Modeling Network Traffic using Game Theory

Part III: Dynamics — Episode 7

Baochun Li
Department of Electrical and Computer Engineering
University of Toronto

Required reading: “Networks, Crowds, and Markets,” Chapter 8

Objective in this episode

- Sending packets through the Internet involves fundamentally game-theoretic reasoning —
  - Rather than simply choosing a route in isolation, individuals must evaluate routes in the presence of the congestion resulting from the decisions made by themselves and everyone else
- **Objective**: to develop models for network traffic using the game-theoretic ideas developed thus far

Consider a transportation network

- Edges are highways, and nodes are exits on highways
- Everyone wants to get from A to B
- Each edge has a designated travel time (in minutes) when there are x cars traveling on this highway

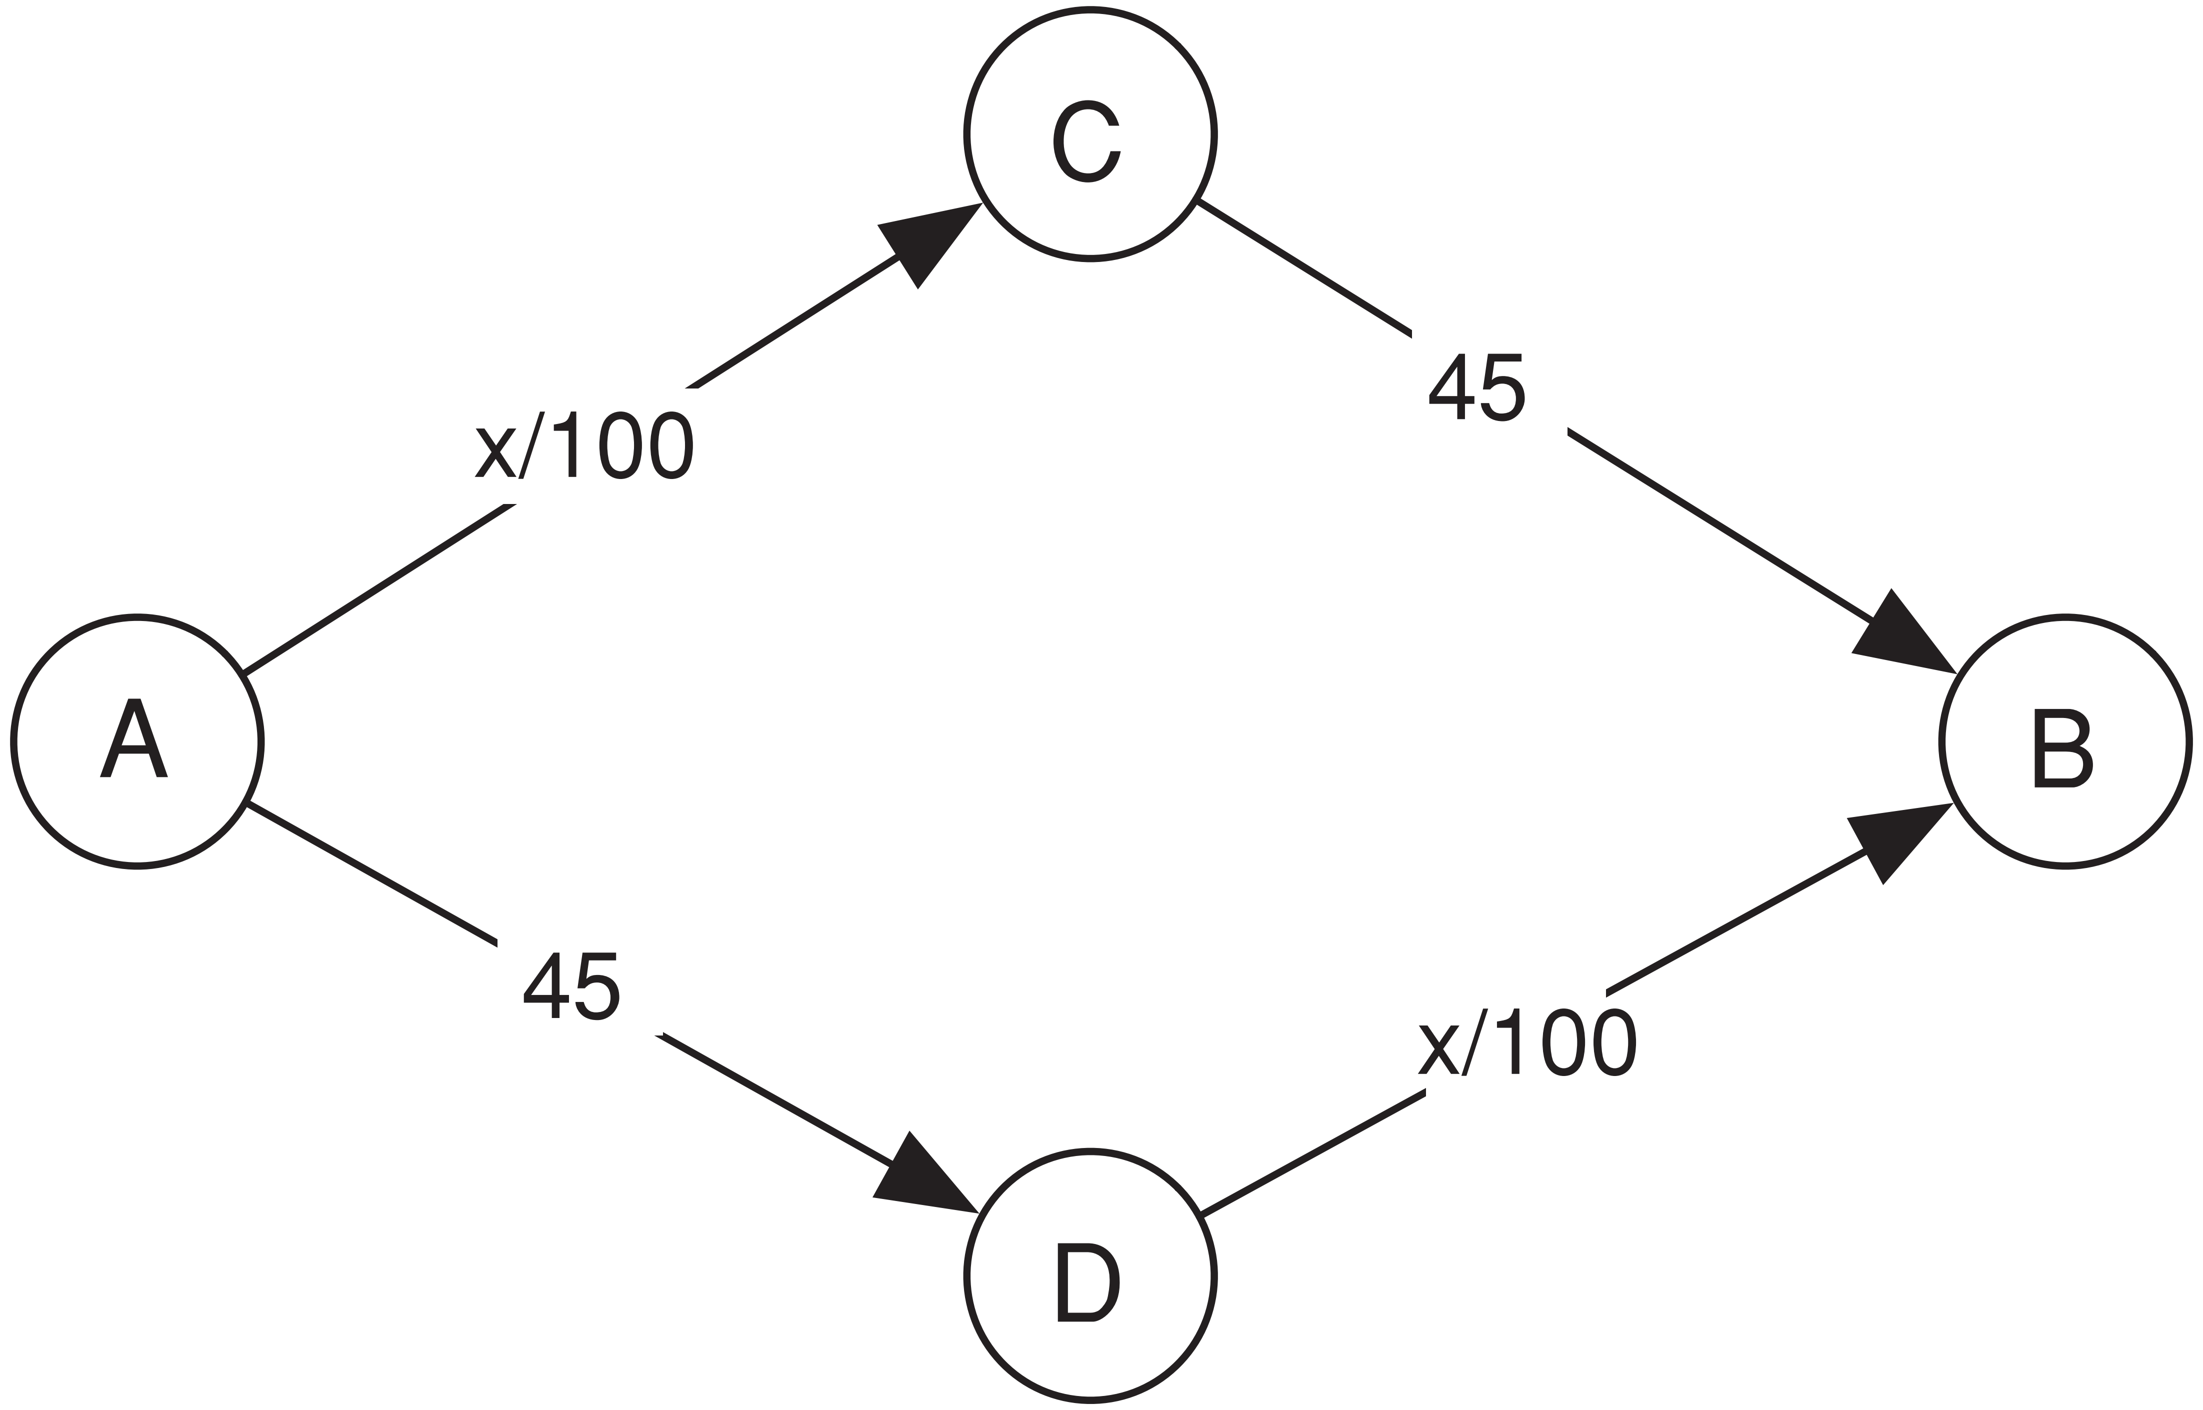

Consider a transportation network

- Suppose that 4,000 cars want to get from A to B
- If the cars divide up evenly between the two routes, the total travel time is 65 minutes
- So what do we expect to happen?

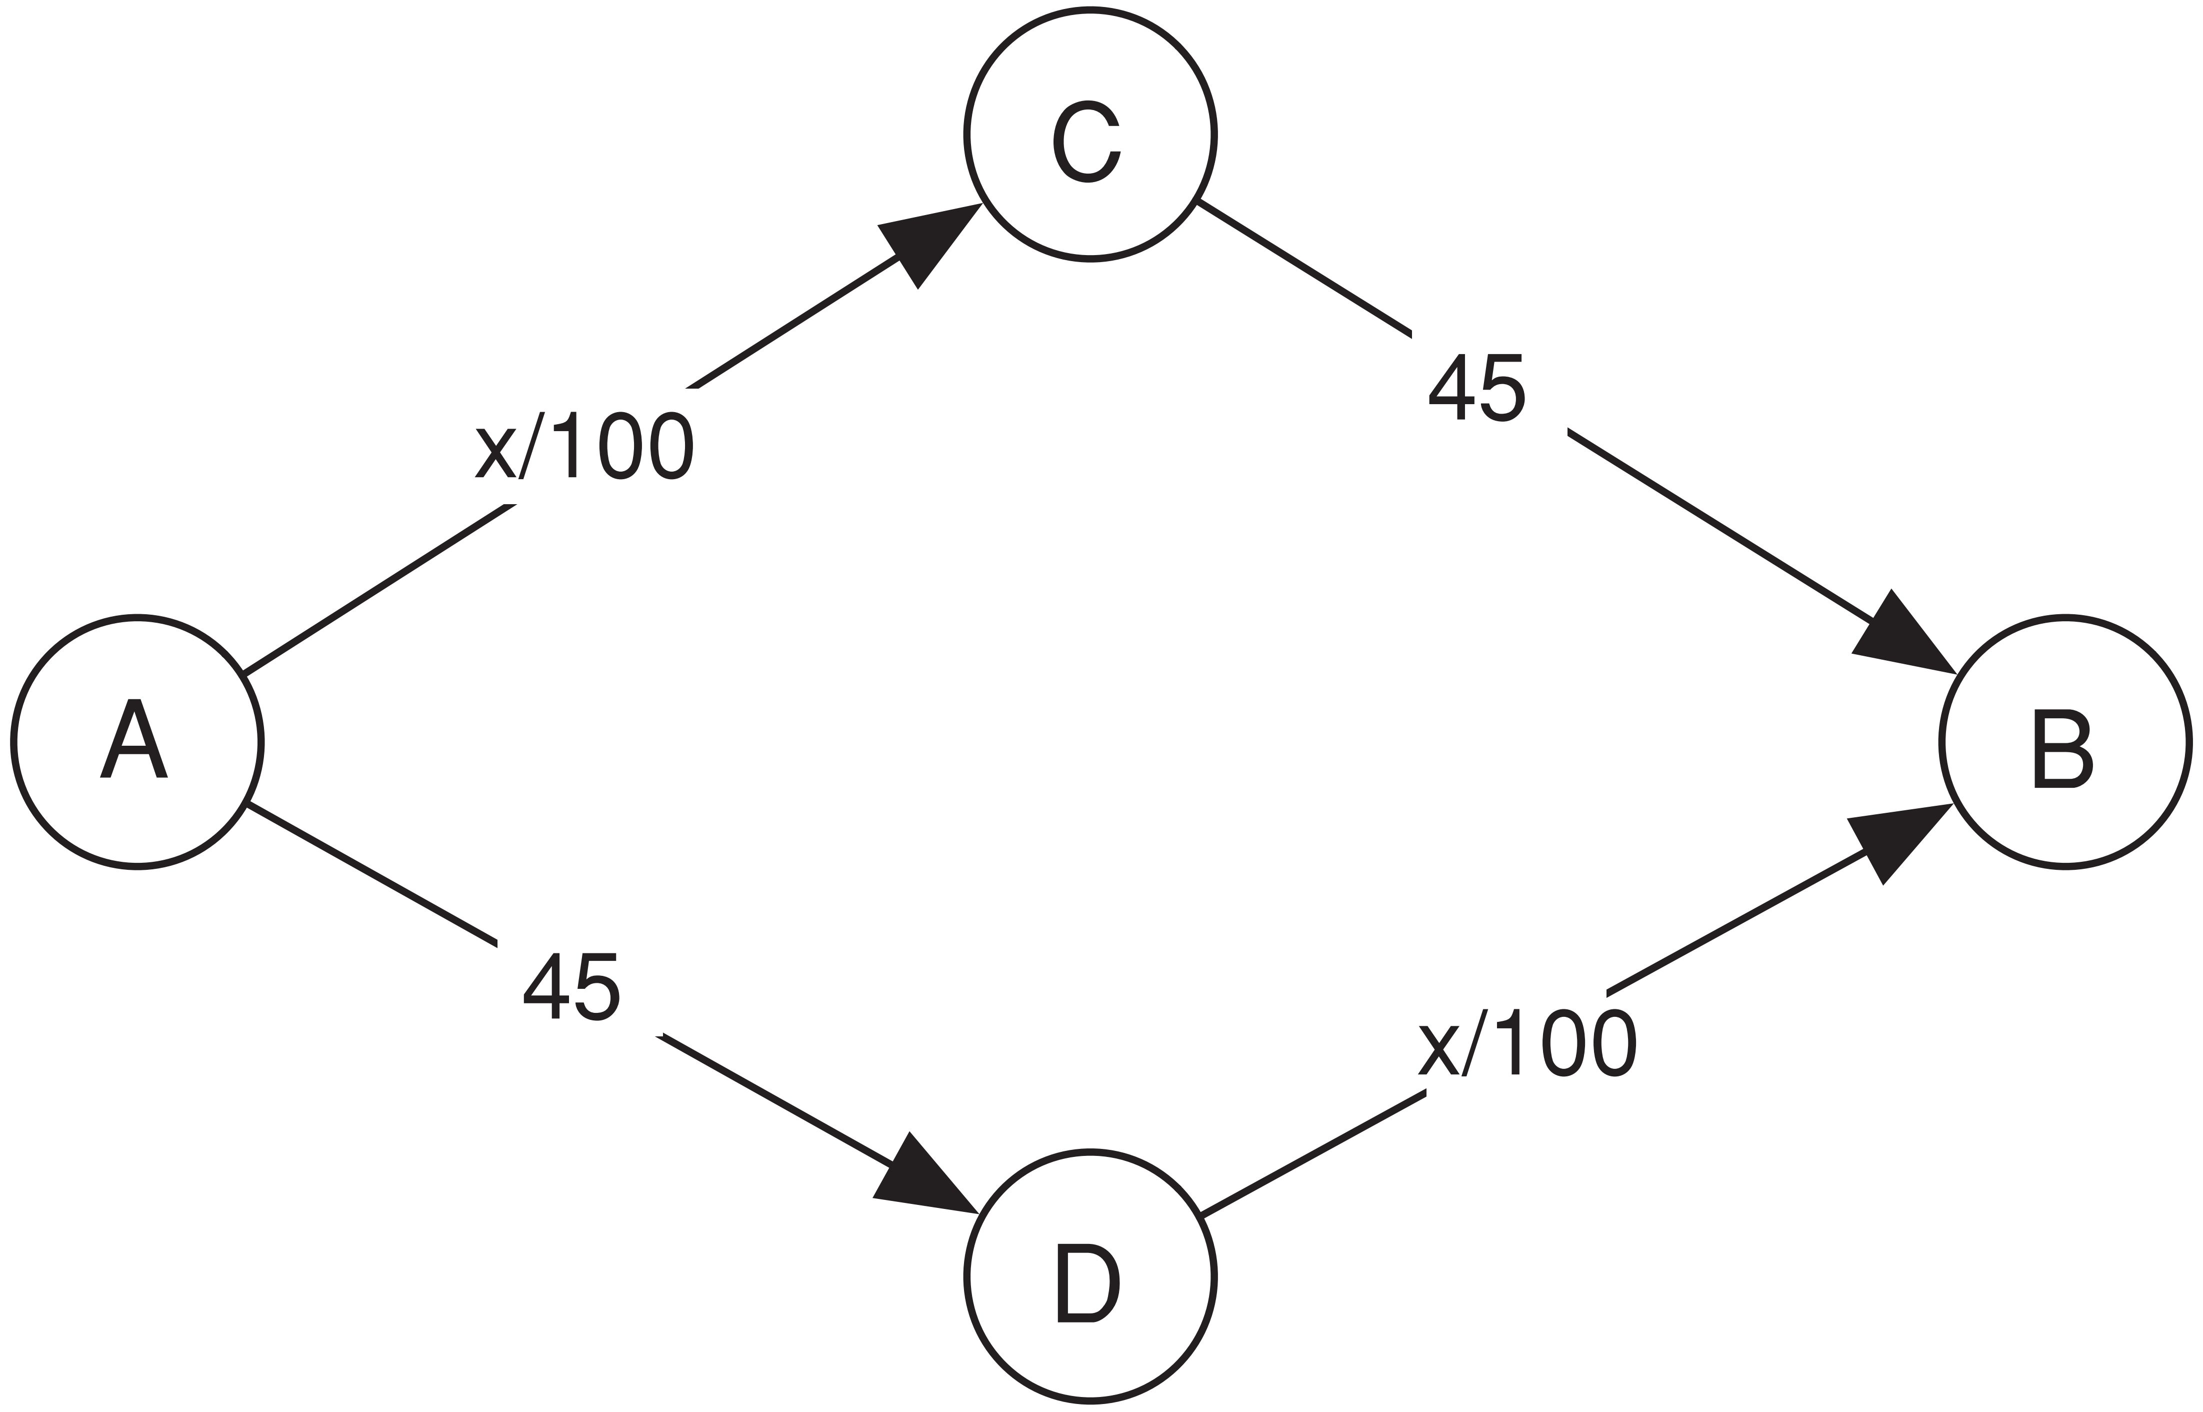

The traffic game

- The traffic model we've described is really a **game** in which the players correspond to the drivers, and each player's possible (two) strategies consist of the possible routes from A to B
  - The payoff for a player is the negative of his or her travel
  - We have focused primarily on games with two players, whereas the current traffic game will generally have an enormous number of players (4,000 in our example)
- There is no dominant strategy in this game
  - Either route has the potential to be the best choice for a player if all the other players are using the other route!

Equilibrium traffic

- The traffic game does have Nash equilibria —
  - Any list of strategies in which the drivers balance themselves evenly between the two routes (2,000 on each) is a Nash equilibrium, and these are the only Nash equilibria
- Why does equal balance yield a Nash equilibrium?
  - With an even balance between the two routes, no driver has an incentive to switch over to the other route

Equilibrium traffic

- Why do all Nash equilibria have equal balance?
  - Consider a list of strategies in which x drivers use the upper route and the remaining $4000 - x$ drivers use the lower route
  - Then if x is not equal to 2000, the two routes will have unequal travel times, and any driver on the slower route would have an incentive to switch to the faster one — not a Nash equilibrium!

Braess's Paradox

- Suppose that the city government decides to build a new, very fast highway from C to D — with a travel time of 0!
- It would stand to reason that people's travel time from A to B ought to get better after this edge from C to D is added
- There is indeed a unique Nash equilibrium in this new highway network
- But it leads to a **worse** travel time for everyone!

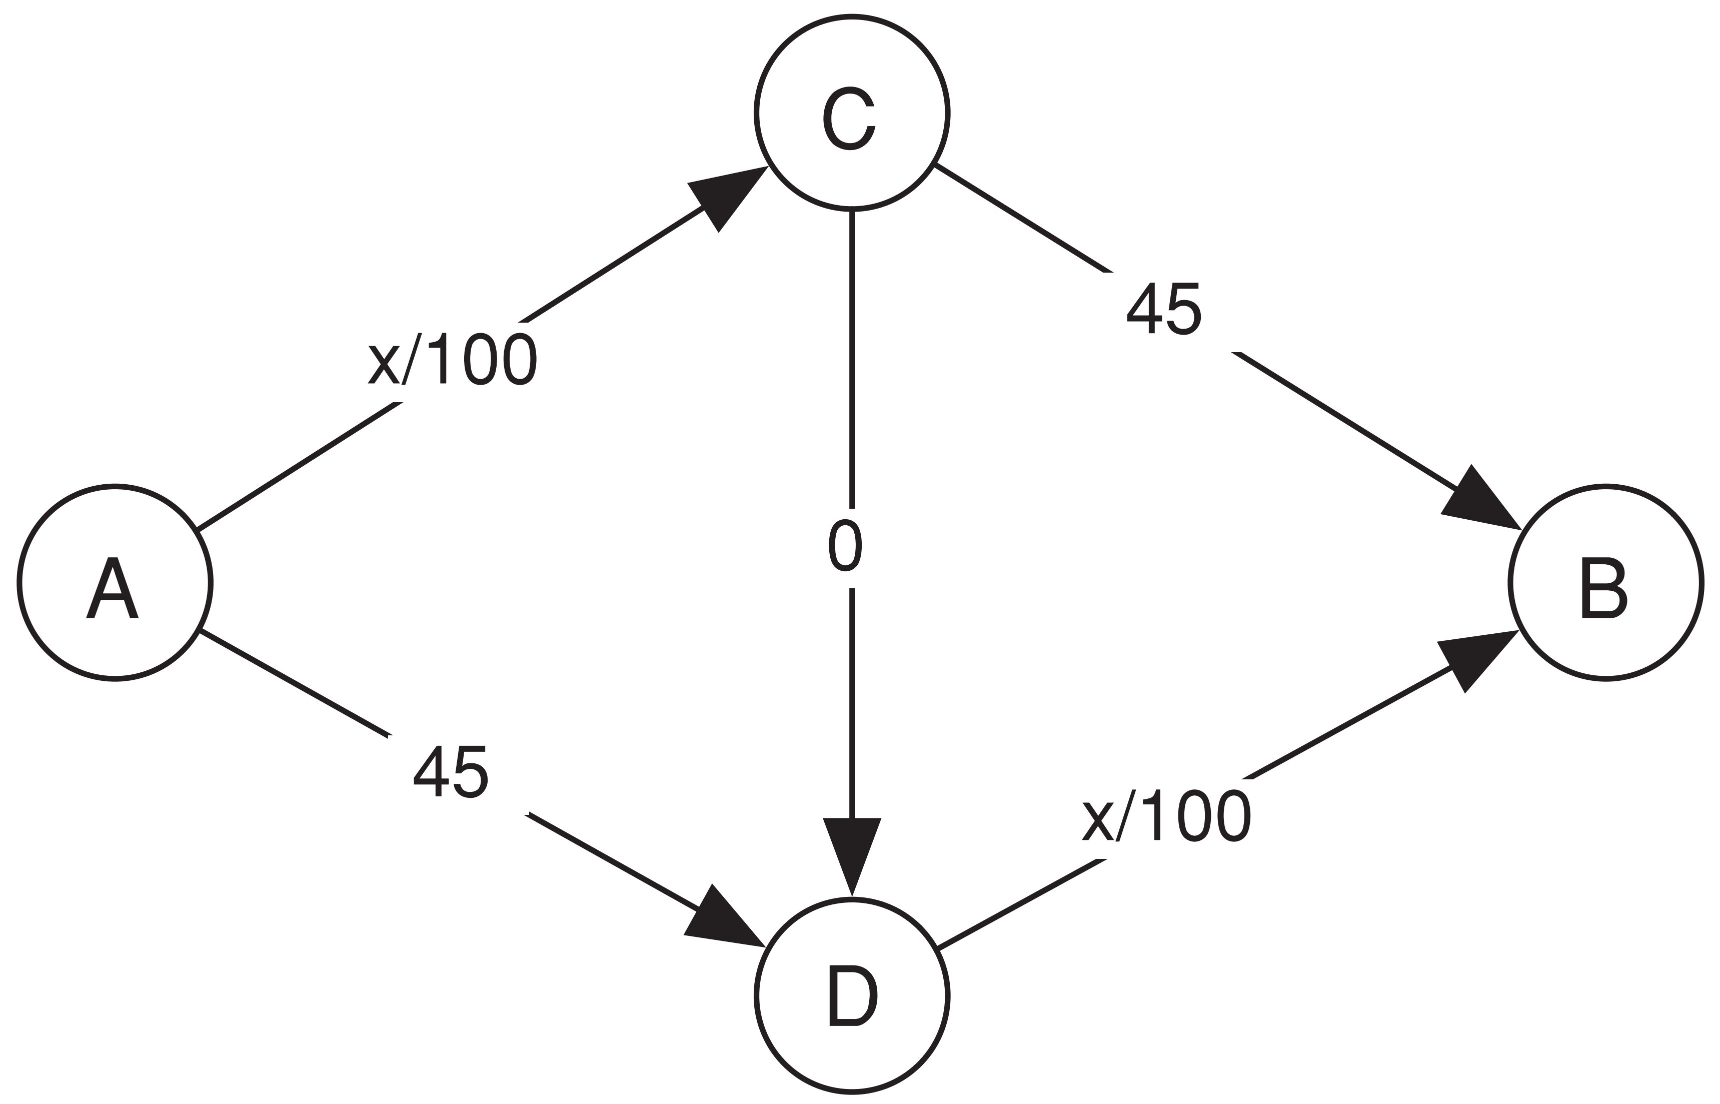

Braess's Paradox

- At equilibrium, every driver uses the route through both C and D
- As a result, the travel time for every driver is 80
- Why this is an equilibrium?
  - No driver can benefit by changing their route: with traffic snaking through C and D the way it is, any other route would now take 85 minutes!

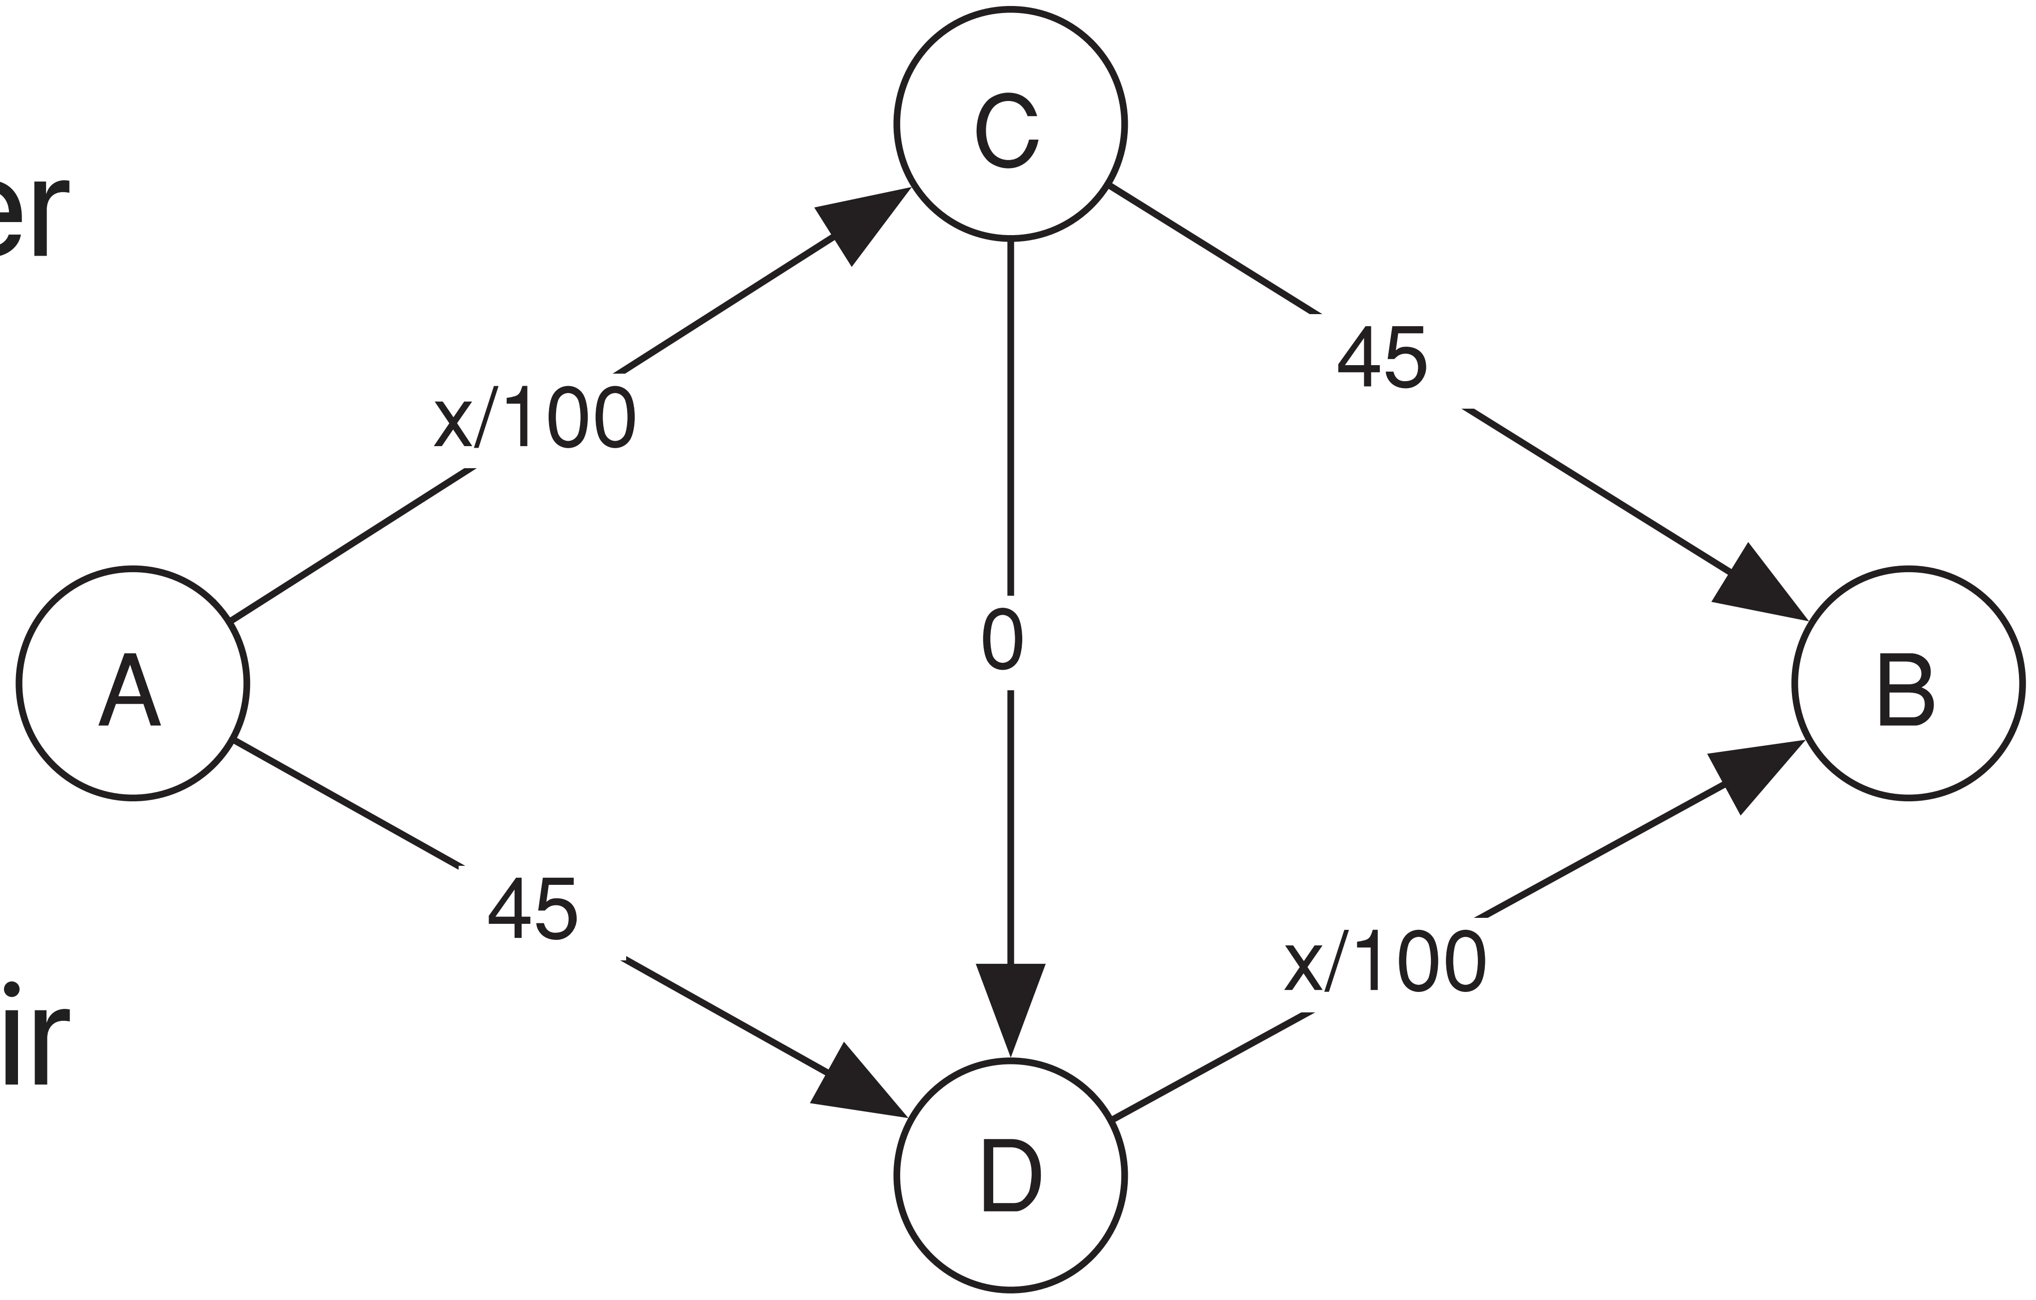

Braess's Paradox

- The creation of the edge from C to D has in fact made the route through C and D a **dominant strategy** for all drivers: regardless of the current traffic pattern, you gain by switching your route to go through C and D!
- In the new network, there is no way, given individually self-interested behaviour by the drivers, to get back to the even-balance solution that was better for everyone
- Adding resources to a transportation network can sometimes hurt performance at equilibrium! — **Braess's Paradox**
  - Just like adding the option of "confess" doesn't improve the payoff in the Prisoner's Dilemma
  - Network traffic **at equilibrium** may not be **socially optimal**!

The social cost of traffic at equilibrium

- In any general network, how **far** from optimal traffic can be at equilibrium?
  - The network can be any directed graph
  - There is a set of drivers, and different drivers may have different starting points and destinations
  - each edge e has a travel-time function: $T_e(x) = a_e x + b_e$
  - A traffic pattern is simply a choice of a path by each driver
  - The **social cost** of a given traffic pattern is the sum of the travel times incurred by all drivers when they use this traffic pattern

Social welfare maximizers

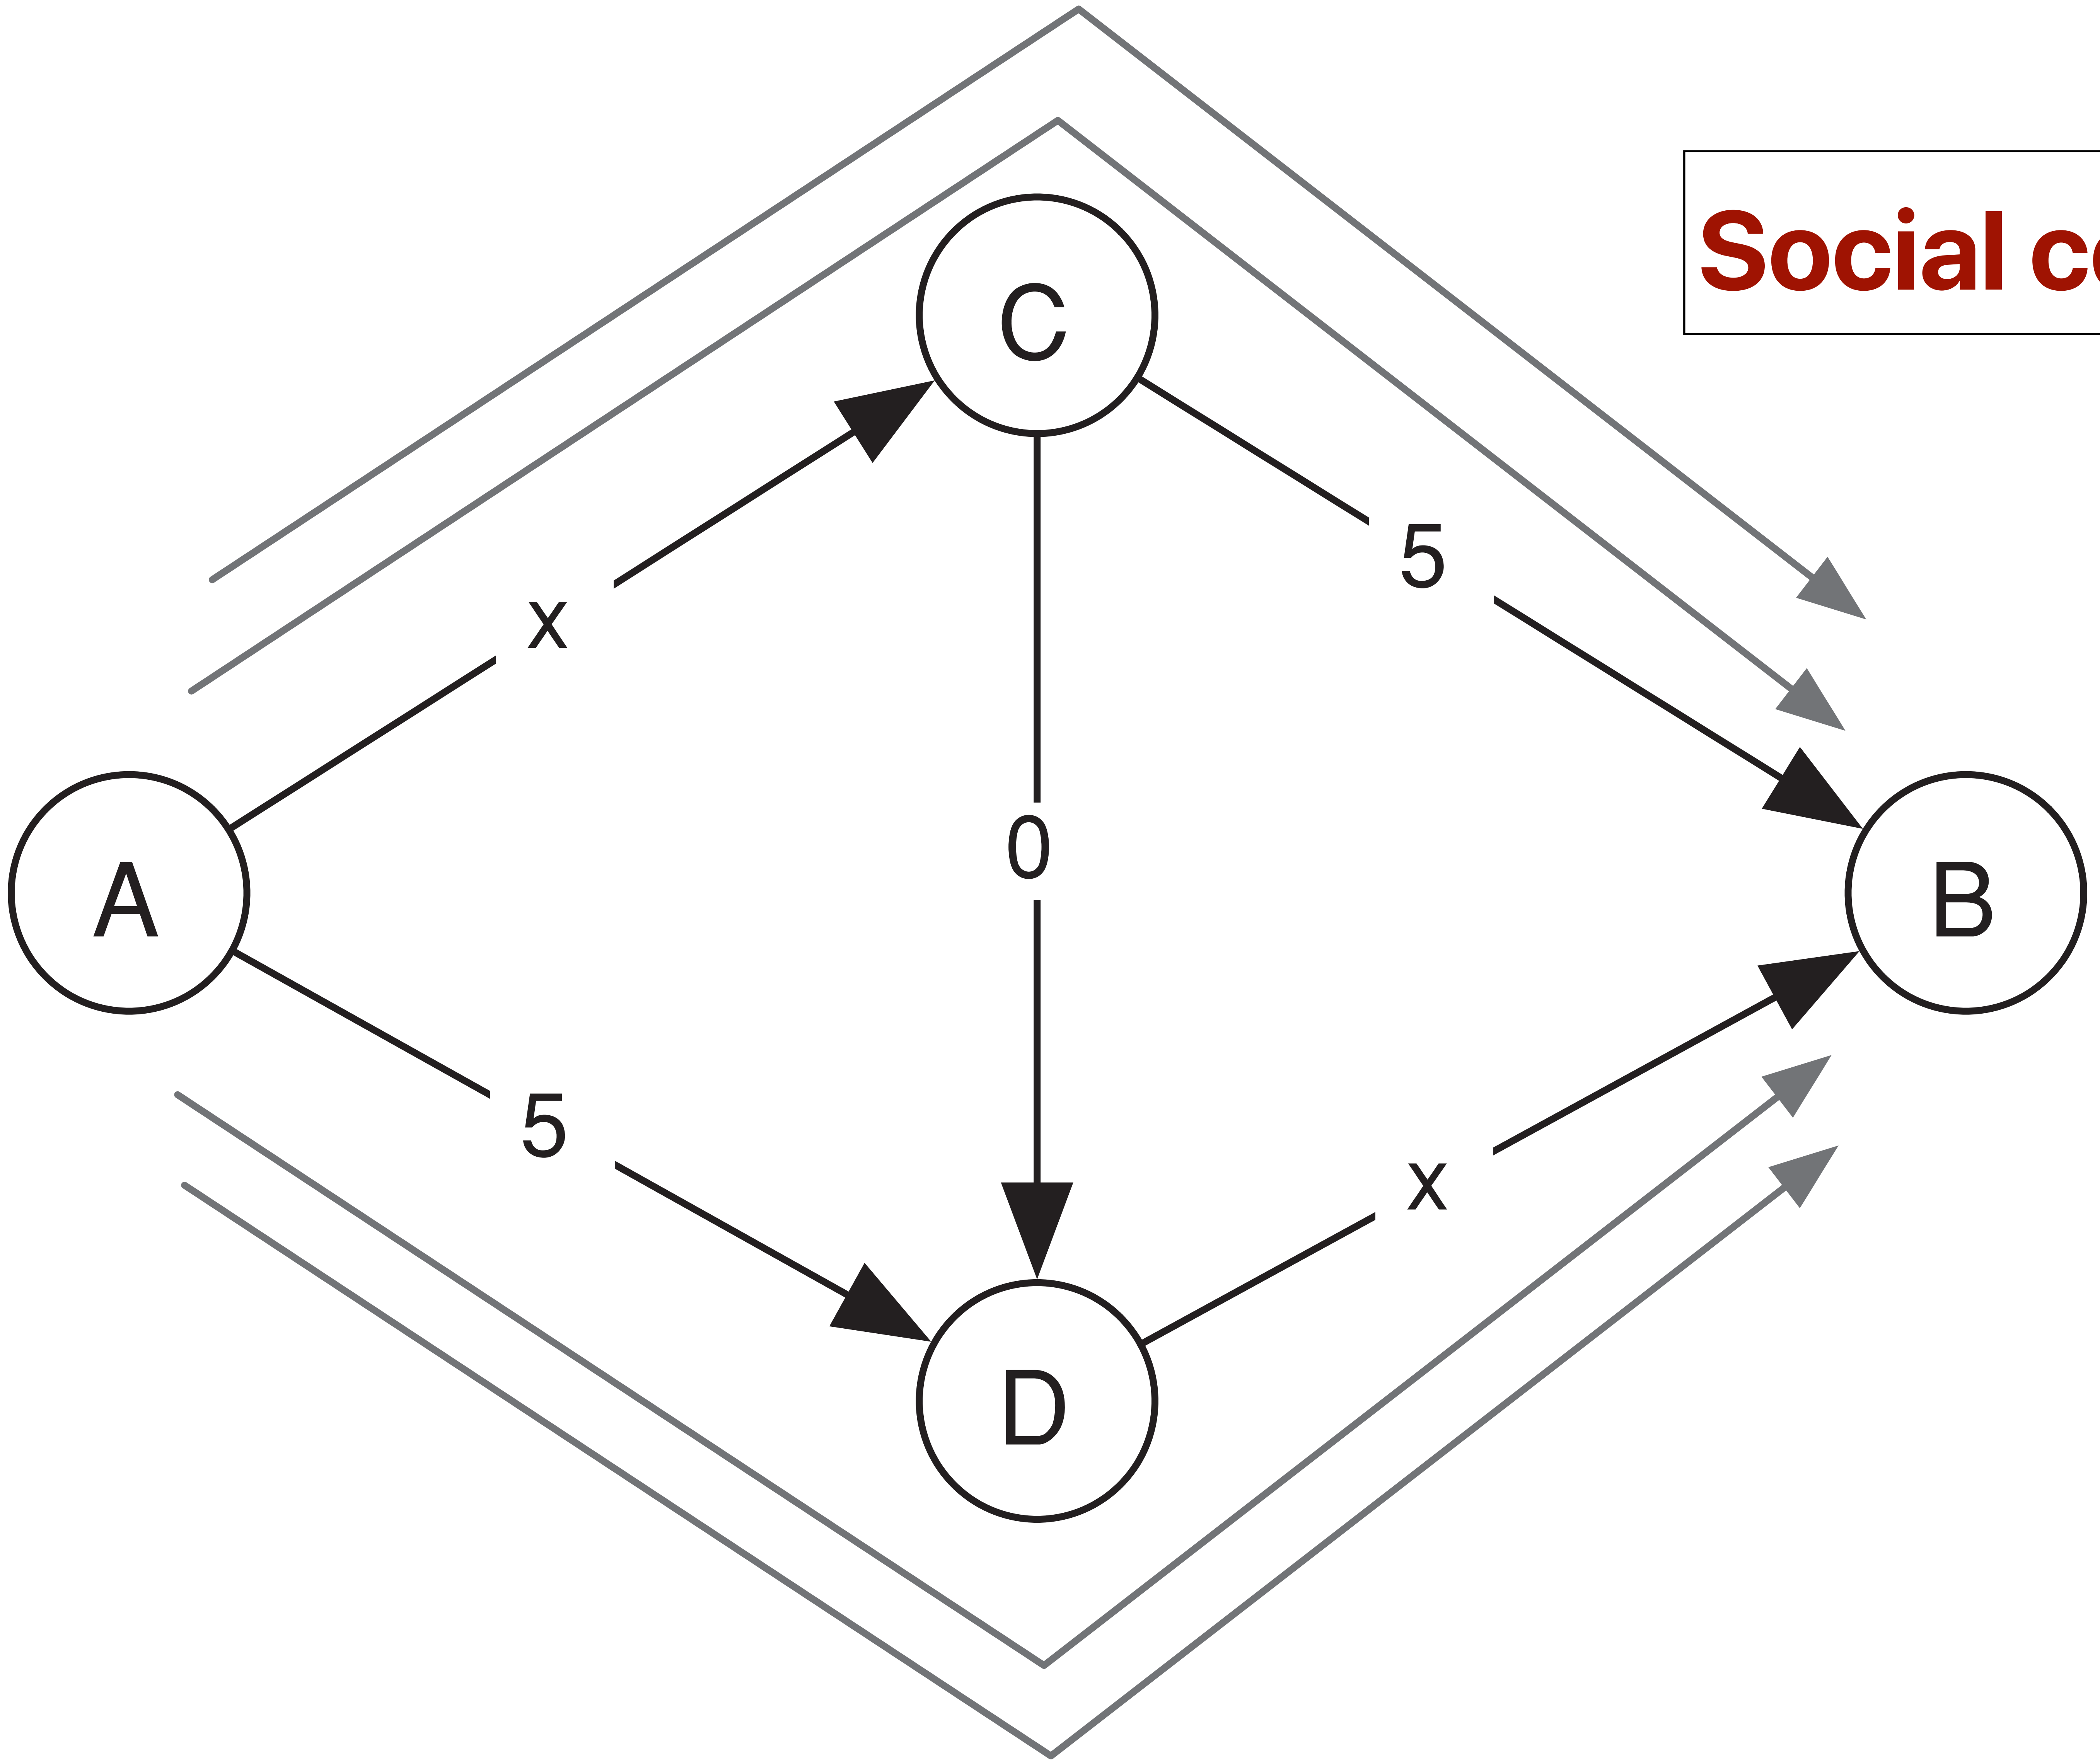

Social cost = 7 × 4 = 28

Unique Nash equilibrium

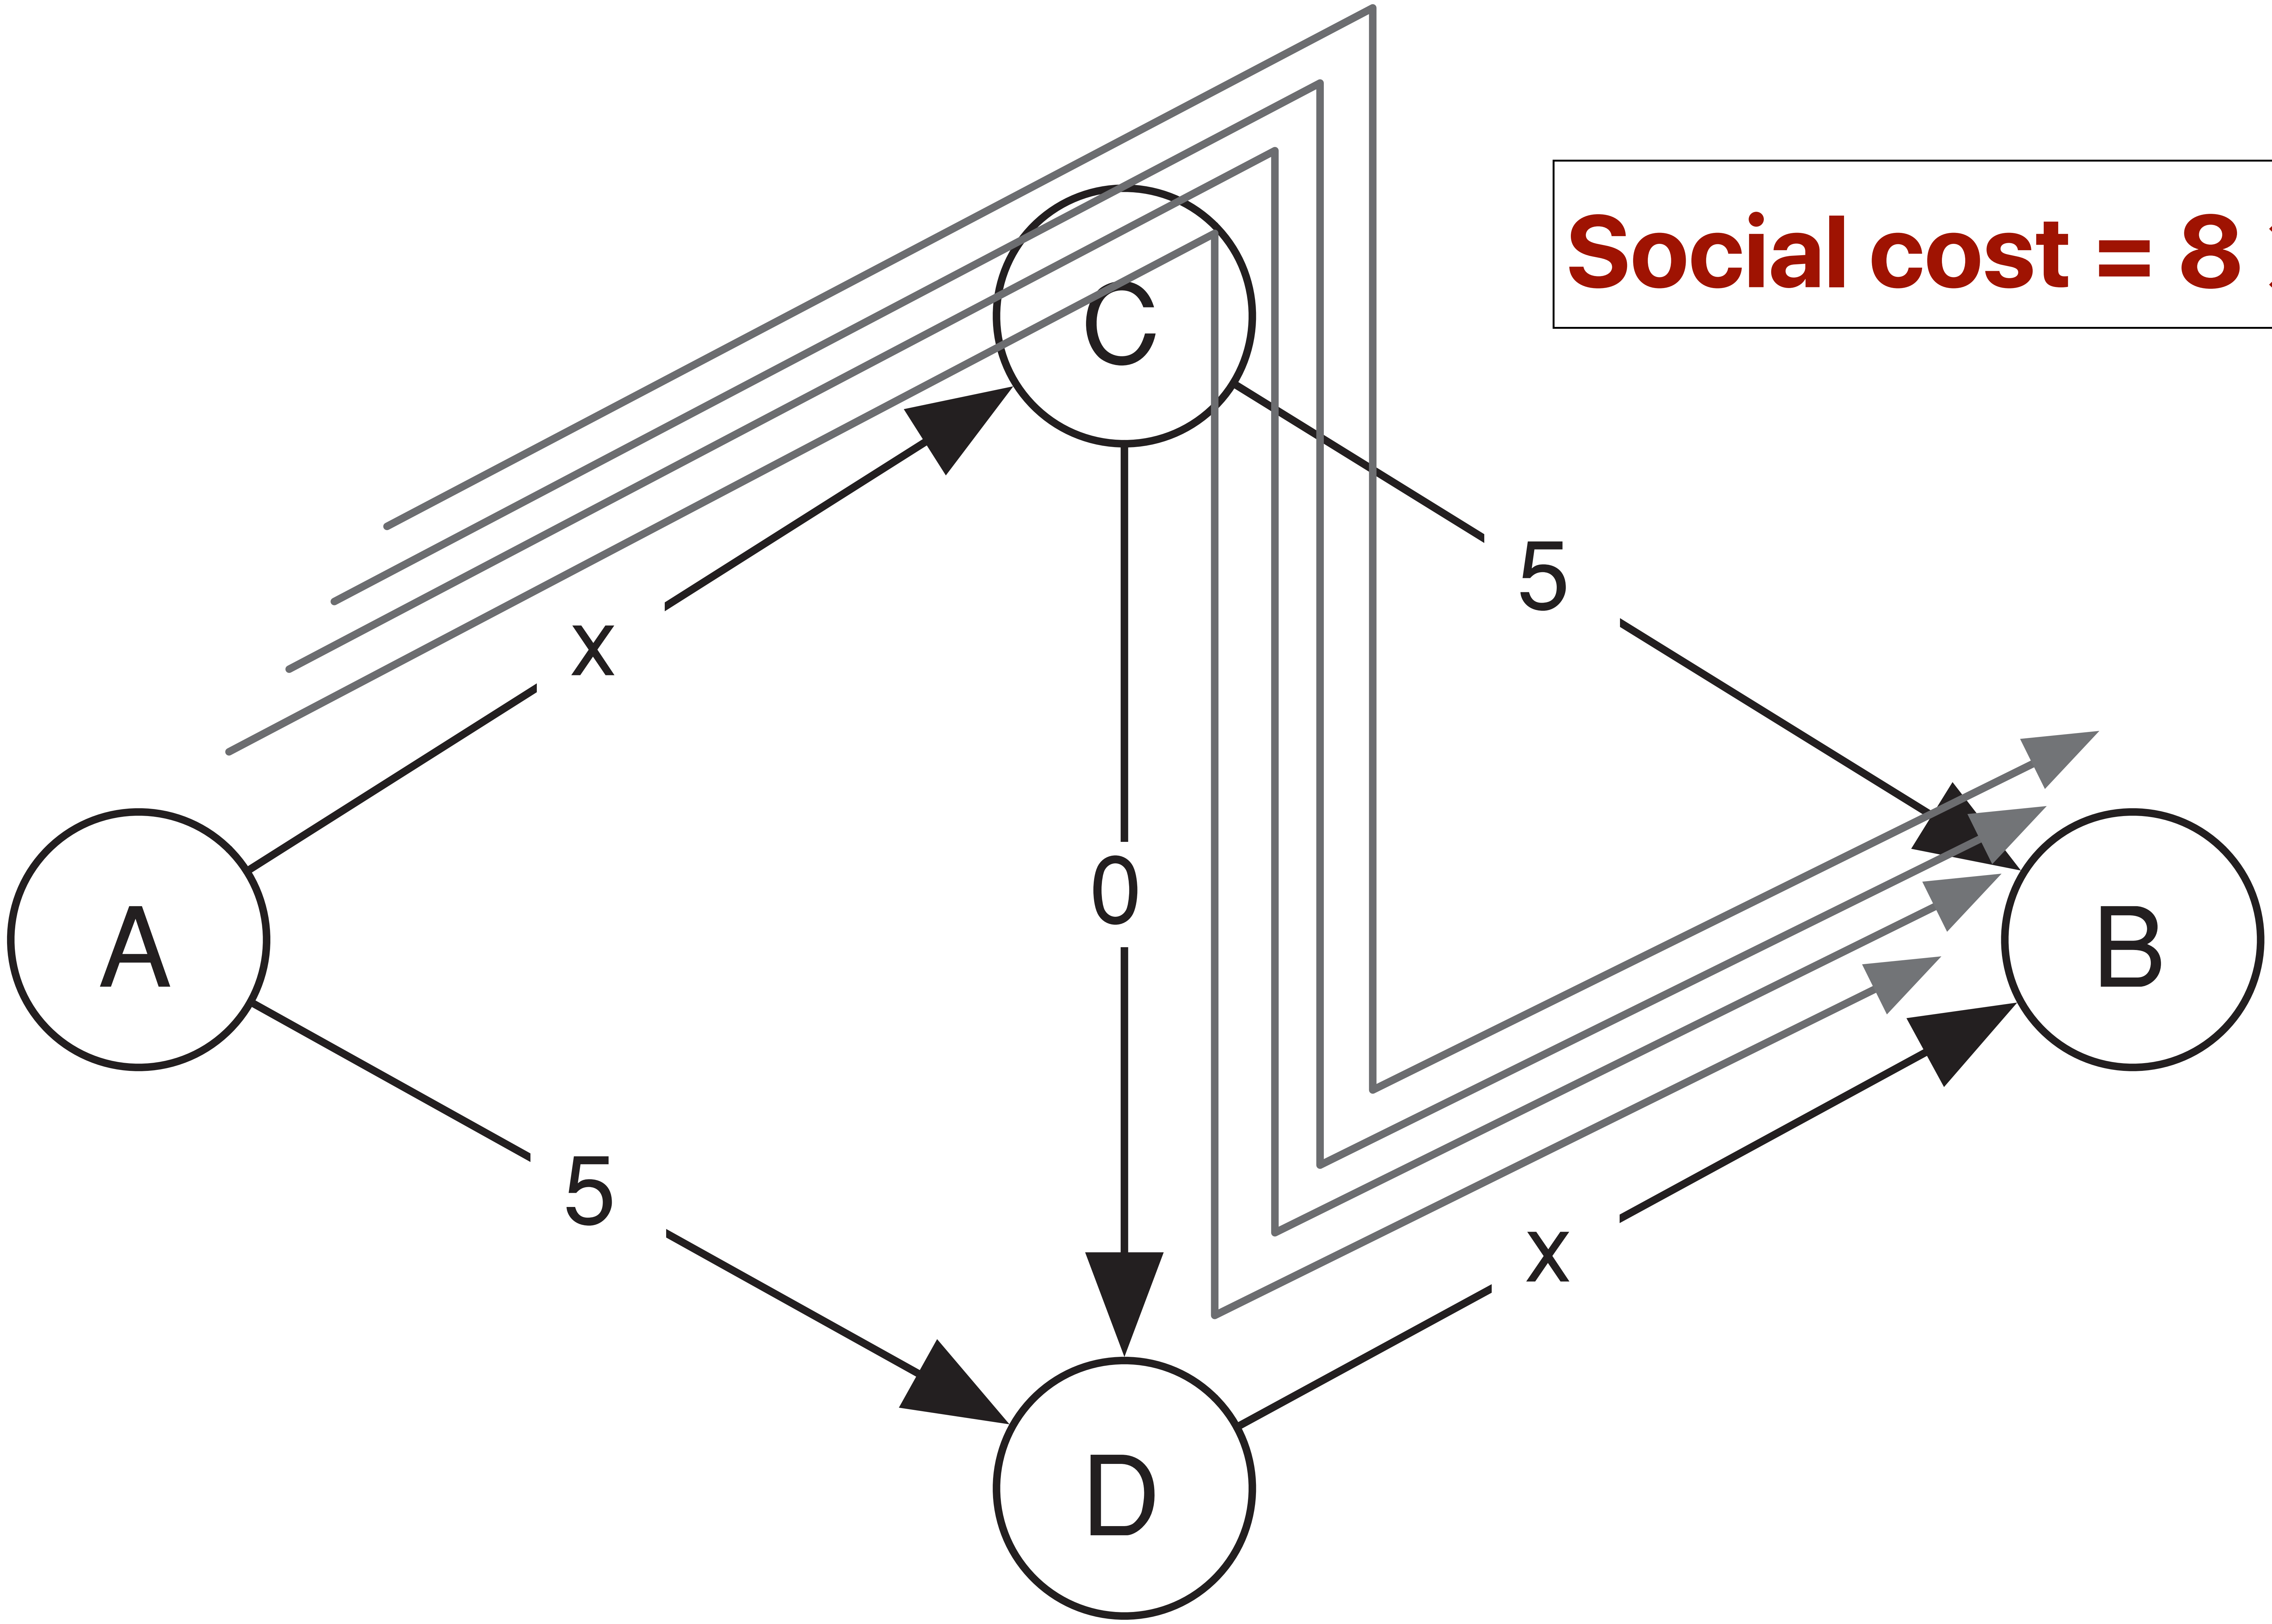

Social cost = 8 × 4 = 32

Two questions

- Q1: Is there always an equilibrium traffic pattern?
  - Pure strategy equilibria may not exist
  - But the answer is **yes** in the traffic game

Two questions

- Q2: Is there always an equilibrium traffic pattern whose social cost is not much more than the social optimum?
  - **Roughgarden and Tardos:** There is always an equilibrium whose social cost is at most 4/3 that of the optimum
  - If we change 45 to 40 in our example, we get the upper bound of "penalty": 4/3 (an increase in travel times from 60 to 80)
  - It is as bad as it can get!

Roughgarden and Tardos, "How Bad is Selfish Routing?" Journal of the ACM, 2002.

How to find a traffic pattern at equilibrium?

- Using an algorithm called **best-response dynamics**
  - Start from any traffic pattern
  - If it is an equilibrium, done!
  - Otherwise, there is at least one driver whose **best response** is some alternate path (with a strictly lower travel time)
  - Check again to see if it is an equilibrium, and continue iterating until the algorithm stops at an equilibrium
- But why should the algorithm stop?
  - In the Matching Pennies game, if only pure strategies are allowed, two players with switch between H and T forever

Analyzing best-response dynamics

- Idea: define a **measure** that shows **progress** in the algorithm
- **Social cost** may not be a suitable progress measure
  - As it may become worse in the Braess's Paradox
  - We need some other measure to show progress
- **Potential energy** of a traffic pattern
  - If an edge **e** currently has **x** drivers on it, then the potential energy of this edge is:

$$\text{Energy}(e) = T_e(1) + T_e(2) + \cdots + T_e(x)$$

  - Interpreted as the "cumulative" quantity in which we imagine increasing batches of drivers crossing the edge
  - The potential energy of a traffic pattern is the sum of the potential energies of all the edges

Potential energy in best-response dynamics

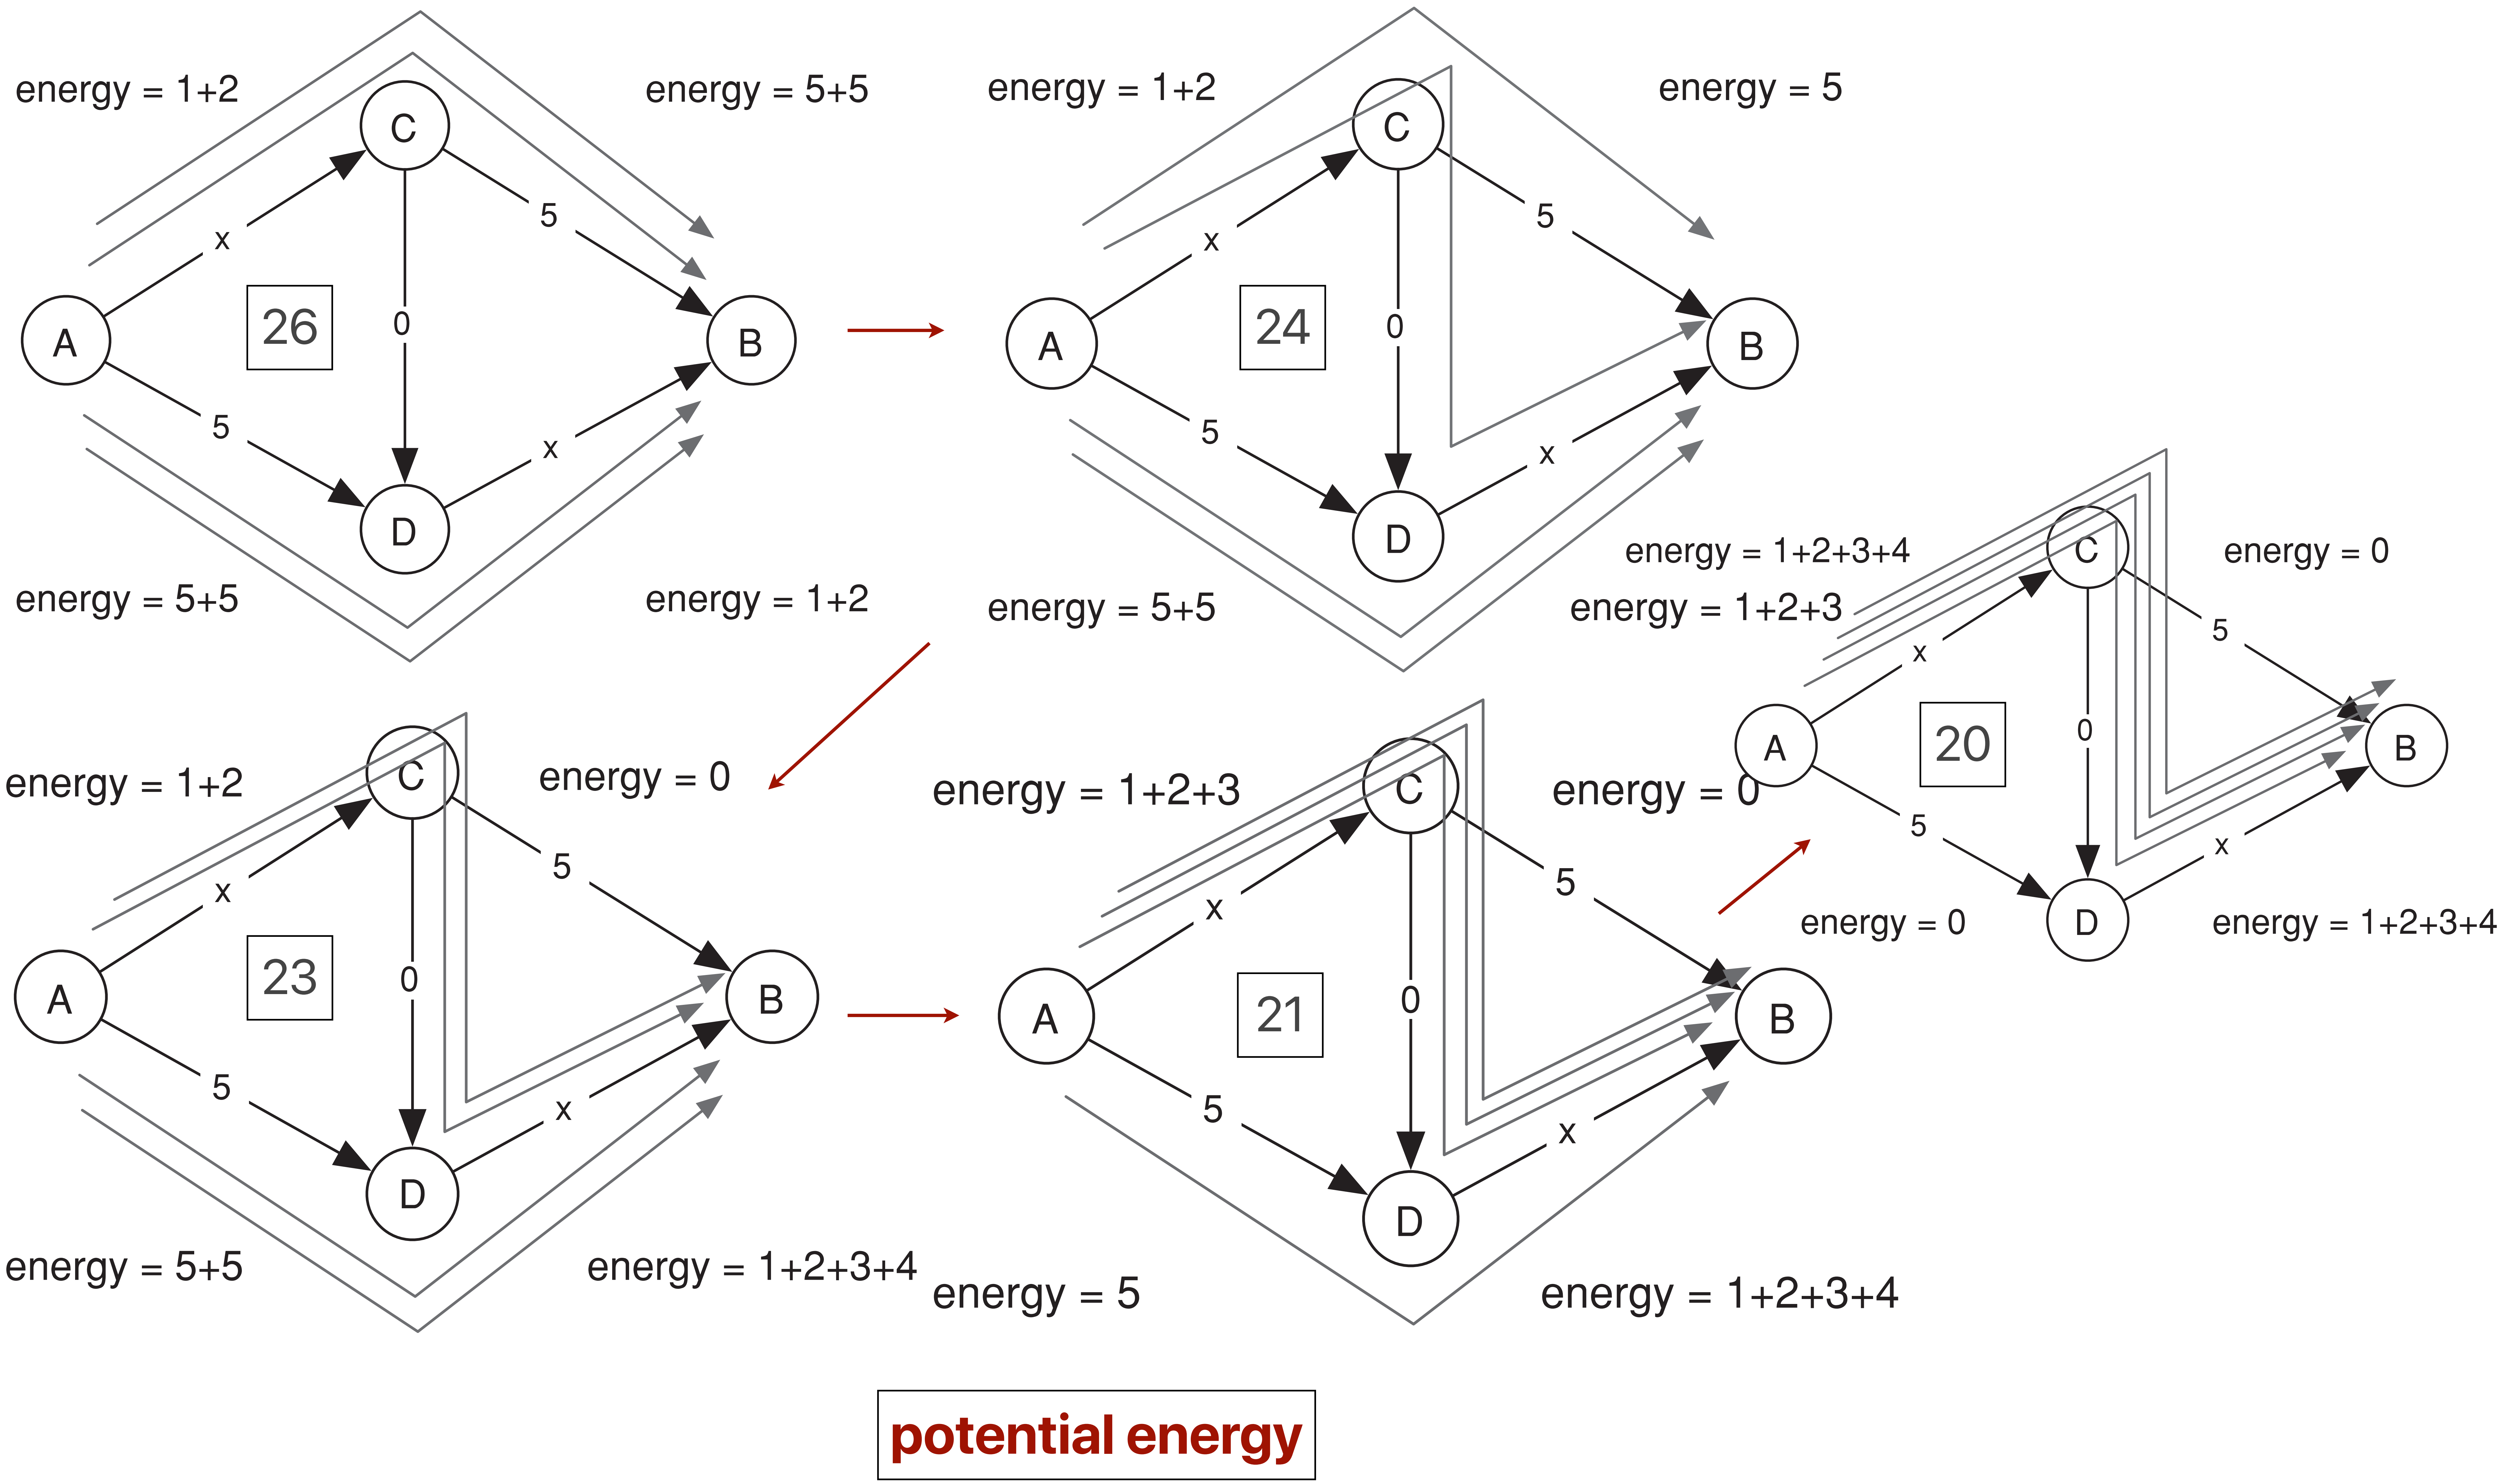

The potential energy strictly decreases

- From one traffic pattern to the next, the only change is that one driver abandons his current path and switches to a new one
- First releases potential energy as the driver leaves the system
  - The change in potential energy on edge e is $T_e(x)$, exactly the travel time that the driver was experiencing on e
- Then adds potential energy as he rejoins
  - The increase of $T_e(x+1)$ is exactly the new travel time the driver experiences on this edge
- The net change in potential energy is simply his new travel time minus his old travel time
- With **best-response dynamics**, this must be **negative**!

Best-response dynamics must terminate

- The potential energy can only take a finite number of possible values, one for each possible traffic pattern
- If it is strictly decreasing with each step of best-response dynamics, it is “consuming” this finite supply of possible values
- Best-response dynamics must come to a stop by the time the potential energy reaches its minimum possible value (if not sooner)
- Once it terminates, we must be at an equilibrium! ∎

Comparing equilibrium traffic to social optimum

$$\text{Energy}(e) = T_e(1) + T_e(2) + \cdots + T_e(x)$$

$$\text{Total-Travel-Time}(e) = xT_e(x)$$

$$\begin{aligned}
T_e(1) + T_e(2) + \cdots + T_e(x) &= a_e(1 + 2 + \cdots + x) + b_e x \\
&= \frac{a_e x(x+1)}{2} + b_e x \\
&= x\left[\frac{a_e(x+1)}{2} + b_e\right] \\
&\geq \frac{1}{2}x(a_e x + b_e) \\
&= \frac{1}{2}xT_e(x).
\end{aligned}$$

Comparing equilibrium traffic to social optimum

$$\frac{1}{2}[\text{Social-Cost}(Z)] \leq \text{Energy}(Z) \leq \text{Social-Cost}(Z)$$

We start from a socially optimal traffic pattern **Z**, allow best-response dynamics to run till it stops at an equilibrium traffic pattern **Z'**, then we have:

$$\text{Social-Cost}(Z') \leq 2[\text{Energy}(Z')] \leq 2[\text{Energy}(Z)] \leq 2[\text{Social-Cost}(Z)]$$

Required reading: “Networks, Crowds, and Markets,” Chapter 8