensive form games (see last lecture slides)
  - Game theory on networks: Braess' paradox (Ch.8.1-8.2)
- Next few lectures
  - Auctions: Ch.9
  - Matching markets
  - Market-clearing prices
- Announcements
  - Quiz 2 on Friday, Oct.23: covers Chapter 4 (closures in a social-affiliation network)
    - see practice question on web page

Extensive Form (Dynamic) Games

- Normal form (matrix) games seem limited
  - Players move simultaneously and outcome determined at once
  - No observation, reaction, etc.
- Most games have a *dynamic* structure (turn taking)
  - Chess, tic-tac-toe, cards games, soccer, corporate decisions, markets…
  - See what "opponent" does before making move
  - Sometimes you only see partial information (won't discuss this)

- Example: Rogers, Telus competing for market in two remote areas
  - Each firm, Rogers and Telus, can tackle one area only
  - Total revenue in Area A: 12, Area B: 6
  - If firm is alone in one area, get all of that area's revenue
  - If both firms target same area, "first mover" gets 2/3, second 1/3
  - Rogers prepped: makes first move, Telus chooses *after* Rogers
  - Critical: Telus *observes* Rogers' choice before moving!

Game Tree (Extensive Form)

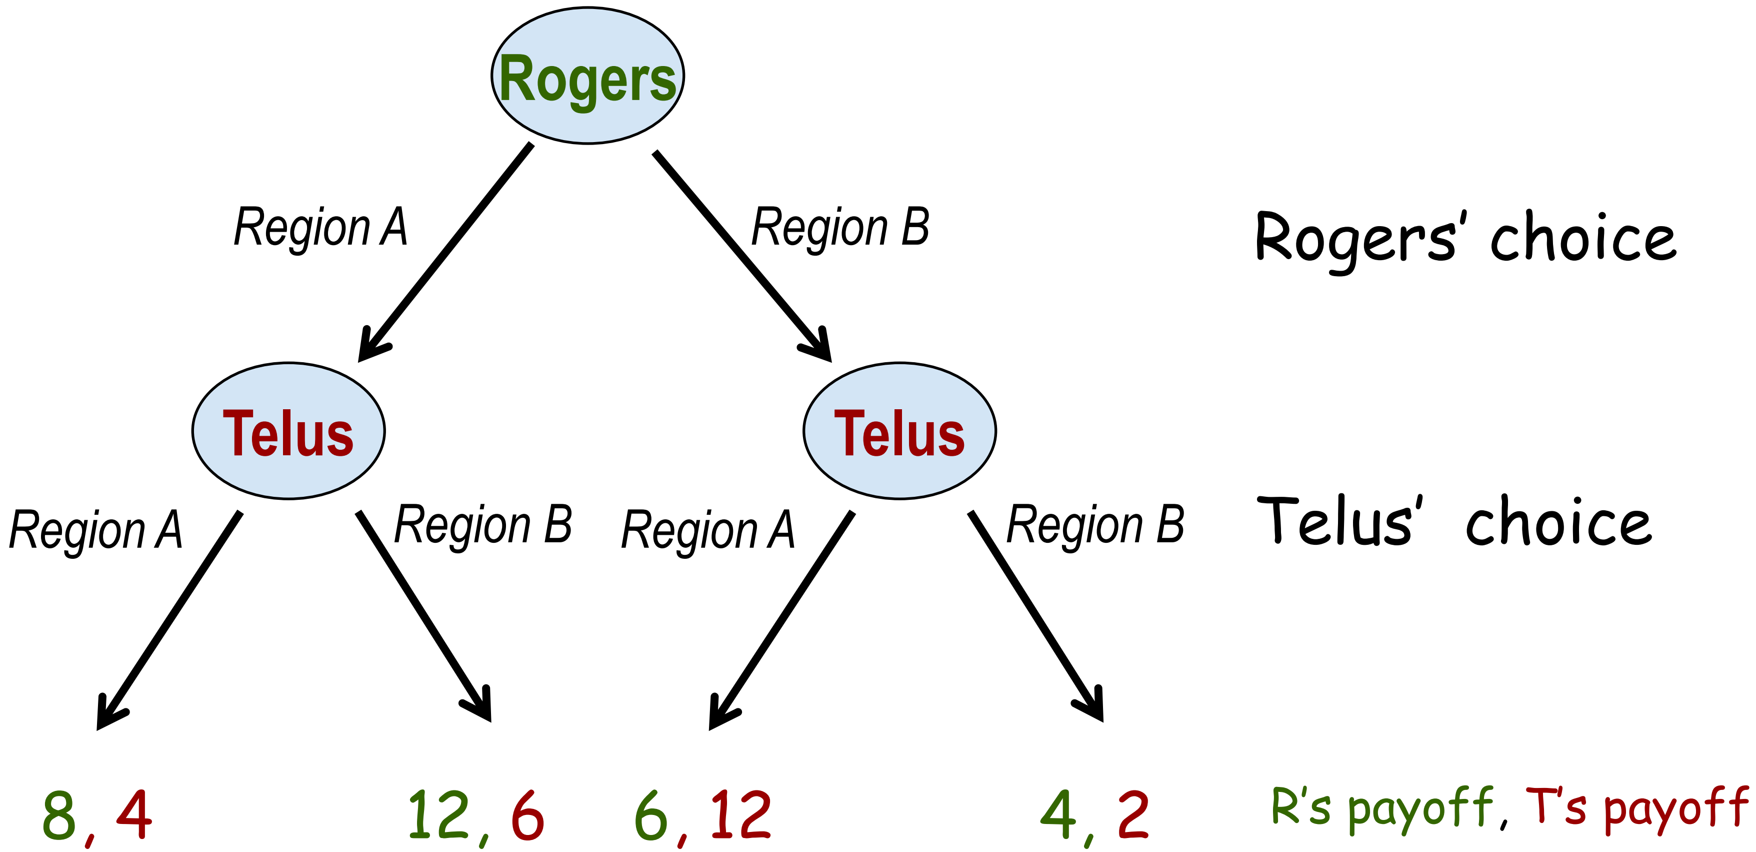

Strategies Adopted

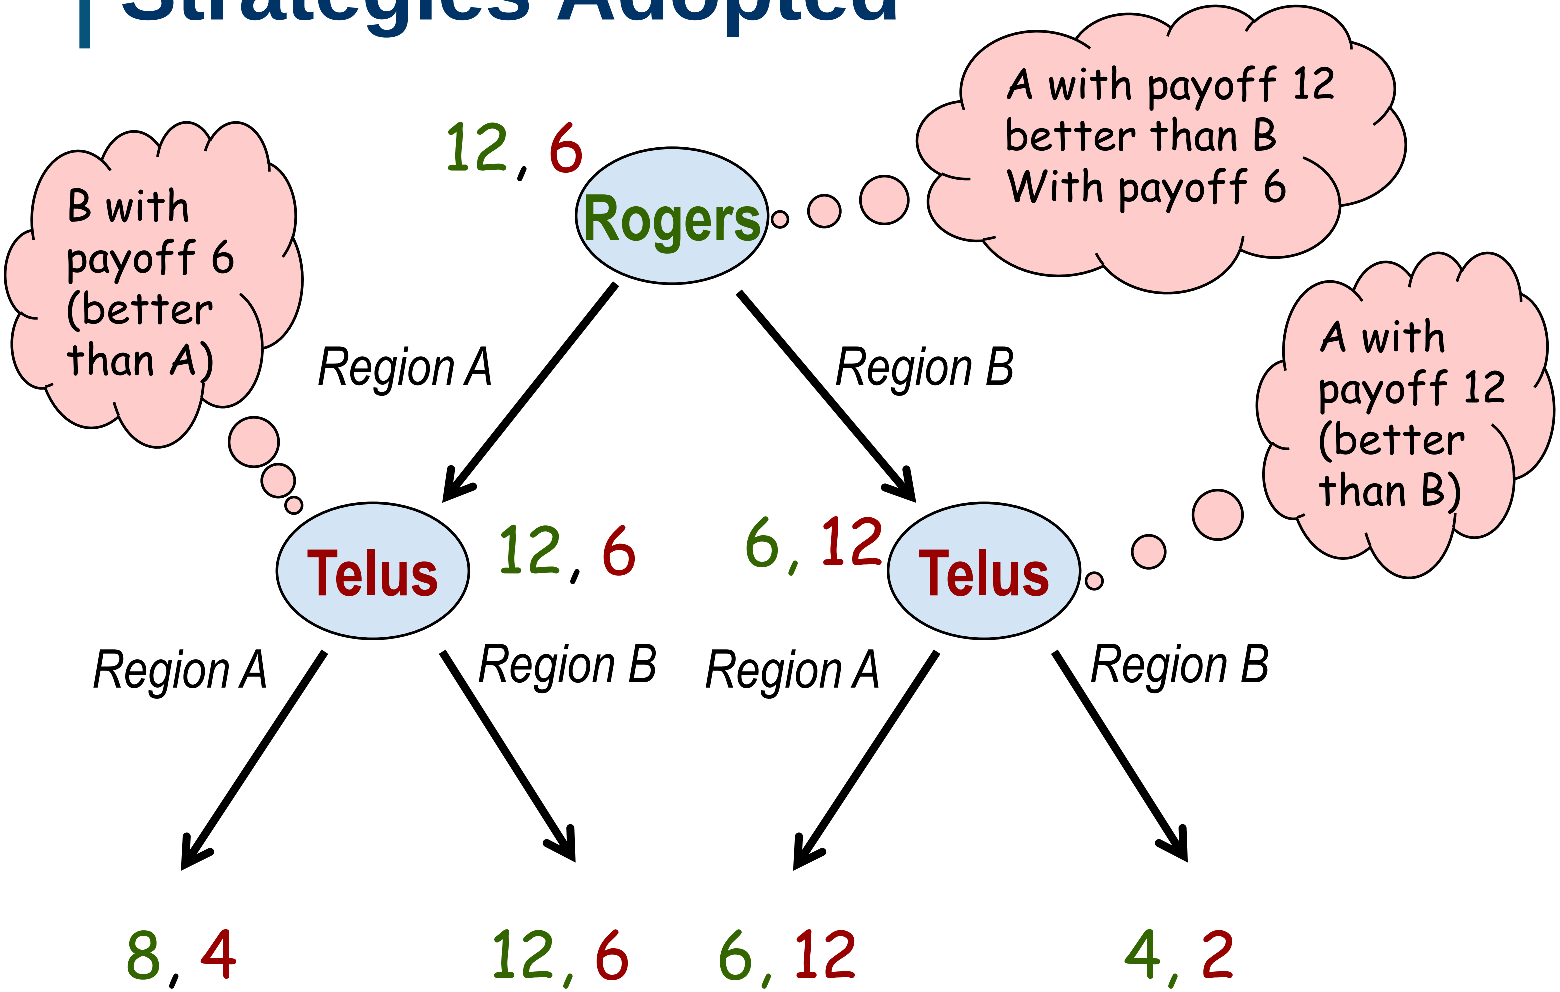

Backward Induction

- Start from choice nodes at bottom of tree
  - Player at that node chooses best move
  - E.g., Telus chooses B at node “Rogers did A”, A at node “Rogers did B”
  - This dictates which terminal node (payoffs) will be reached
  - We can now assign payoffs to that node (from chosen terminal node)
- Work up the tree: compute choices at other nodes once choices/payoffs at all child nodes made
  - Player at that node chooses best move
  - E.g., Rogers chooses A at root node
  - Dictates which child (“Rogers did A”) will be reached, hence which payoff
- At end of procedure:
  - Each choice node labelled with action/choice and payoff vector
  - Payoff for the game is the payoff vector at the root node
  - Path through tree given by choices, tells us how the game will unfold

Tic-Tac-Toe Game Tree

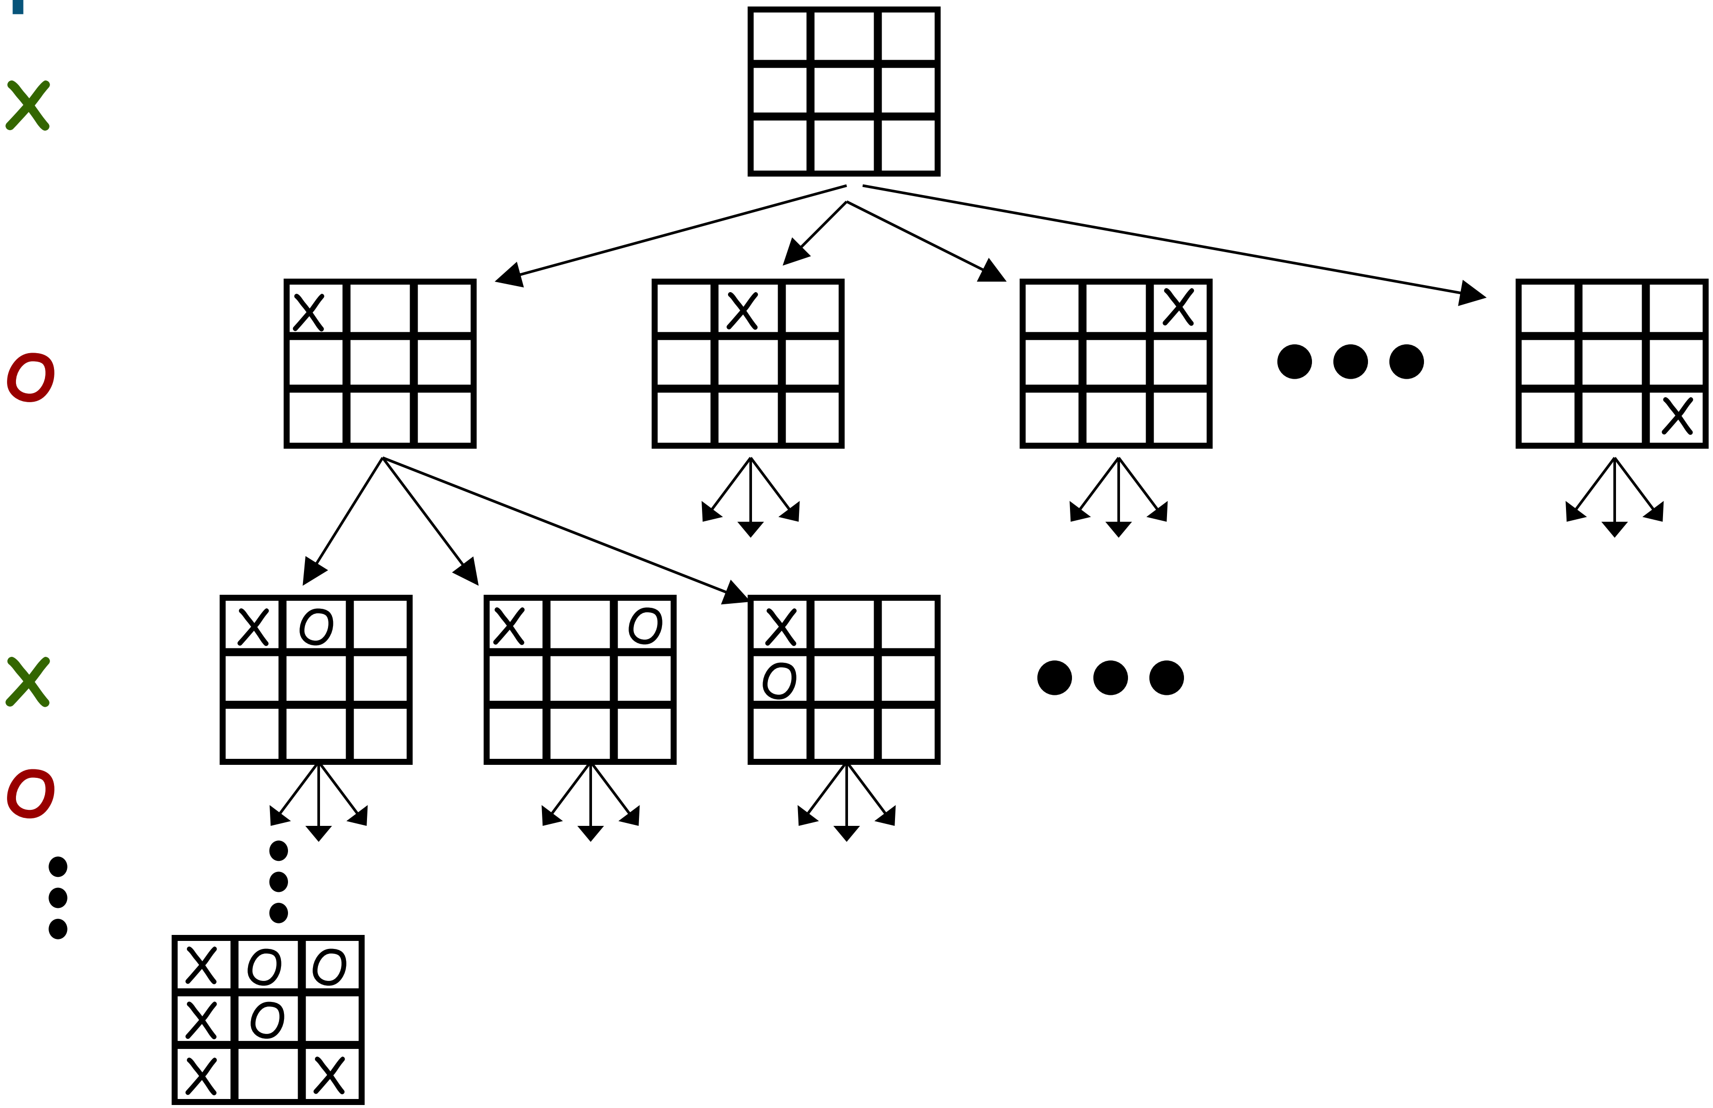

Payoff: X=1, O=-1

Conversion to Normal Form

- Extensive form game can be converted to normal (matrix) form by extending the space of strategies
- Notice Telus has *four* strategies, not just *two*
  - Two "independent" choices of A or B: "if R did A", and "if R did B"
    - *AA,AB: A if R does A, A if R does B (i.e., do A no matter what)*
    - *AA, BB: A if R does A, B if R does B*
    - *BA, AB: B if R does A, B if R does B*
    - *BA, BB: B if R does A, B if R does B (i.e., do B no matter what)*
- Matrix game has two "identical" NE: *A/BA,AB* and *A/BA,BB*
  - Difference: what Telus does at node that won't be reached
  - Same outcome as backward induction

		Telus			
		AA, AB	AA, BB	BA, AB	BA, BB
Rogers	A	$8, 4$	$8, 4$	$12, 6$	$12, 6$
	B	$6, 12$	$4, 2$	$6, 12$	$4, 2$

Conversion Hides Information

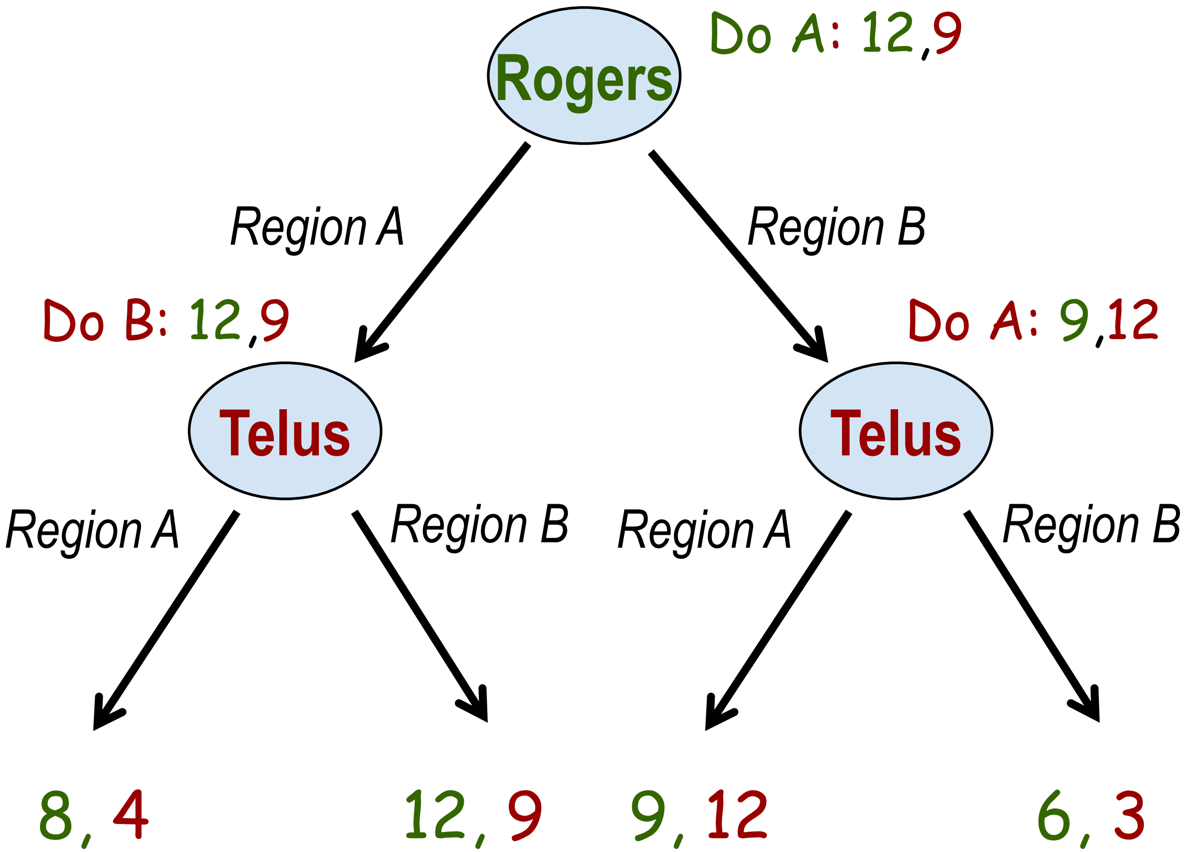

- Change payoffs so that region B is more valuable (9 instead of 6)
- Backward induction gives same choice of strategies (Telus gets higher payoff for region B, but otherwise the same)

Conversion Hides Information

- Notice what happens in normal form
  - The two Nash eq. ***A/BA,AB*** and ***A/BA,BB*** remain in place
  - But a new one emerges: ***B/AA,AB***
- New Nash equilibrium
  - Telus threatens to move into A if Rogers does
  - This is enough to make Rogers move to B (alone): full amount (9) of smaller payoff is better than 2/3-share (8) of larger payoff
- Why doesn't this arise in tree (backward induction)?
  - The threat is not credible (Telus sacrifices its own payoff)

Telus

Rogers	AA, AB	AA, BB	BA, AB	BA, BB
A	8, 4	8, 4	12, **9**	12, **9**
B	**9**, 12	**6, 3**	**9**, 12	**6, 3**

Credible vs. Non-credible Threats

- Normal form supports equilibria where second player can threaten to make a move that hurts their own payoff in order to prevent first player from taking an earlier move
  - Threat is enough to prevent the move in normal NE
  - Hence the self-inflicted damage to player will not occur

- What if first player takes the move anyway, would second player do what it threatened?
  - *The threat is not credible*, not supported by backward induction
  - We call equilibria in the game tree: *subgame perfect equilibria*
  - SPE must be Nash eq in normal form, but not all NE are SPE!

What about Pre-commitment?

- If player *could* pre-commit, the “threatening NE” is valid
- How do you make an irrevocable (or costly) commitment?
  - sign a contract
  - set loose a computer program
  - doomsday machine (Dr. Strangelove)
  - Important: first player must *know about* pre-commitment

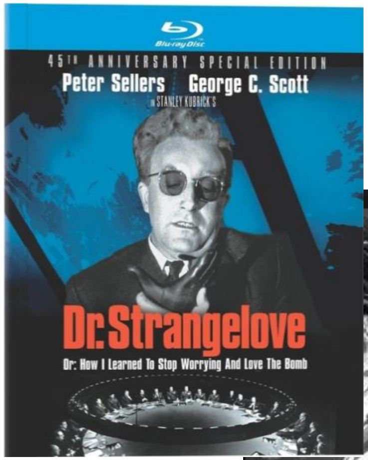

- But this is not the same game
  - the space of strategies (precommitment moves), and the ordering of action are different than in the original game
- So subgame perfect equilibria are the most natural form of NE for dynamic (extensive form) games

Exercise: Voting Game

- Take our three politicians who want a raise
  - all want a raise, all want to vote no; but want raise more!
  - 2 out of 3 votes needed to pass
  - now make the game dynamic: politician 1 votes first, politician 2 votes second, politician 3 votes thirds
- Do the following
  - draw the game tree
  - identify the strategies available to each player (how many does each have?)
  - construct the normal form representation
  - what is/are its subgame perfect equilibria?
  - what is/are Nash equilibria in normal form?
  - what threats could politician 1, 2 or 3 make?

A Traffic Network

- Let's look at a game with network structure in action space
  - *Stylized highway network:* travel time varies with traffic
  - if *x* cars on a link (segment) travel time is as labeled
    - varies on A—C and D—B but fixed on A—D and C—B

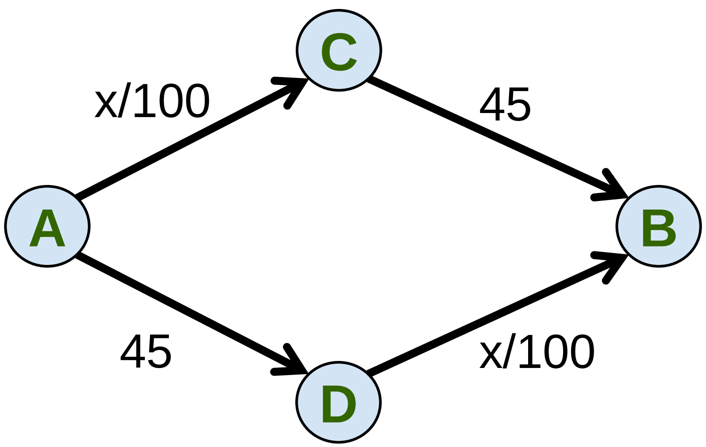

Example: Suppose 4000 drivers must get from A to B each morning. 3000 take route A-C-B, and 1000 take route A-D-B:

- route A-C-B (C): 75 mins
- route A-D-B (D): 55 mins

Traffic Flow in Equilibrium

- Suppose 4000 cars travel from A to B each morning
  - *What is equilibrium traffic flow?*
- Model as a game with 4000 players
  - each driver can choose route A—C—B or A—D—B
  - each driver desires minimum *personal* travel time: payoff is *-mins*

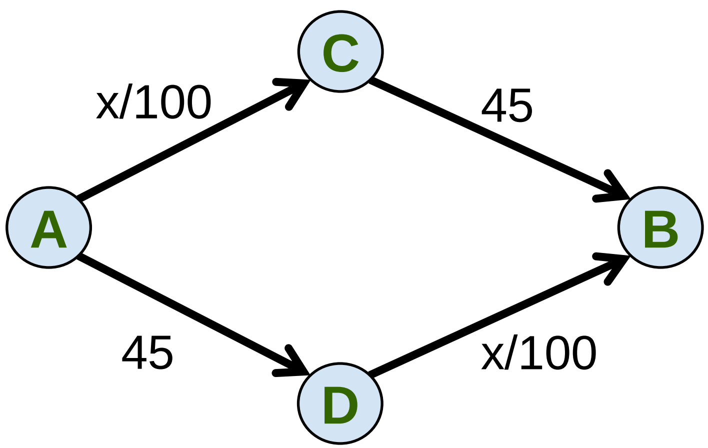

Many Nash equilibria!

But all are "equivalent":
- 2000 drivers take C
- 2000 drivers take D
- all travel times: 65 mins

Traffic Flow in Equilibrium

- (Any) profile <*2000 C, 2000 D*> is a NE
  - each route is equally fast: 65 mins, no incentive to switch
  - in fact, if a driver switches (e.g., from *C* to *D*): her travel time goes up from 65 mins to 65.01 mins
  - How many NE? $\binom{4000}{2000} \approx 1.6 \times 10^{1202}$

- Why is <*n C, 4000-n D*> not a NE if $n \neq 2000$?
  - Any driver on slower route will want to switch to faster route

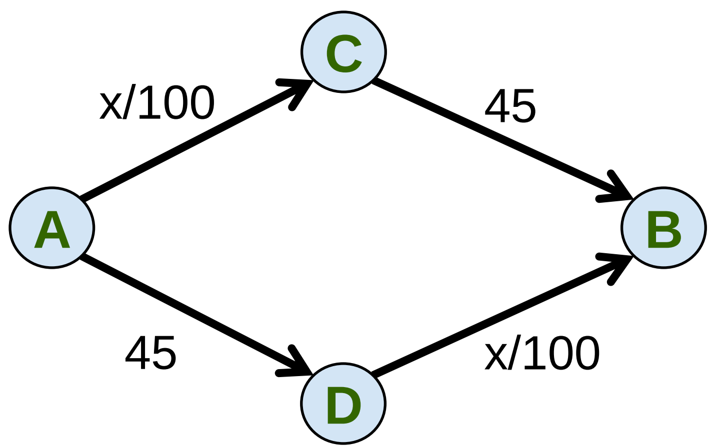

Social Optimality

- The Nash equilibrium (we'll speak as if there's one, since they are all interchangeable) is in fact *socially optimal* and it is the only socially optimal way to arrange traffic
  - it minimizes the population's total (equiv. average) commute time

  - in the NE (2000/2000): everyone has 65 minute commute time
  - if you shift balance to (2001 C, 1999 D):
    - *1999* drivers see commute time drop by 0.01 (64.99 mins)
    - but ***2001*** drivers see commute time rise by 0.01 (65.01 mins)
    - total commute time goes up by 0.02 mins
  - (2100 C, 1900D): total increase of 200 mins
  - (3000 C, 1000D): total increase of 20,000 mins (about 2 weeks)
  - (4000 C, 0D): total increase of 80,000 mins (almost 2 months)

How might NE emerge in practice?

- With 10^{1200} NE: the *ultimate* equilibrium selection problem!
  - 4000 drivers didn't call each other up this AM and coordinate
  - Iterative process? Try out a route… if it's fast you stick, if it's slow you switch?
    - Suppose 4000 start with C: what do they do the next day?
    - All switch to D! Then all switch back to C, …
  - More likely, a probabilistic process, some people more amenable to switching than others… and the slower it is the more likely you are to switch… over time after a process of gradual adjustment leads to something that is an approximate NE
- What's nice about this: self-organization based on self-interest makes everyone better off, indeed as well-off as possible, since it maximizes social welfare

Braess' Paradox

- Our new mayor treats Torontonians to reward them for her/his recent victory
  - adds a new superhighway to reduce everyone's commute time
  - link with much smaller time (we'll call it zero)
  - what happens to traffic patterns?

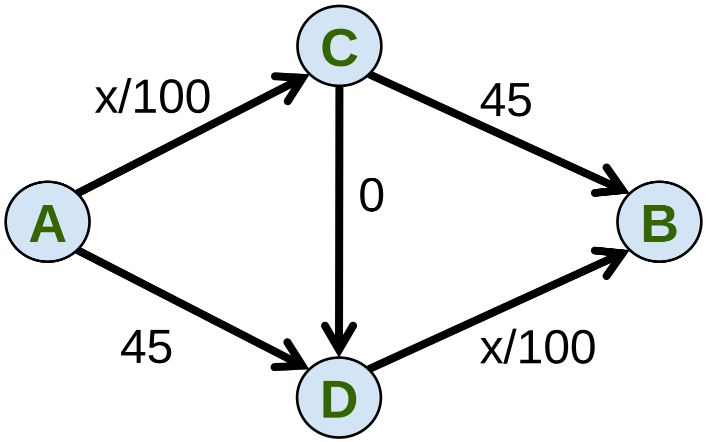

Braess' Paradox

- Unique NE in new game:
  - *everyone* takes A-C-D-B; commute increases from 65 to 80 mins!
- Why?
  - the *most* links A-C and D-B can take is 40 mins (all 4000 on them)
  - so A-C-D always faster than A-D; and D-B always faster than C-B
  - so A-C-D-B is *dominant* for every driver
- "Paradox": adding capacity slowed everyone down
  - named for discoverer (Diettrich Braess, 1968)
  - *observed in real traffic situations (usually closures)*

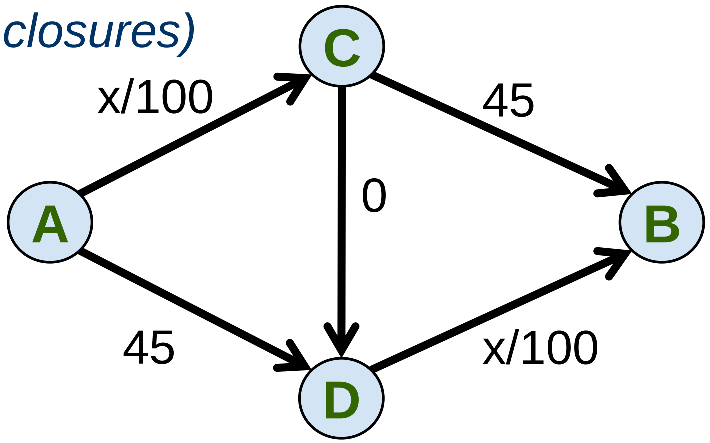

Why does it happen?

- Before new link:
  - all routes from A-B required one 45 min link
  - but *two* of them, so traffic split, easing congestion on A-C and D-B
- After new link:
  - everyone can avoid 45 min link
  - but *only one way to do so*: draws all traffic through C-D
  - leaves both 45 min links (A-D, C-B) unused!

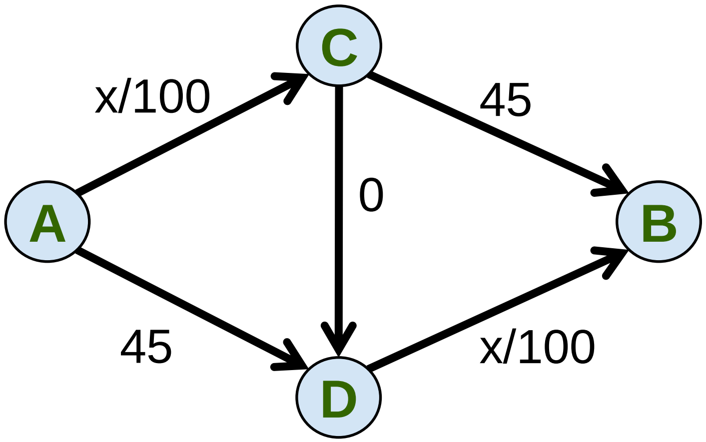

How useful could the new link be?

- What is socially optimal solution?
  - 2250 take A-C; 1750 take A-D;
  - 500 take C-D
  - 1750 take C-B; 2250 take D-B
  - *reasoning: find split of 4000 drivers over links A-C, A-D that minimizes their average time to C or D (2250/1750 split, minimize quadratic function)*
  - *by symmetry, 2250/1750 on D-B, C-B minimizes avg. time getting 4000 drivers from C or D to B*
  - *but we can swap 500 people from C-D for free to achieve both*
- Socially optimal solution can only be "imposed"
  - will not emerge in equilibrium
  - nobody will willingly take A-D *or* C-B

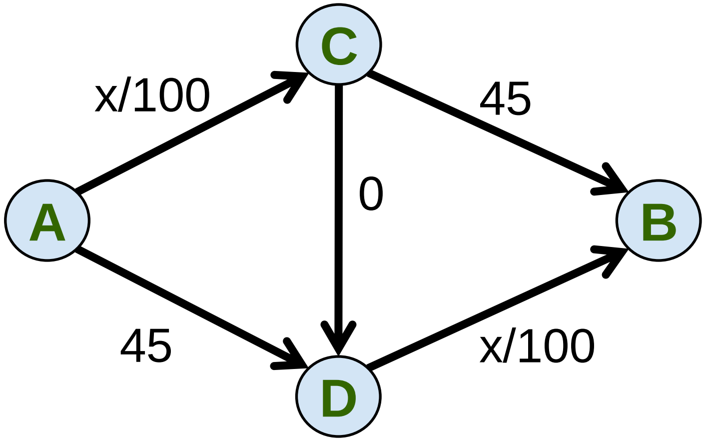

How useful could the new link be?

- Compare new link with ***imposed*** social optimum to no link
  - Without new link: everyone takes 65 mins
  - With new link, social optimum average commute time is 64.69 mins
    - 500 have A-C-D-B (45 mins)
    - 3500 have ADB or ACB (67.5 mins)
    - Total time saved: 1250 mins (avg. 19 sec per driver)
  - Not Pareto improving: 500 people save a lot at expense of 3500 others
    - Which 3500 will put up with it?
- New link, if you *can't* impose social optimum:
  - Average increases from 64.69 to 80 mins
  - So allowing people to act in their own interests (equilibrium) causes society (and in this case, every member of society) to suffer (aka *Tragedy of the Commons*)

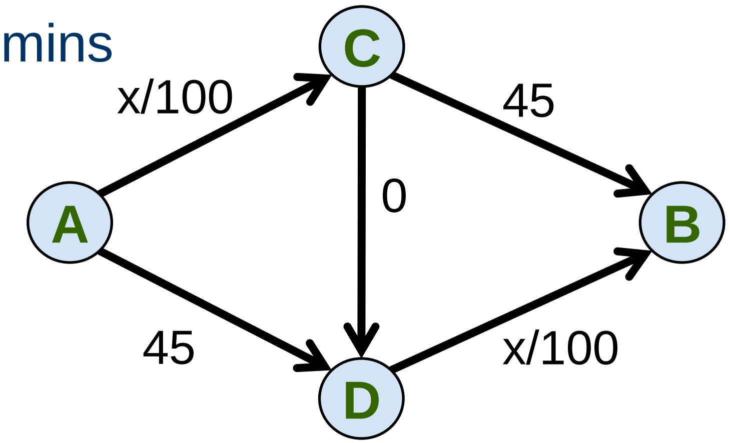

Price of Anarchy

- Computer scientists have studied the following question
  - consider a game (usually of a *specific form*) and ask what social cost we can derive from the *socially optimal profile* (OPT)
  - consider *social cost of the NE* of the same game (SCNE)
  - what is the ratio SCNE:OPT?
- In other words, how much societal benefit do you sacrifice by letting everyone choose their own actions?
  - sometimes called the price of anarchy
- For networks with *linear* cost functions:
  - known that one can lose *at most half* of optimal social welfare (or that OPT is no more than twice what you achieve in (worst) NE)
  - also known that by adding a link, the change in the SCNE is at most 4/3 (i.e., new NE with new link causes avg. time to increase by a factor of at most 4/3 relative to the old NE in the network prior to adding the link)
- Read 8.3 (Advanced) for more technical discussion if interested