Social and Information Networks

CSCC46H, Fall 2025
Lecture 9

Prof. Ashton Anderson
ashton@cs.toronto.edu

Today

A3 out this week, due last day of class

Today

Game Theory: Congestion games
Decision-Based Diffusion
Information Diffusion

Today: Game Theory in the Wild and Influence Through Networks

If people are connected through a network, it's possible for them to influence each other's knowledge, behaviour and actions

Today: why?

- Informational
- Direct benefit
- Social conformity

Getting to UTSC: 401 or Gardiner?

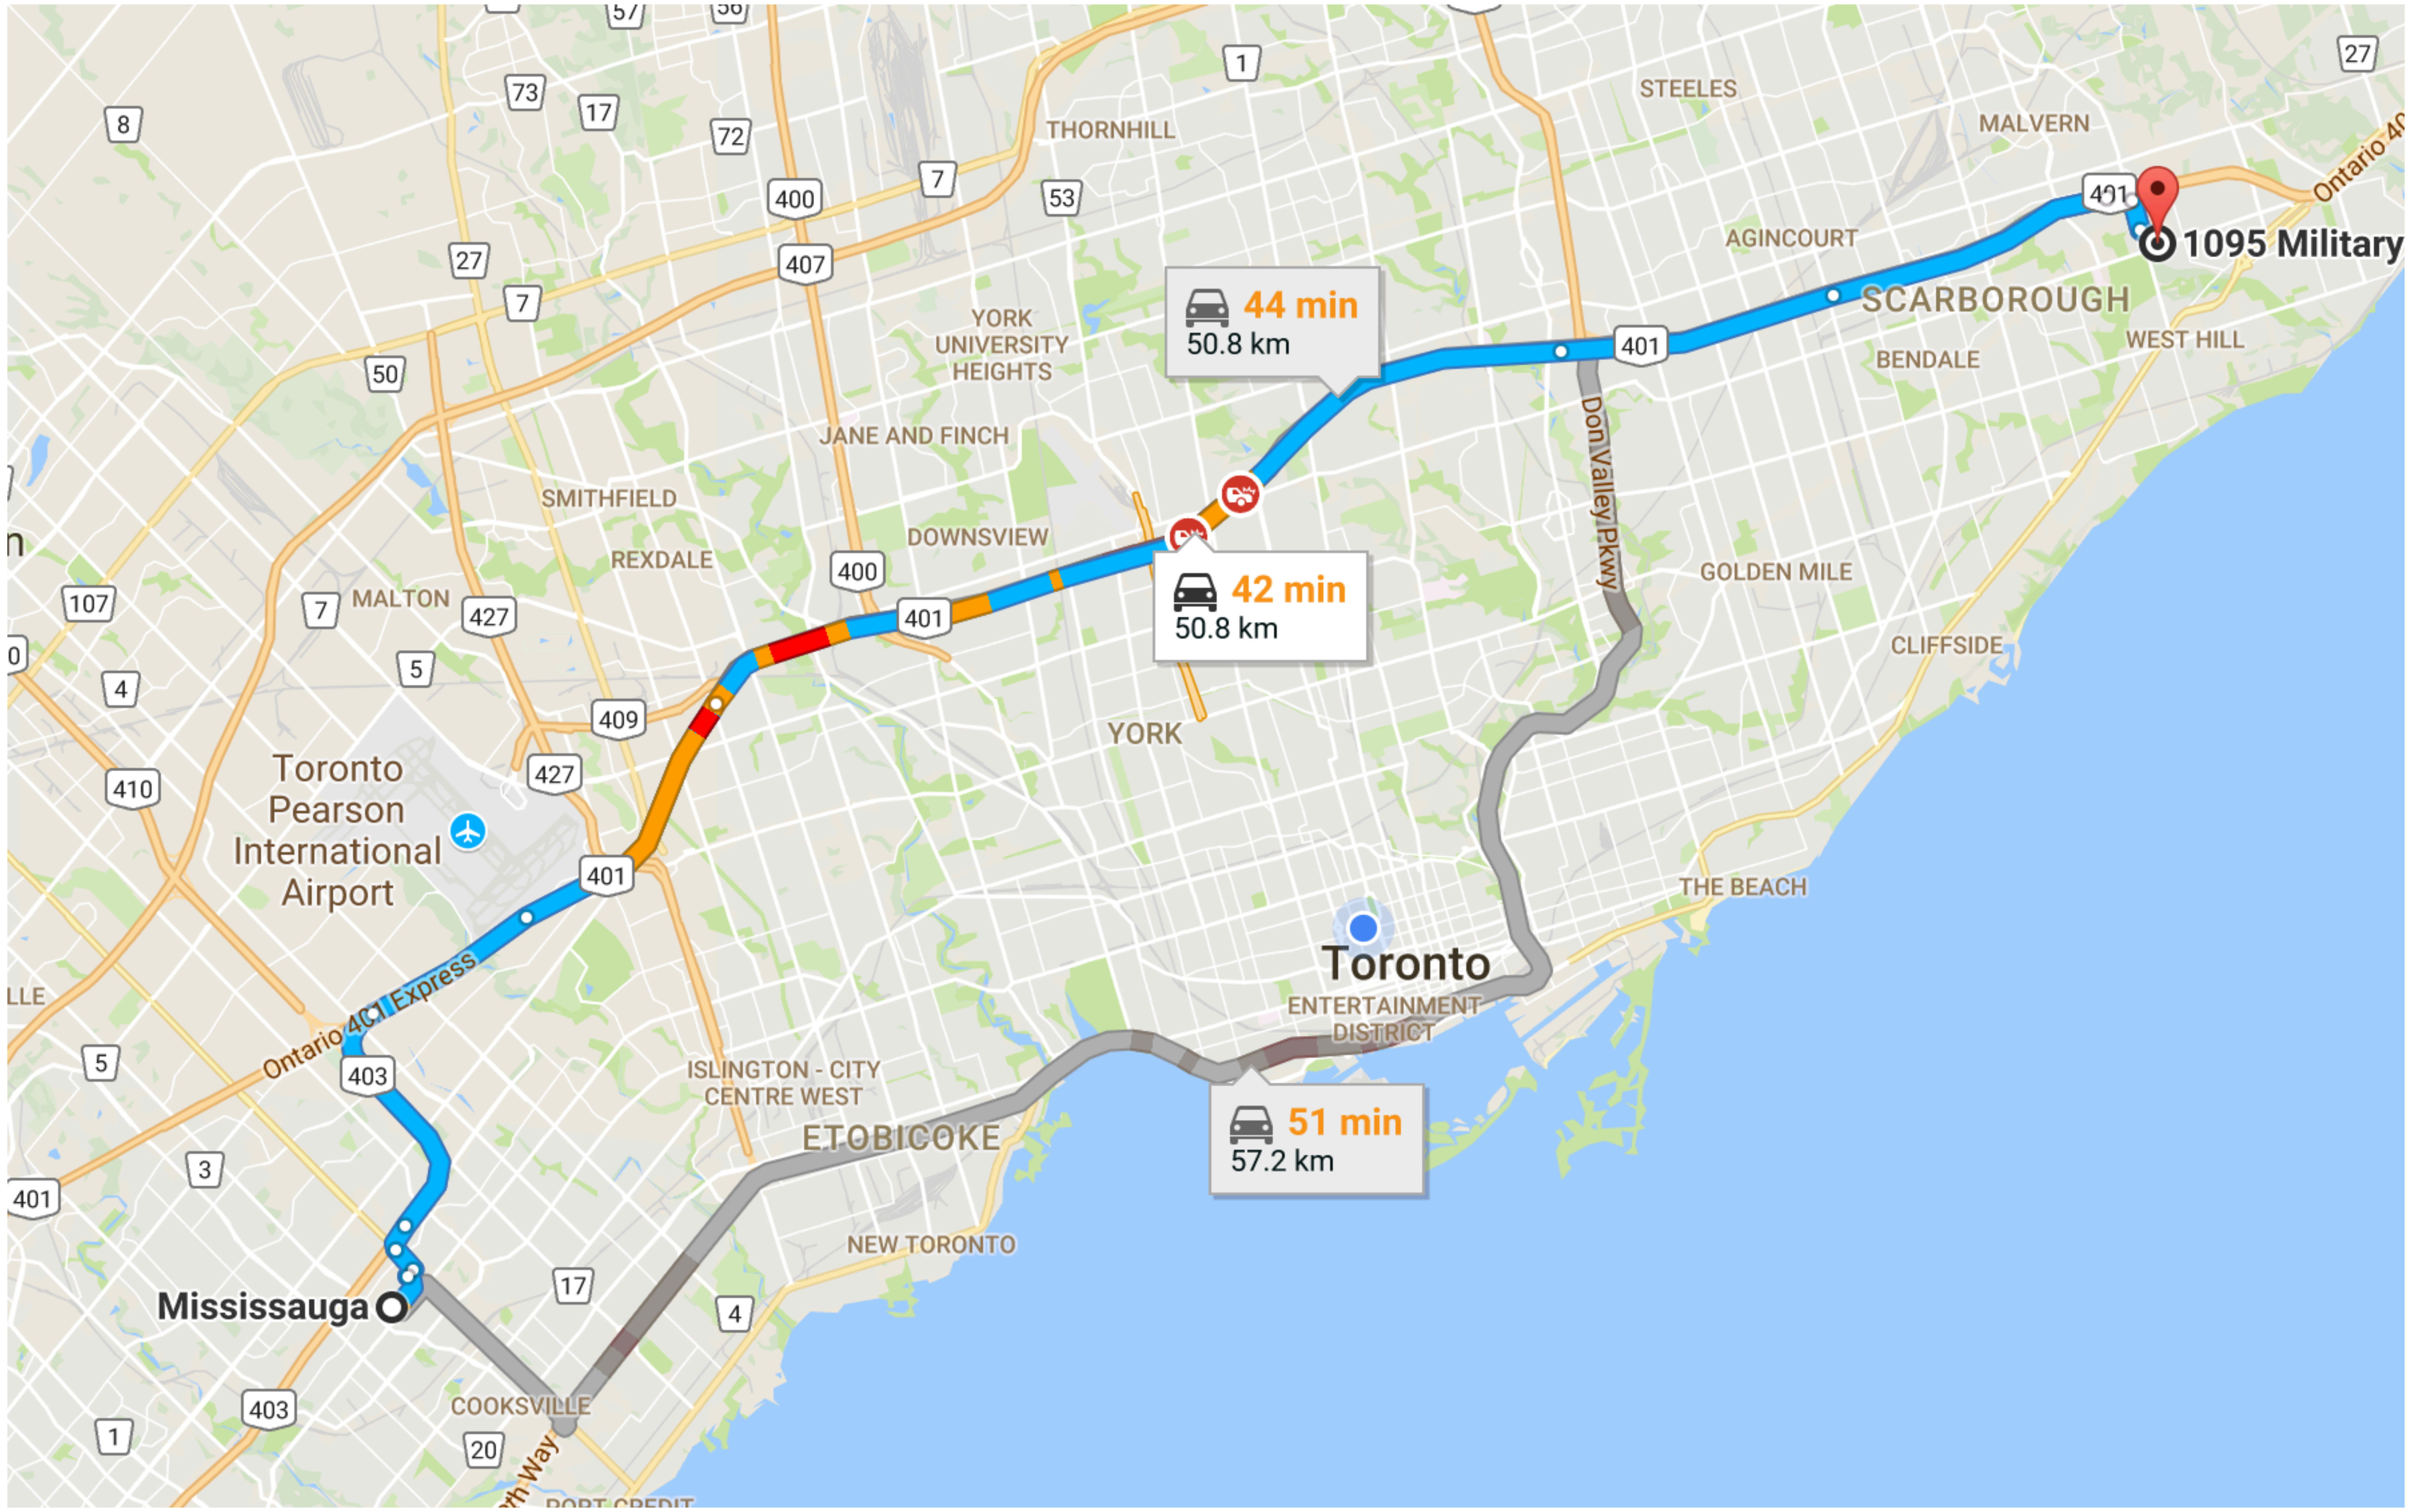

Getting to UTSC: 401 or Gardiner?

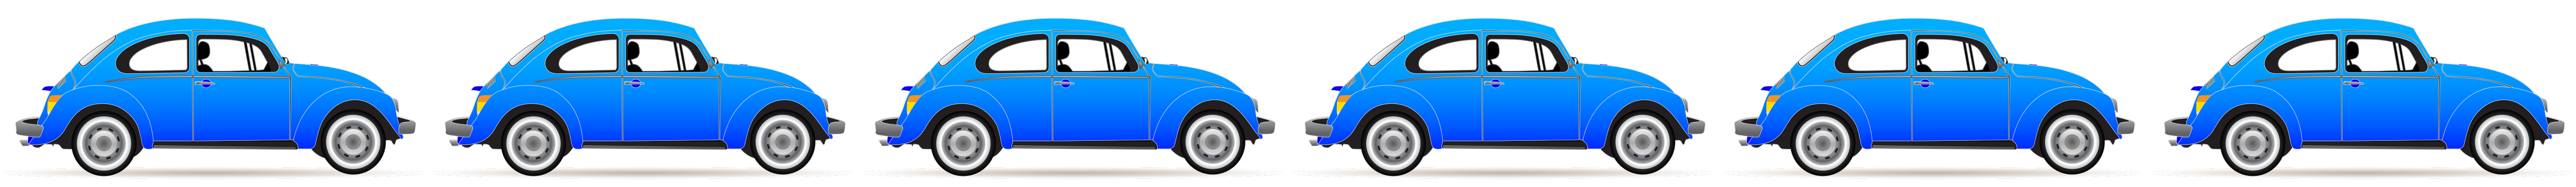

401

Gardiner

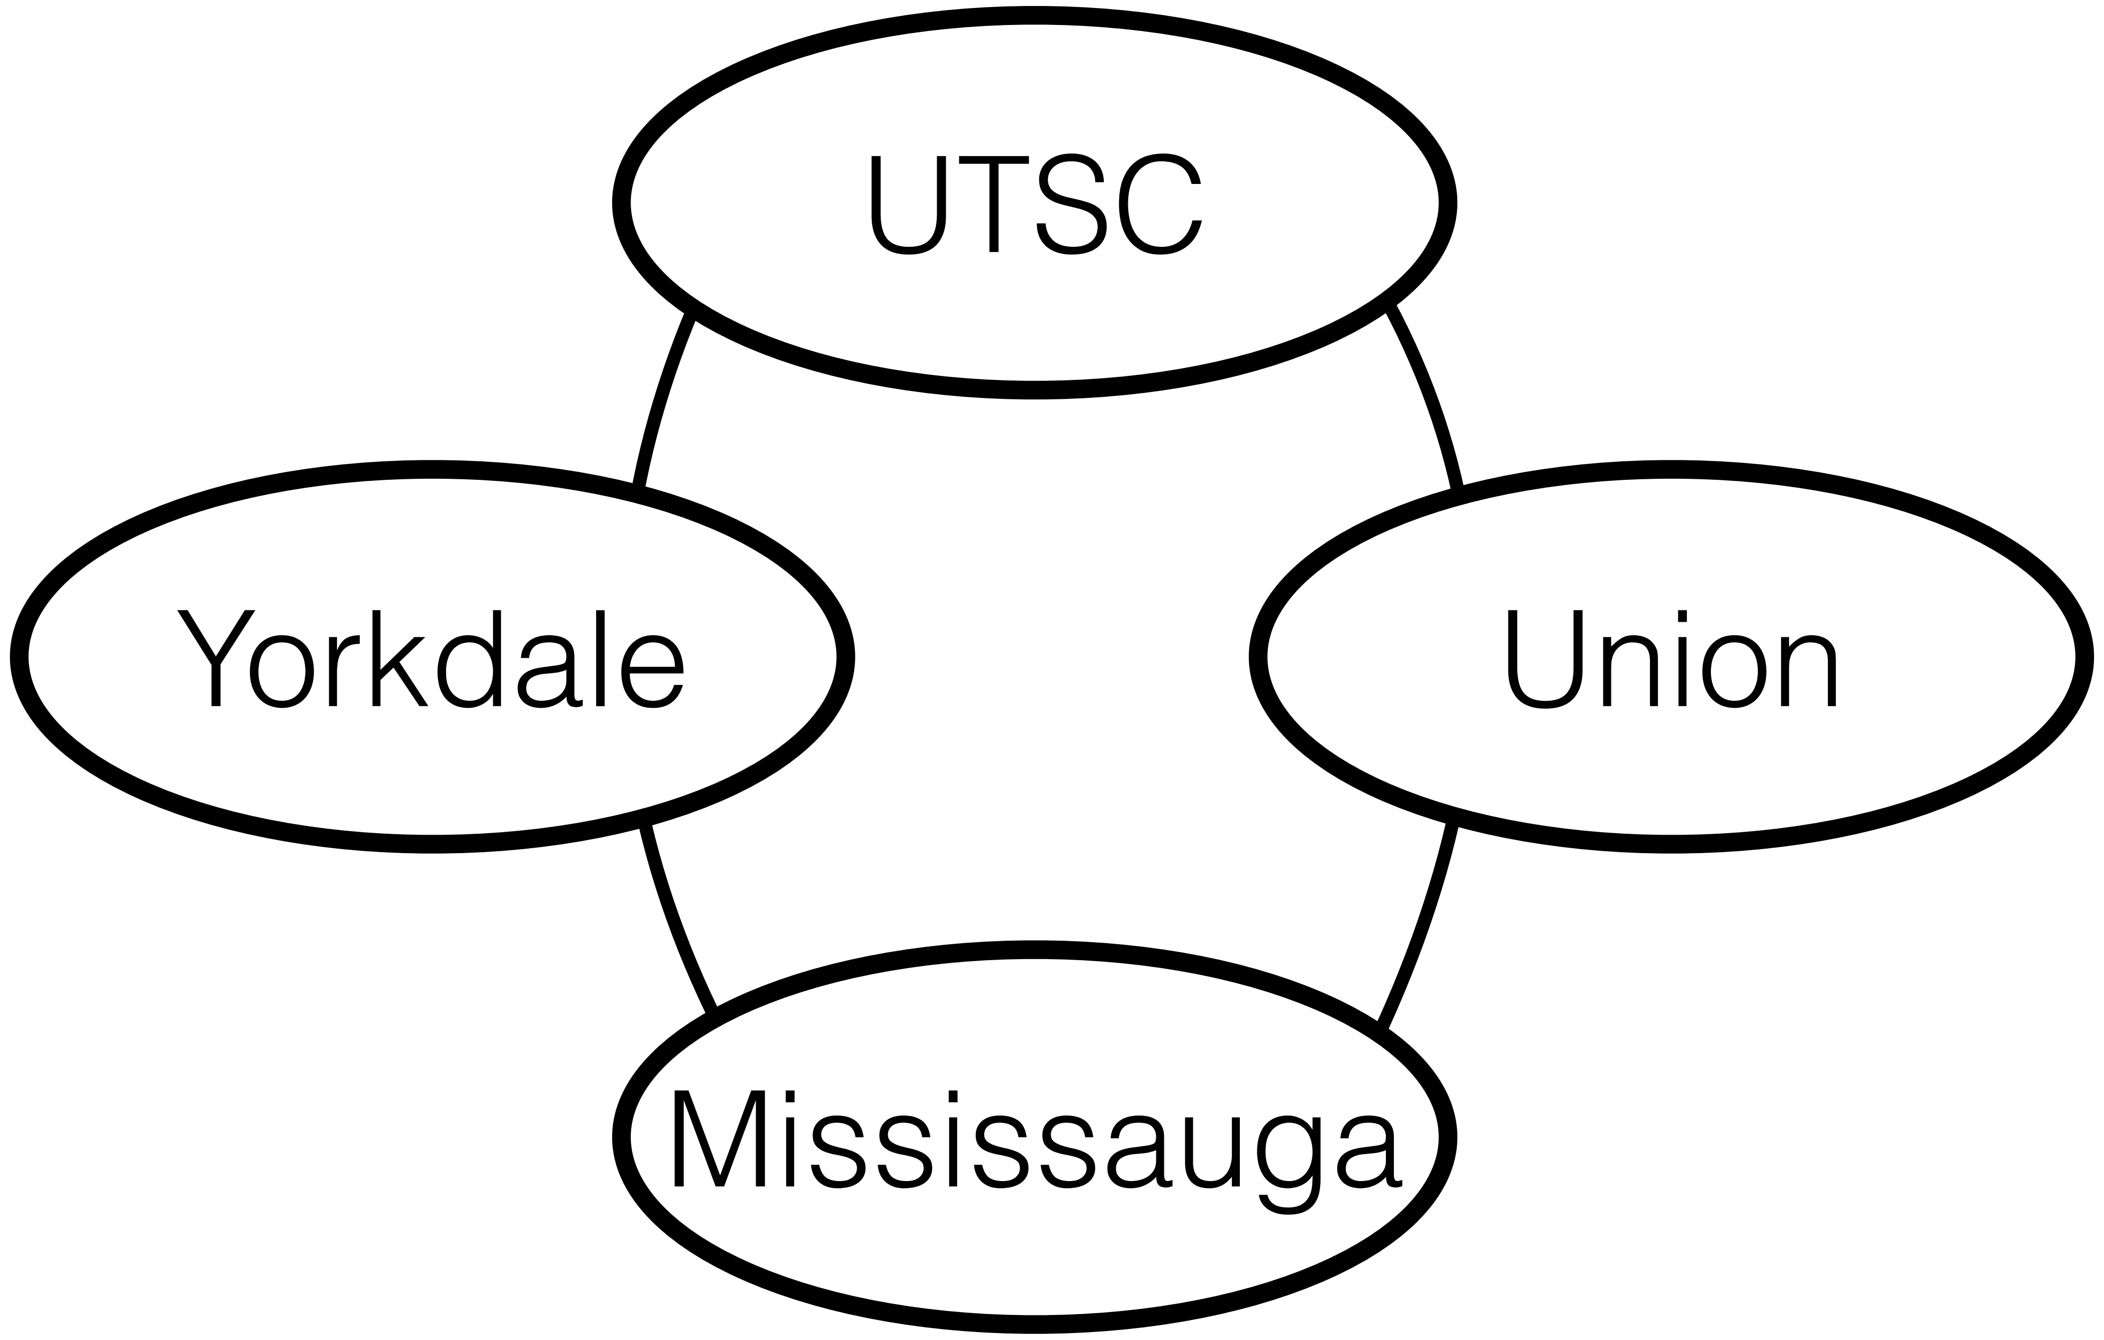

Getting to UTSC: 401 or Gardiner?

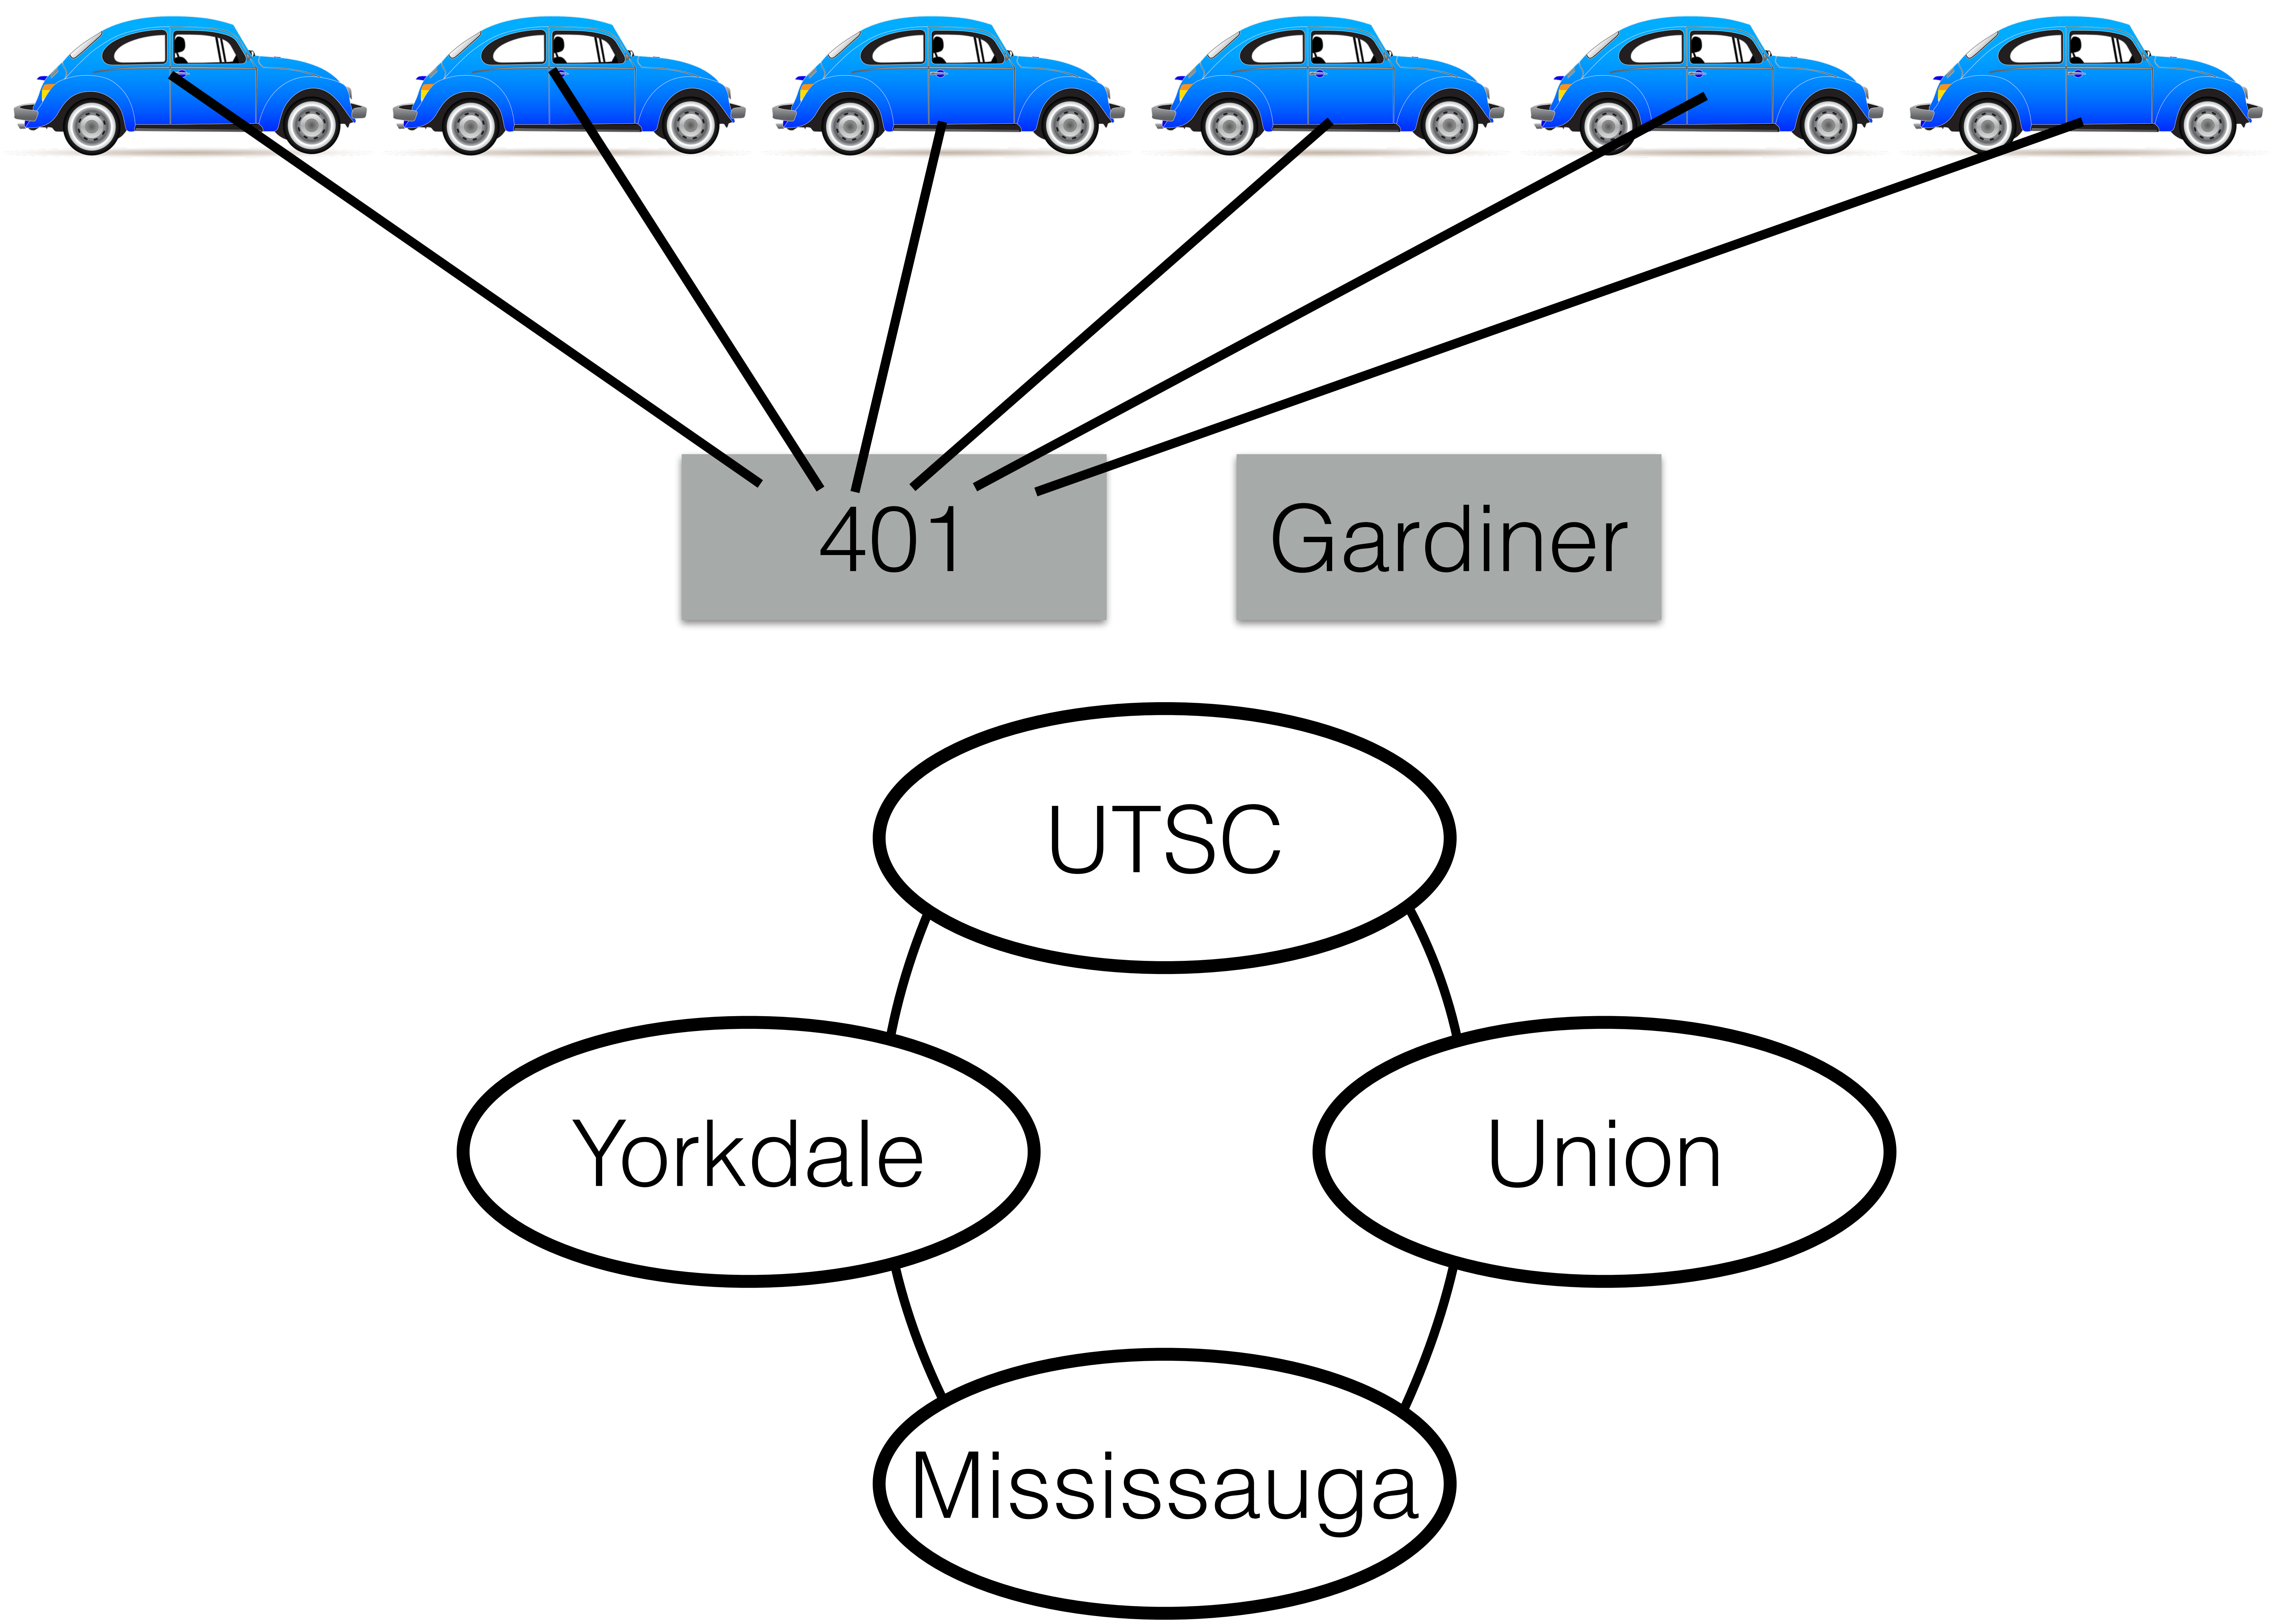

Getting to UTSC: 401 or Gardiner?

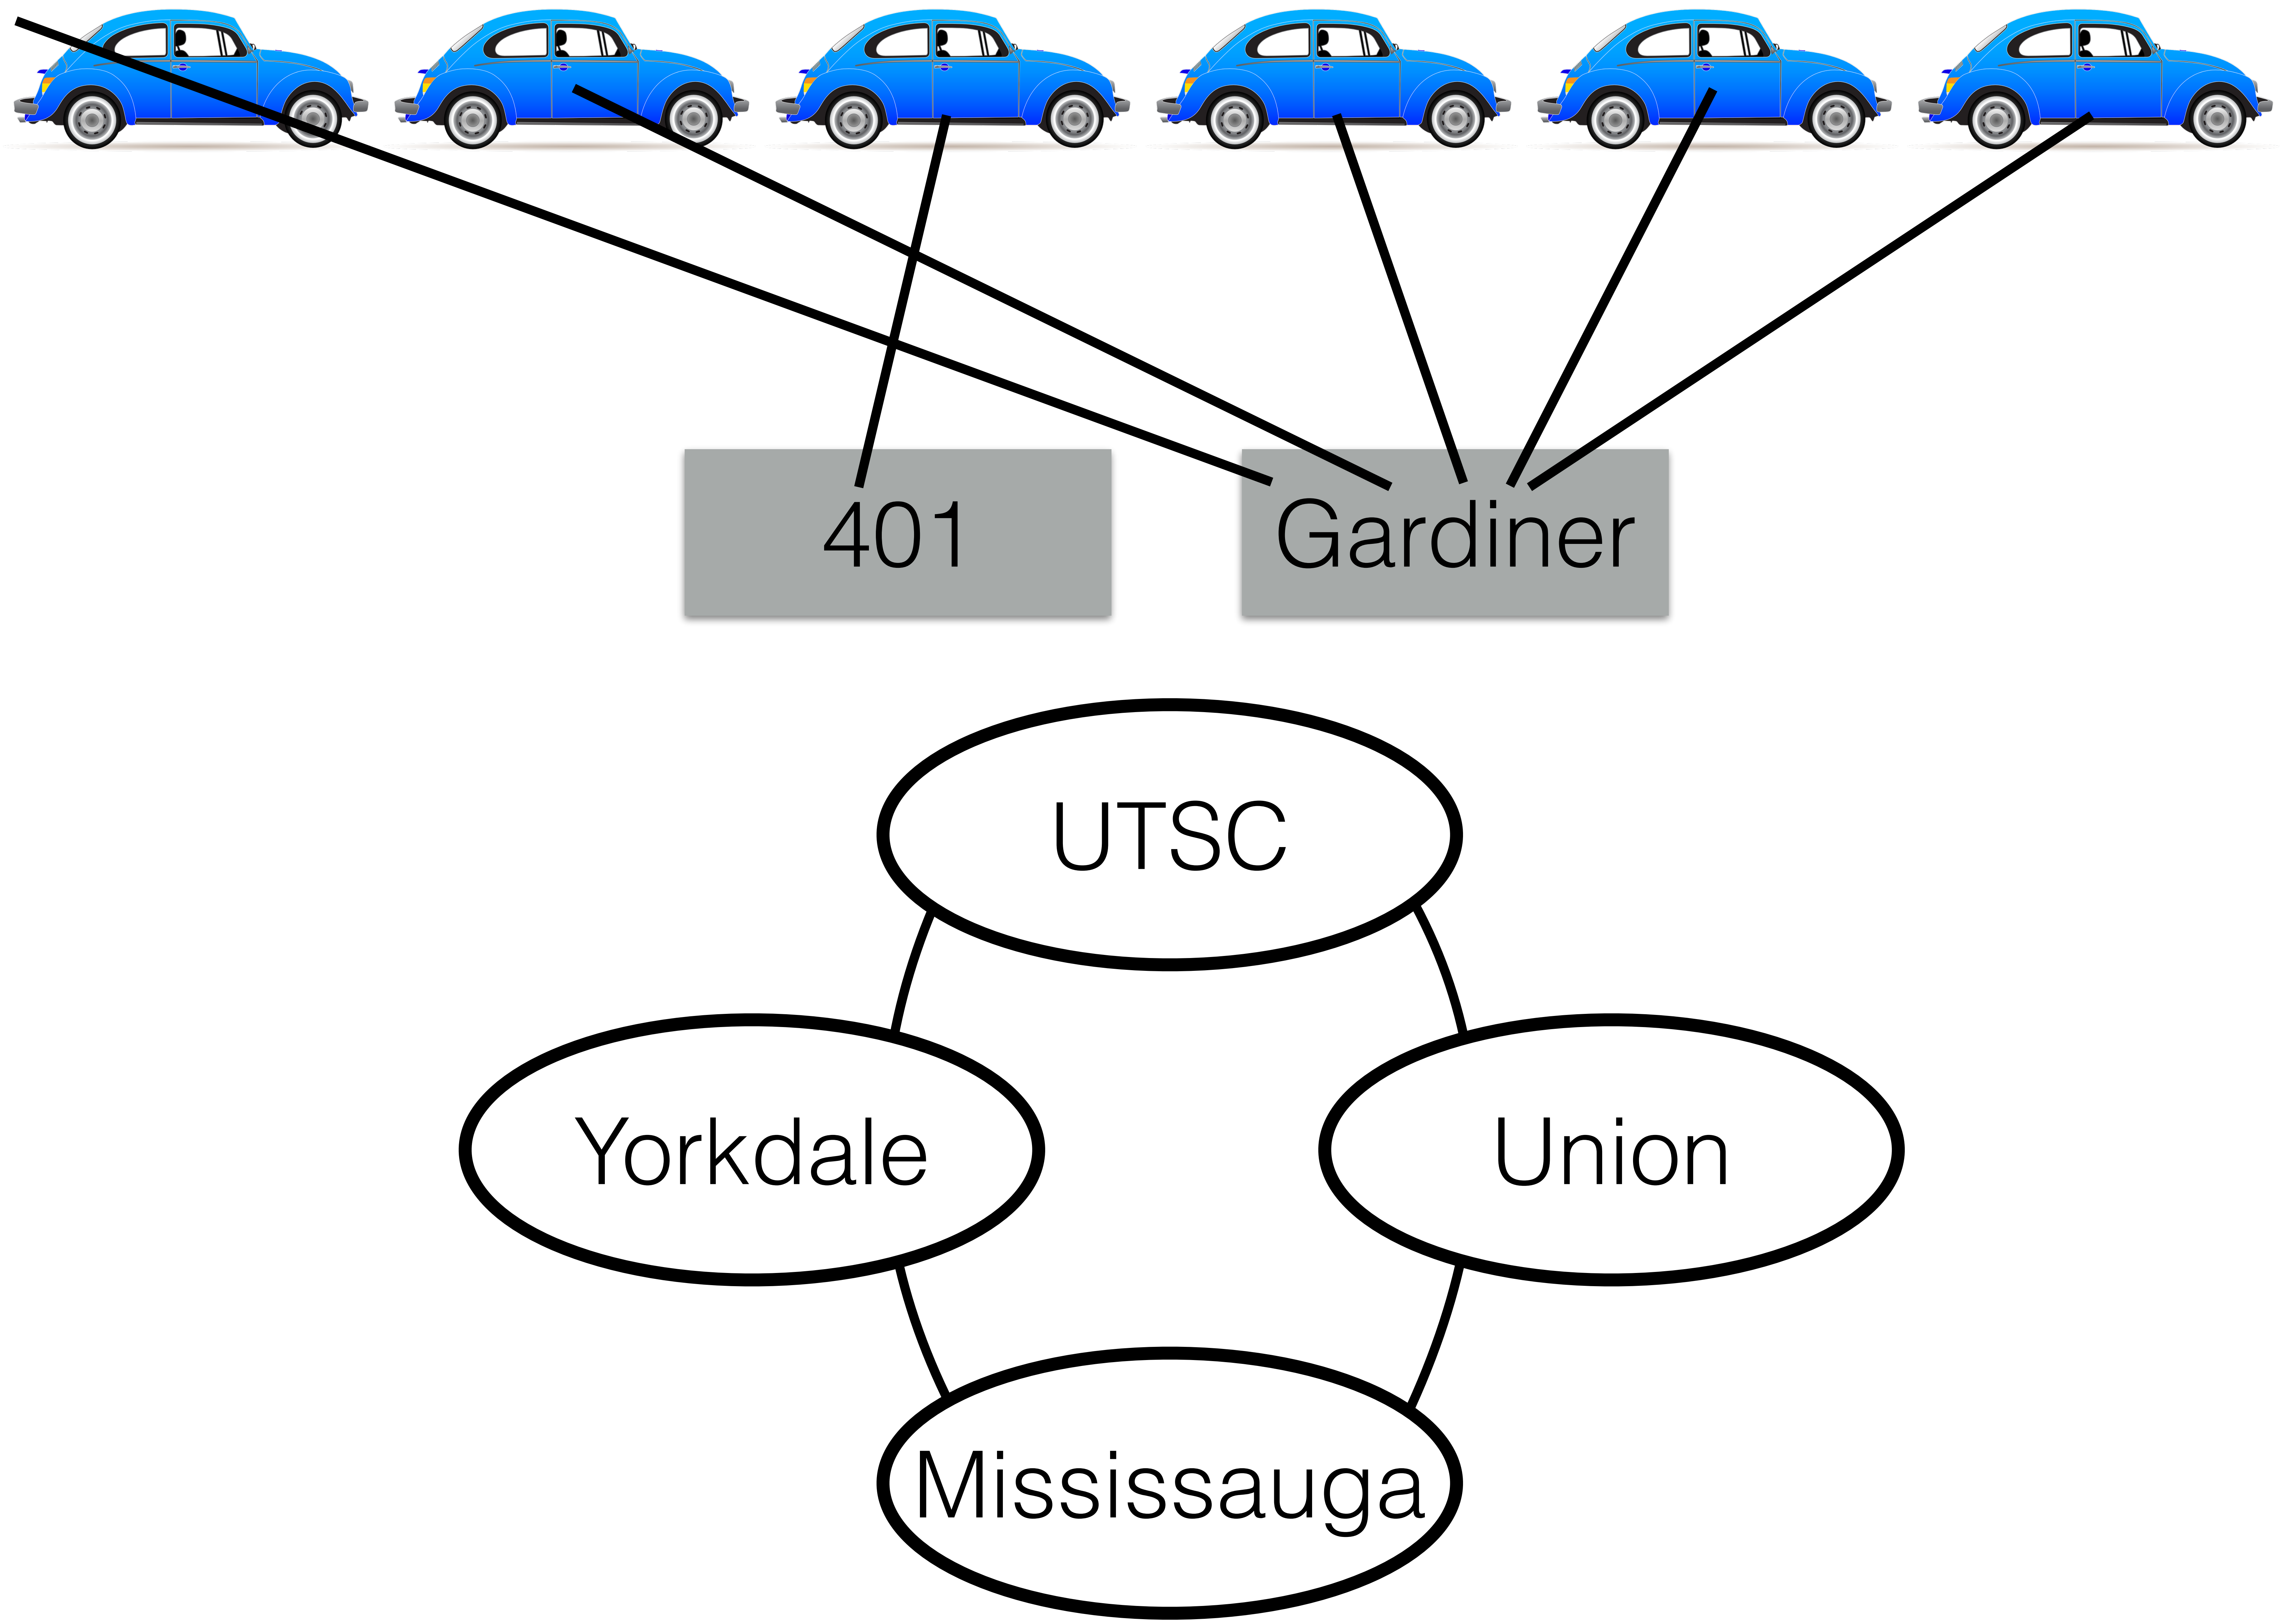

Getting to UTSC: 401 or Gardiner?

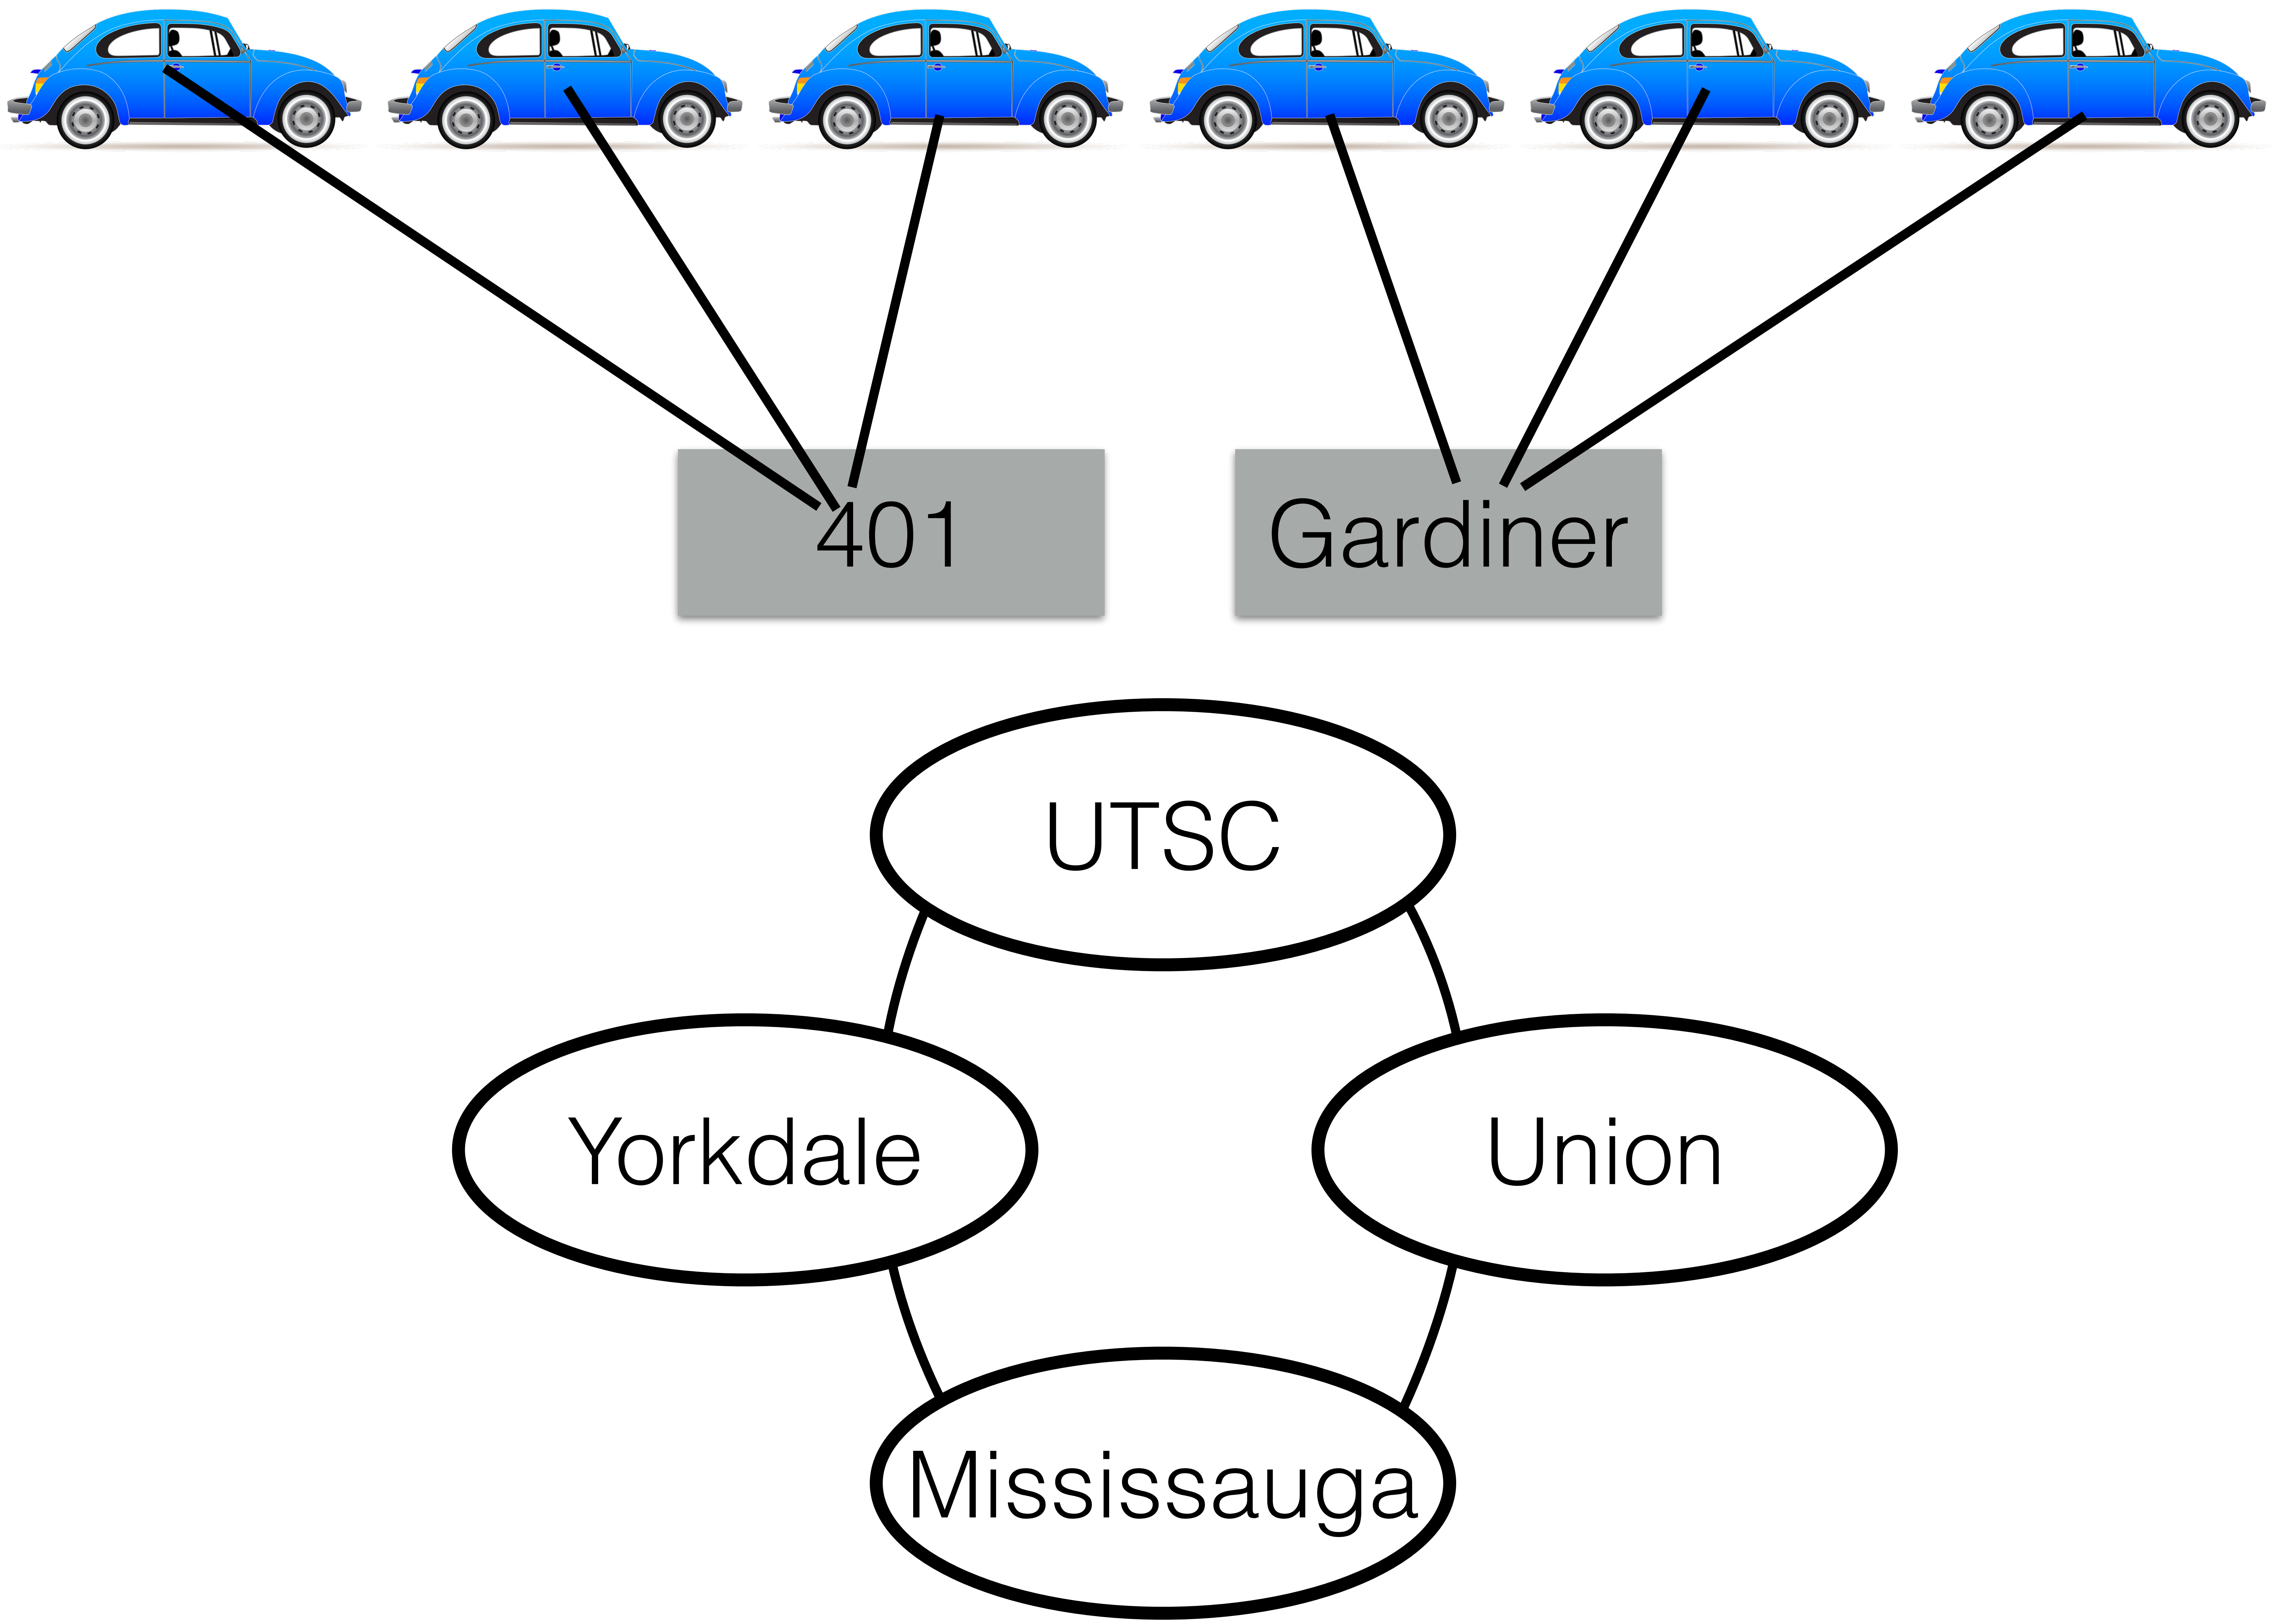

Traffic routing

Let's model this as a simple network, with two kinds of edges:

Constant edges (wide highways that don't get congested)

Traffic-dependent edges (quick routes that can get congested)

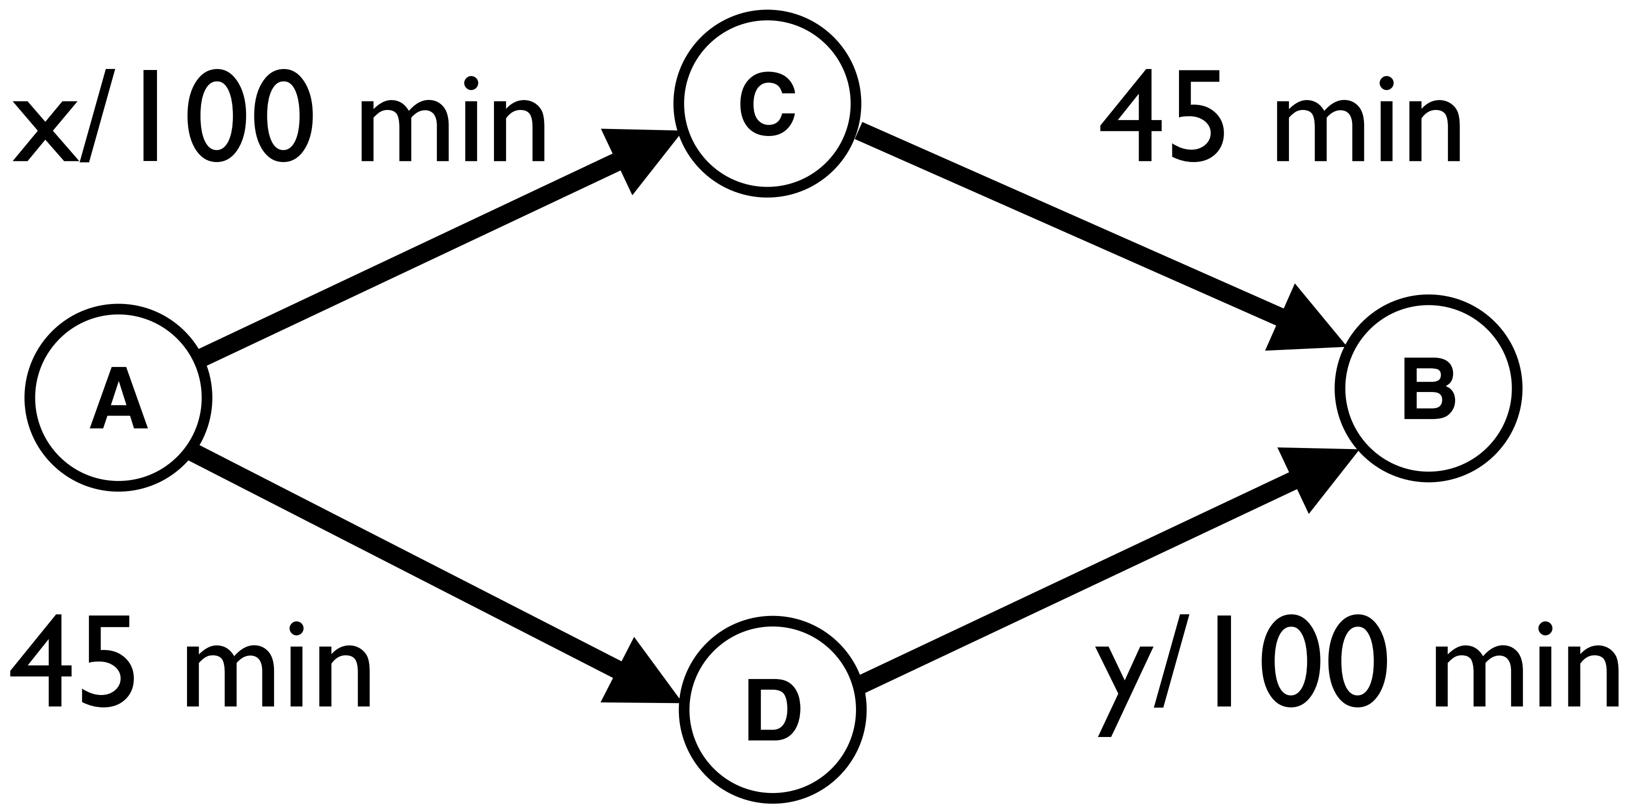

Traffic routing

Let's model this as a simple game on a network, with two kinds of edges:

Constant edges (wide highways that don't get congested)

Traffic-dependent edges (quick routes that can get congested)

There are 4000 drivers. Each one can choose A-C-B or A-D-B.

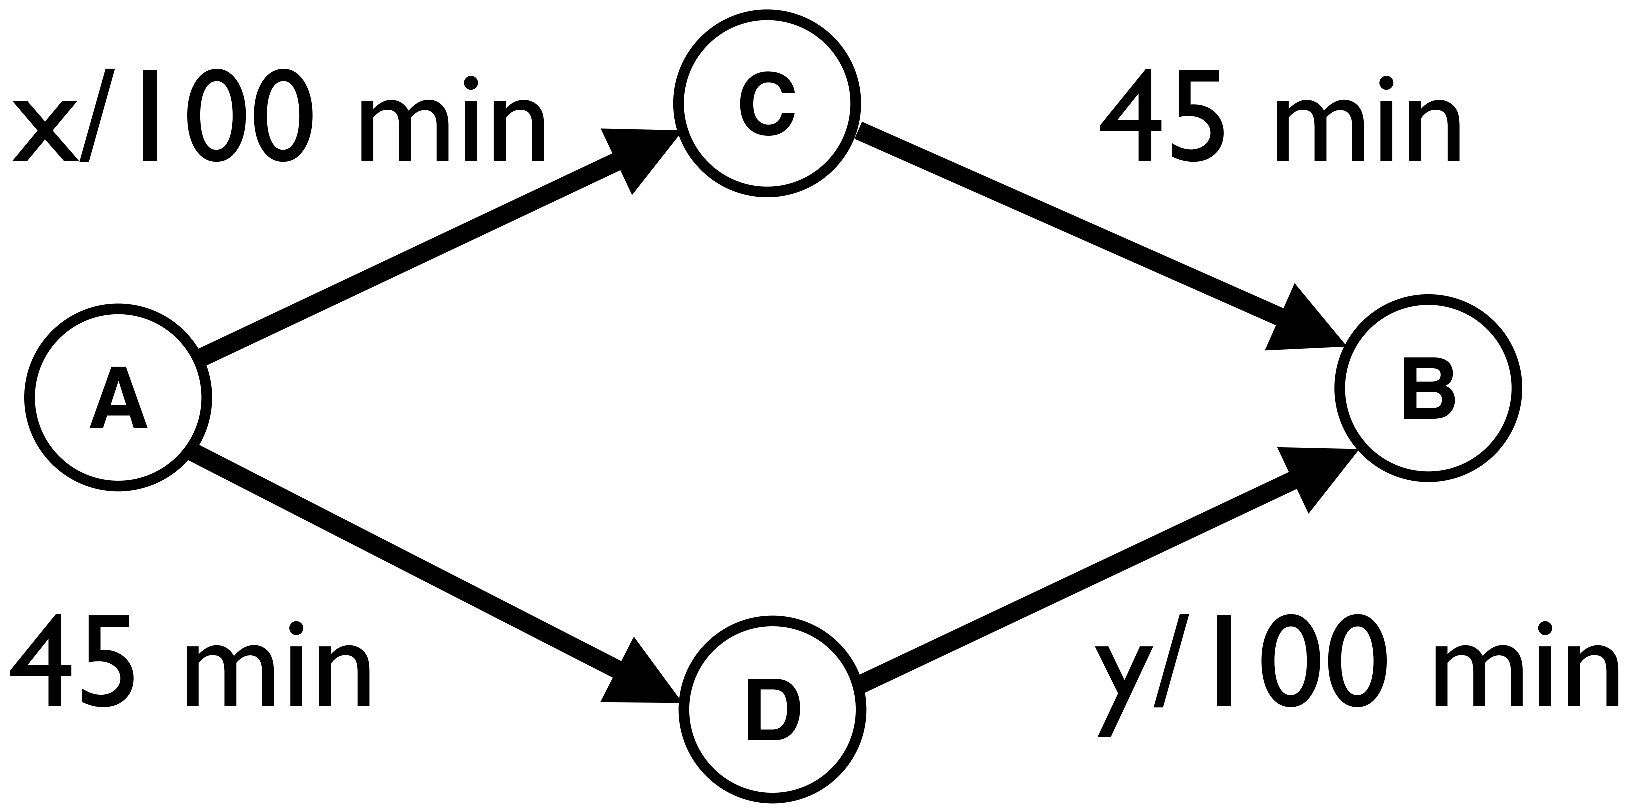

Traffic modeled as a game

Players: Drivers 1,2,3…,4000

Strategies: Two strategies each: A-C-B or A-D-B

Payoffs: ?

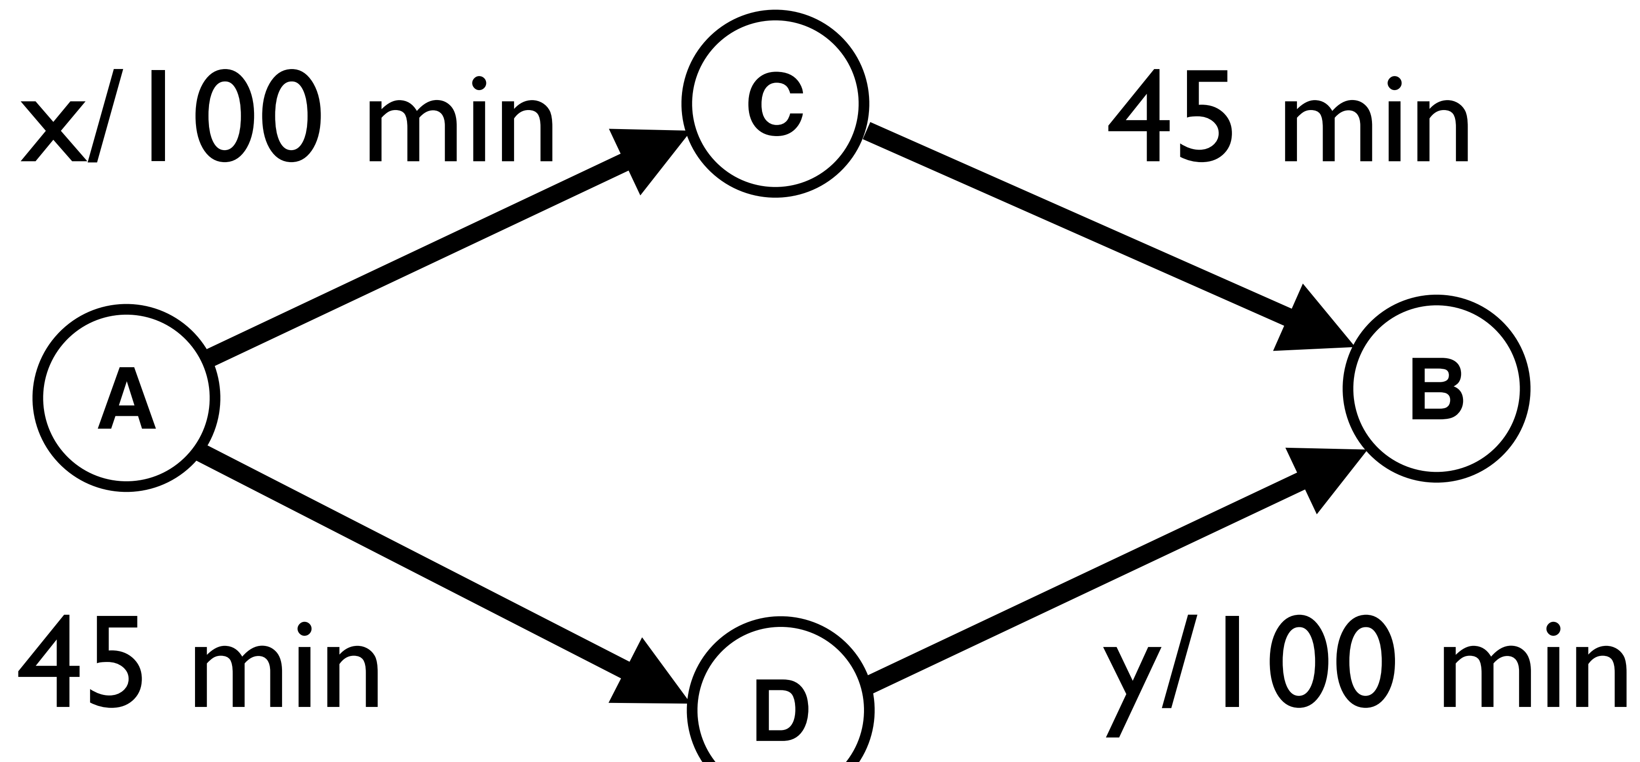

Traffic modeled as a game

Players: Drivers 1,2,3…,4000

Strategies: Two strategies each: A-C-B or A-D-B

Payoffs: Negative drive time

A-C-B time: - (x/100 + 45)

A-D-B time: - (45 + y/100)

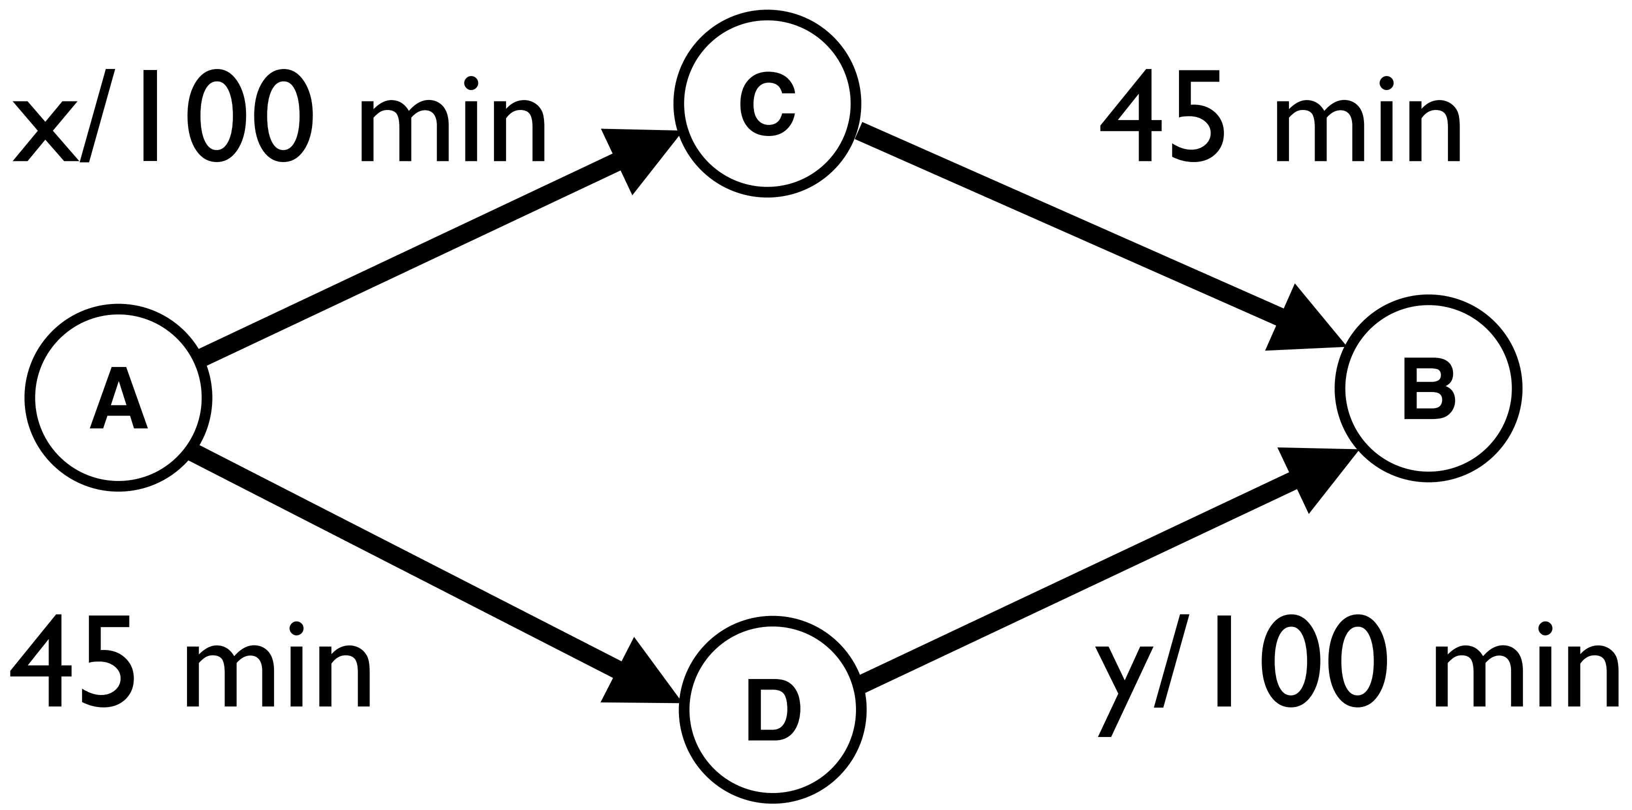

Traffic Equilibrium?

- 4000 drivers
- Two route options: A-C-B or A-D-B.
- Consider a few outcomes (strategy for each player):
  - Payoffs when 4000 choose top (ACB), 0 choose bottom (ADB):
    - Top path: 4000/100 + 45 = 85 min
    - Bottom path: 45 + 0/100 = 45 min
  - Payoffs when 0 choose top, 4000 choose bottom:
    - Top: 0/100 + 45 = 45 min
    - Bottom: 45 + 4000/100 = 85 min

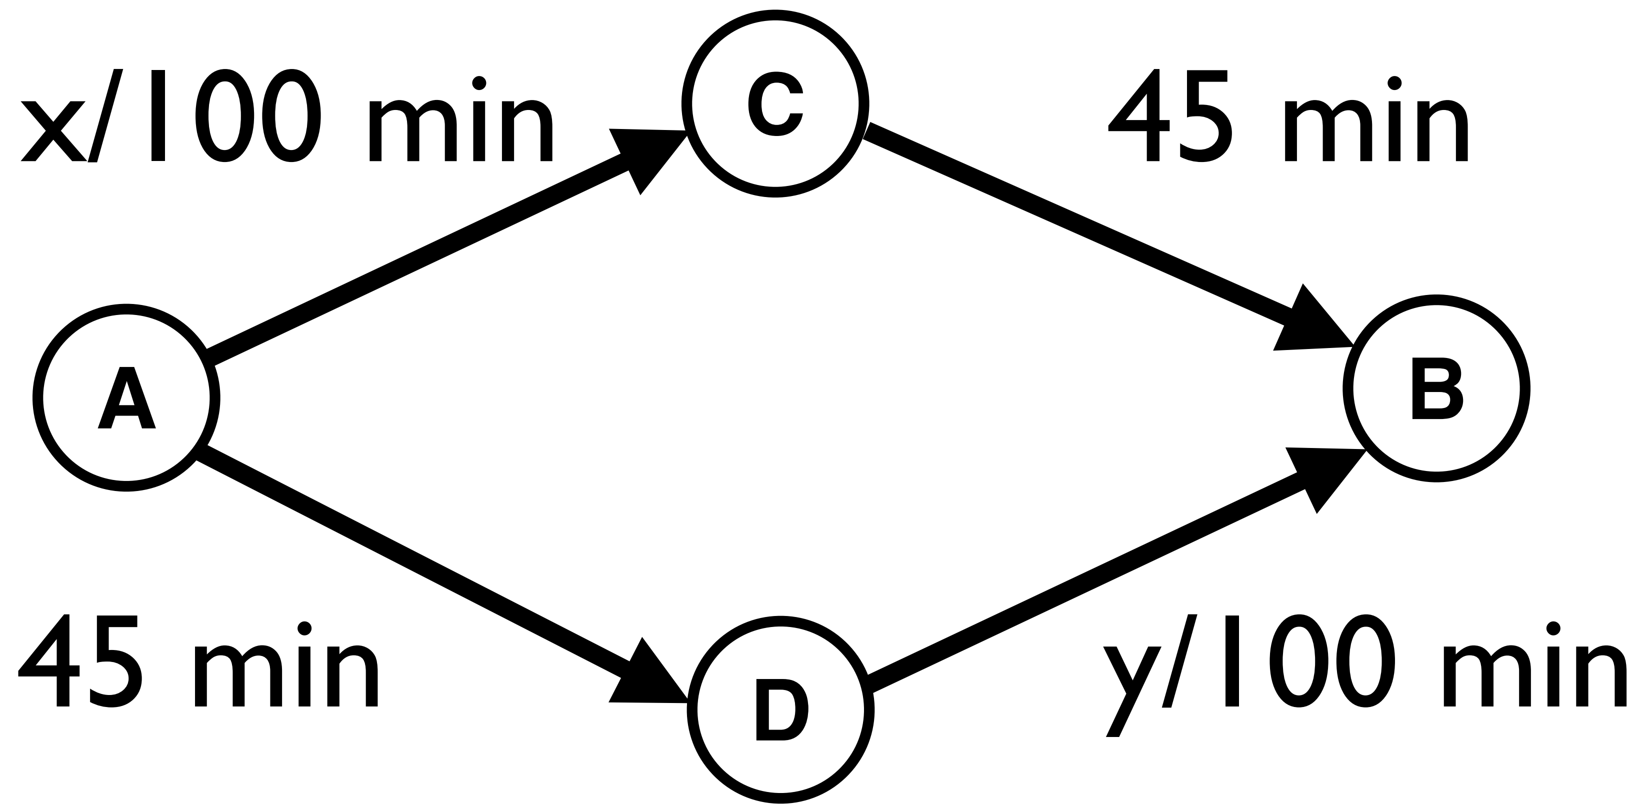

Equilibrium in traffic?

- 4000 drivers
- Two route options: A-C-B or A-D-B.
- Payoffs when 2000 choose top, 2000 choose bottom:
  - Top: 2000/100 + 45 = 65 min
  - Bottom: 45 + 2000/100 = 65 min

This is an equilibrium because no one has an incentive to deviate

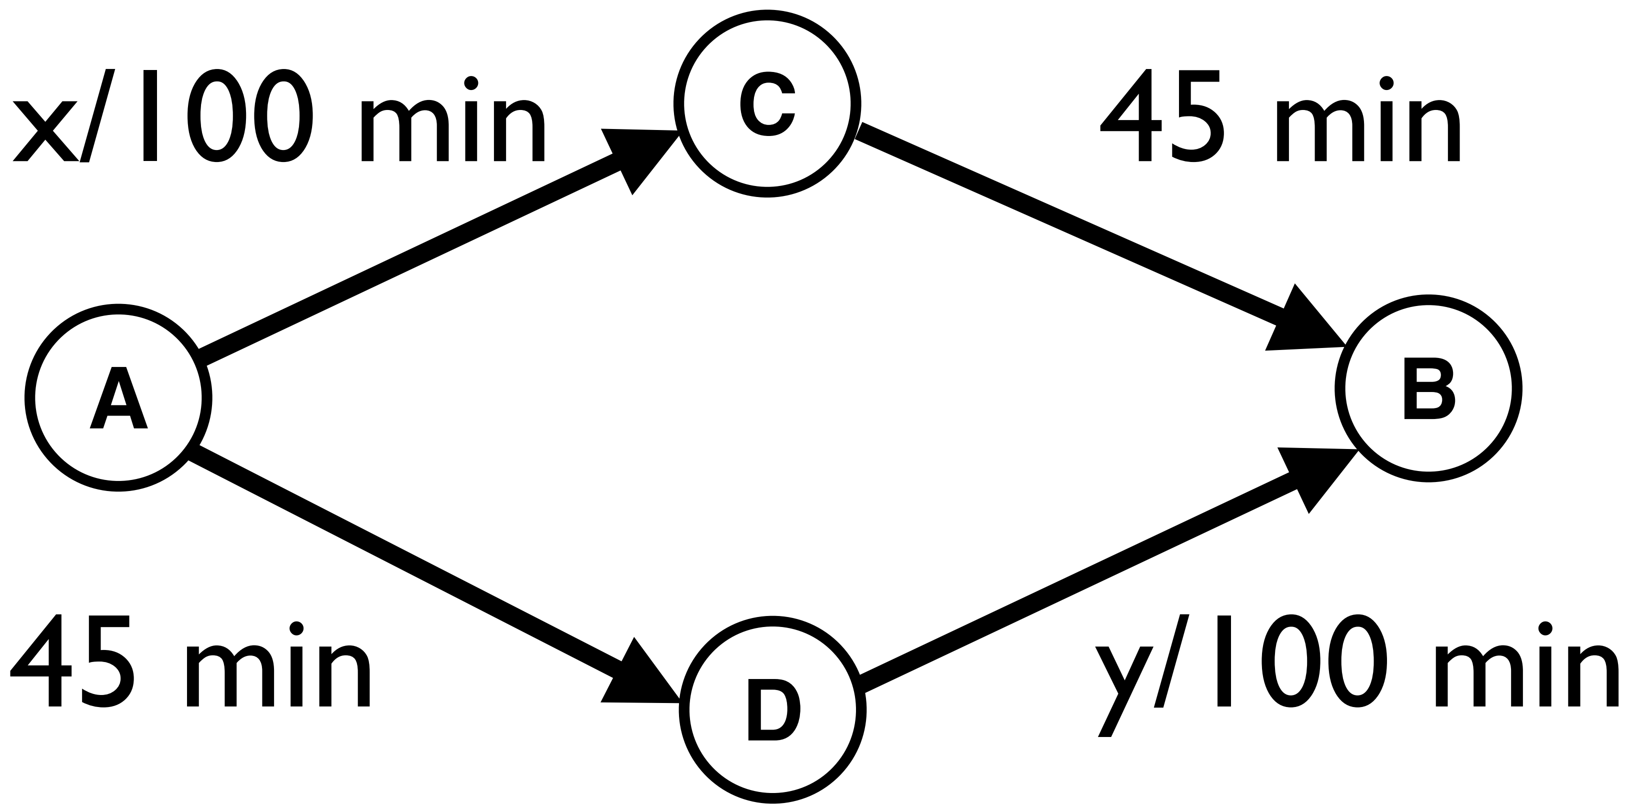

Equilibrium in traffic?

Payoffs when 2000 choose top, 2000 choose bottom:

Top: 2000/100 + 45 = 65 min

Bottom: 45 + 2000/100 = 65 min

This is an equilibrium because no one has an incentive to deviate

If someone currently using A-C-B decides to switch to A-D-B:

Currently: Top: 2000/100 + 45 = 65.00 min

Switch: Bottom: 45 + 2001/100 = 65.01 min

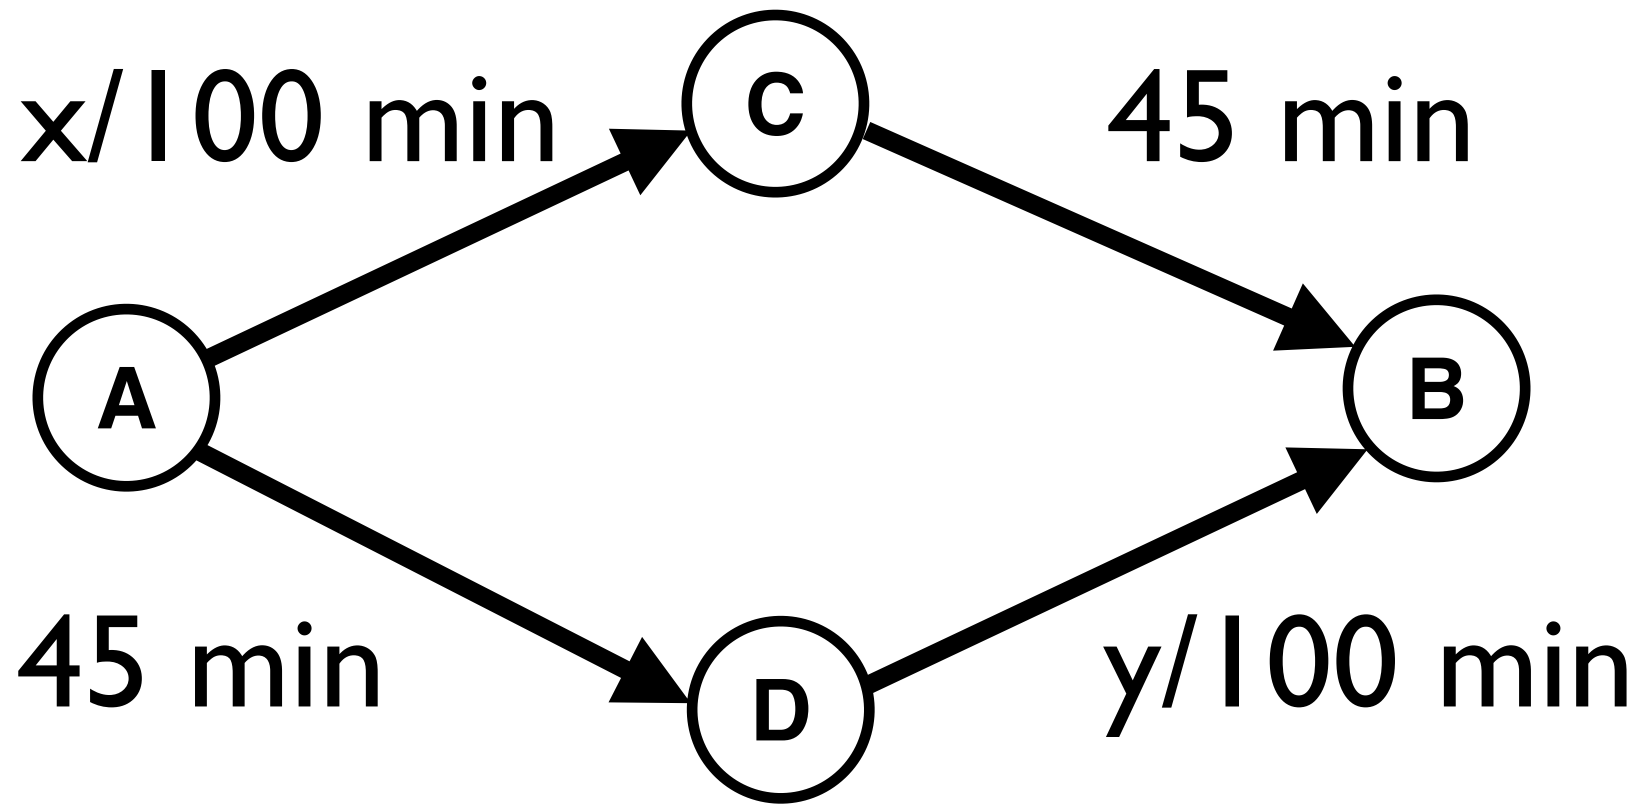

Traffic modeled as a game

Players: Drivers 1,2,3…,4000

Strategies: A-C-B, A-D-B

Payoffs: Negative drive time

A-C-B time: -(x/100 + 45)

A-D-B time: -(45 + y/100)

You want to lower your drive time, so we take the negative drive time as the "payoff"

Notice that this actually describes many equilibria: any set of strategies "2000 choose top, 2000 choose bottom" is an equilibrium (players are interchangeable, so any set of 2000 can be using ACB and any set of 2000 can be using ADB)

For any other set of strategies, deviation benefits someone (therefore isn't an equilibrium)

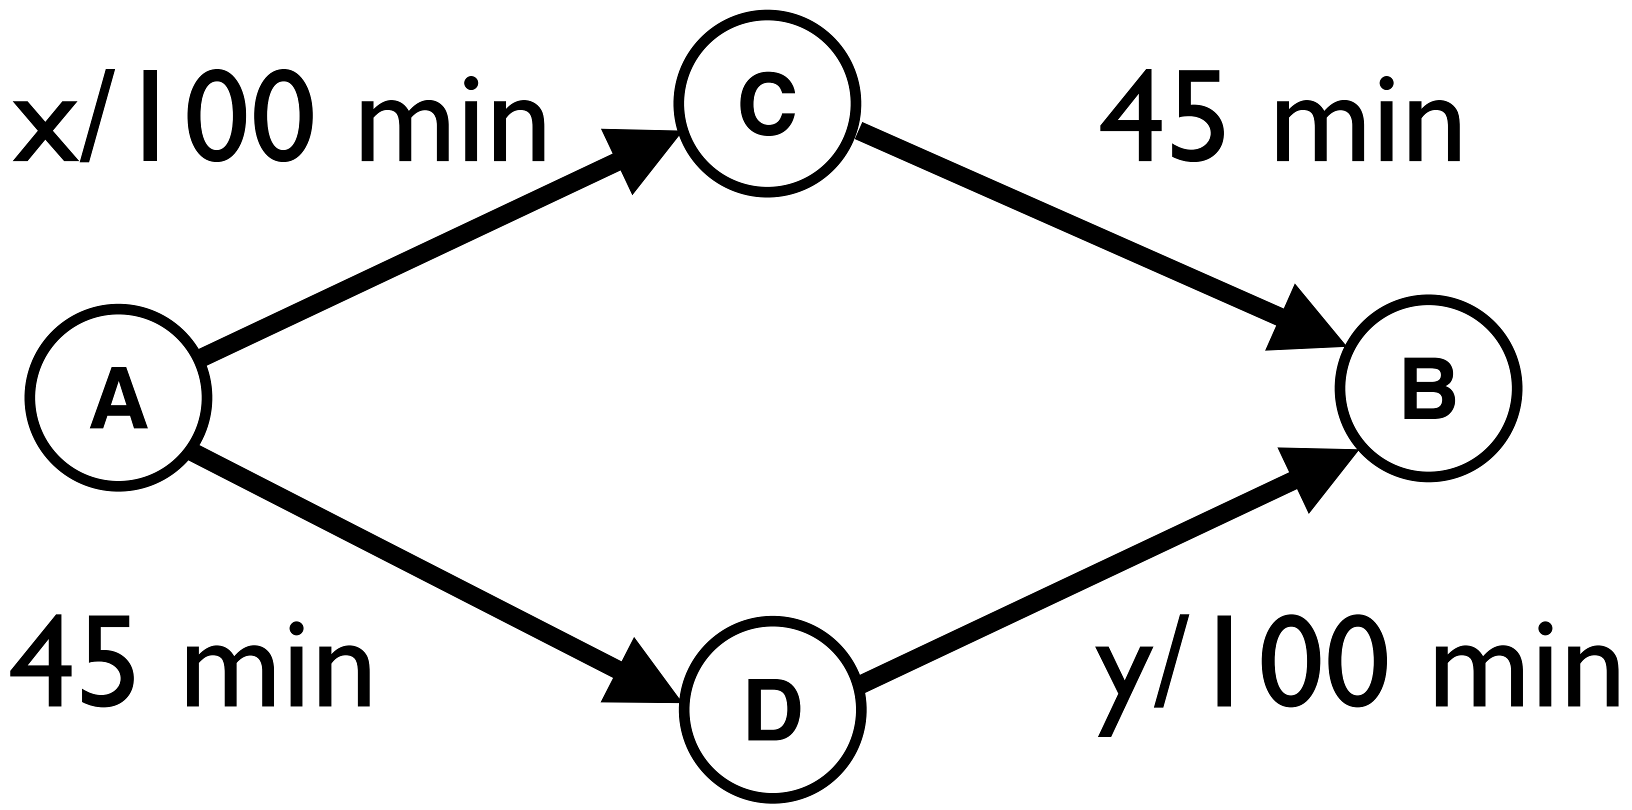

Traffic modeled as a game

Now Elon Musk adds a **teleport**!

Players can take it if they want — or not

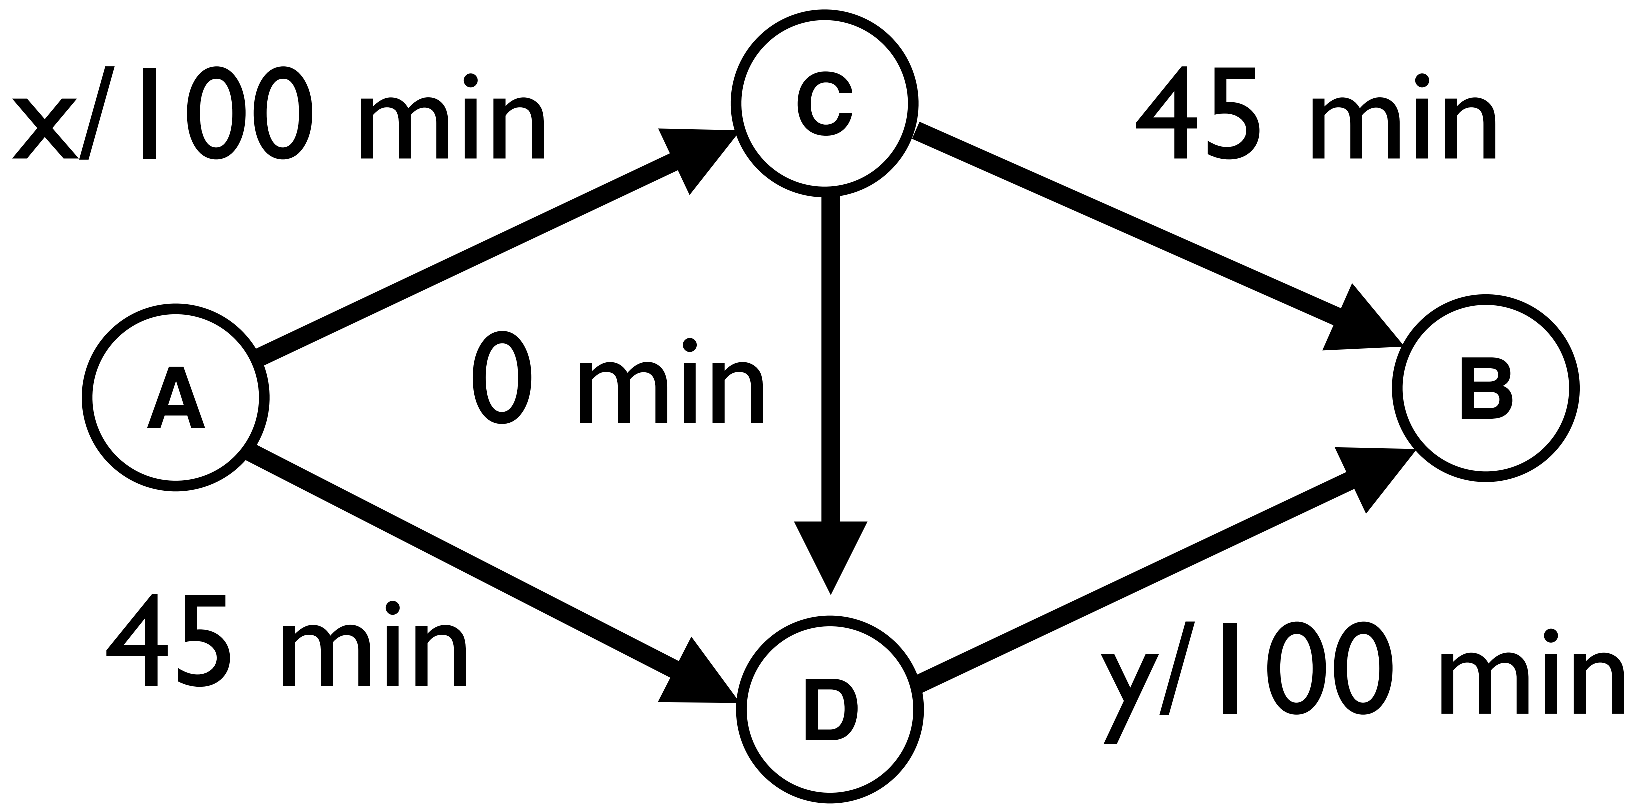

Traffic modeled as a game

Players: Drivers 1,2,3…,4000

Strategies: A-C-B, A-D-B, A-C-D-B

Payoffs: Negative drive time

A-C-B time: - (x/100 + 45)

A-D-B time: - (45 + y/100)

A-C-D-B time: - (x/100 + y/100)

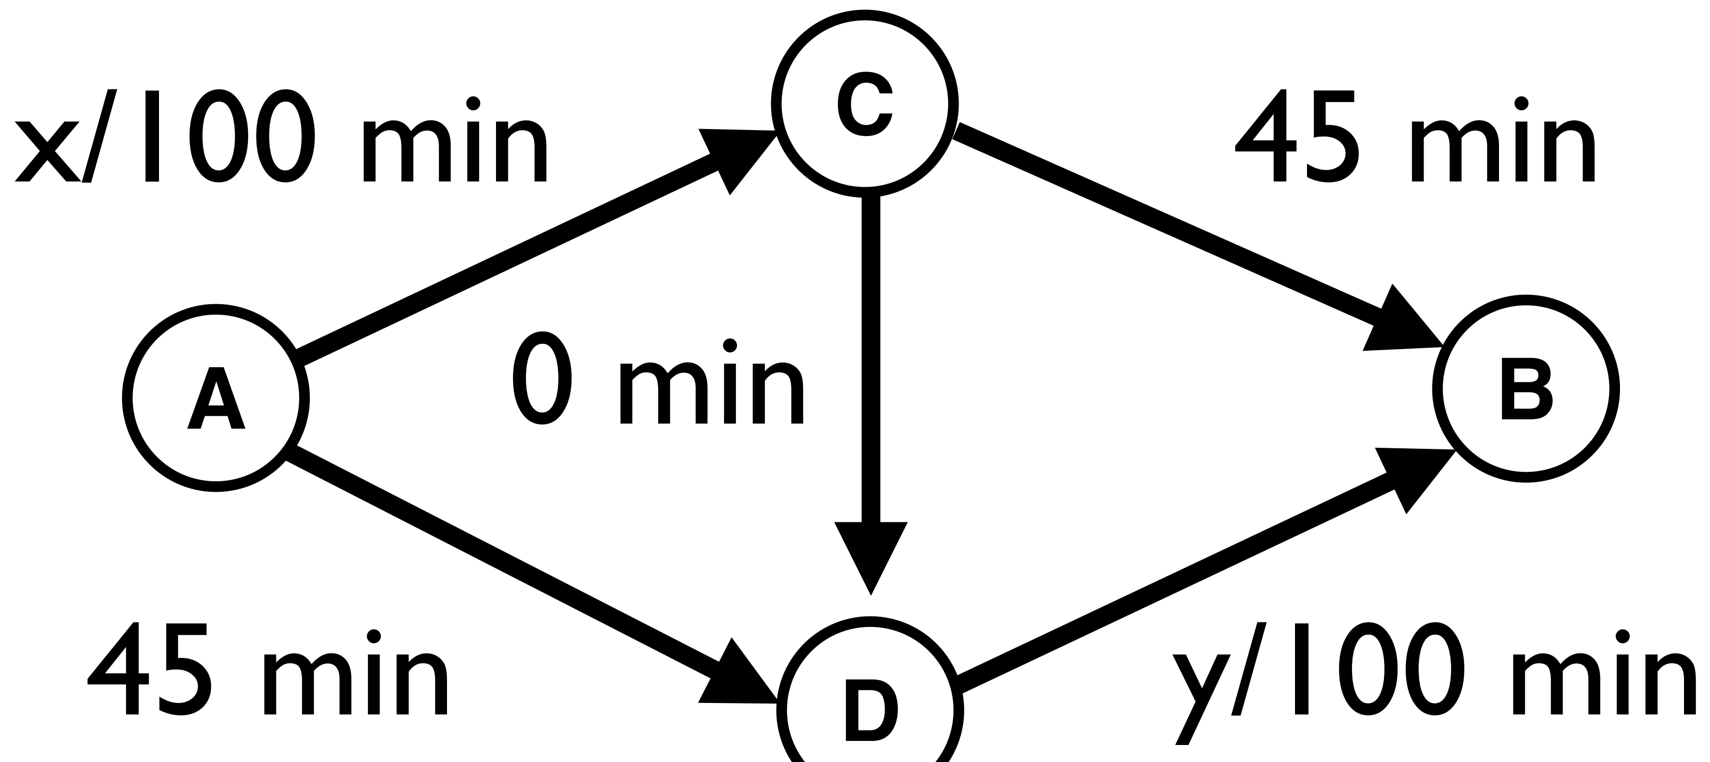

Would you teleport?

Say we are at the equilibrium from before: 2000 ACB, 2000 ADB, 0 ACDB

A-C-B time: - (x/100 + 45)

2000/100 + 45 = 65 minutes

A-D-B time: - (45 + y/100)

2000/100 + 45 = 65 minutes

A-C-D-B time: - (x/100 + y/100)

2000/100 + 2000/100 = 40 minutes

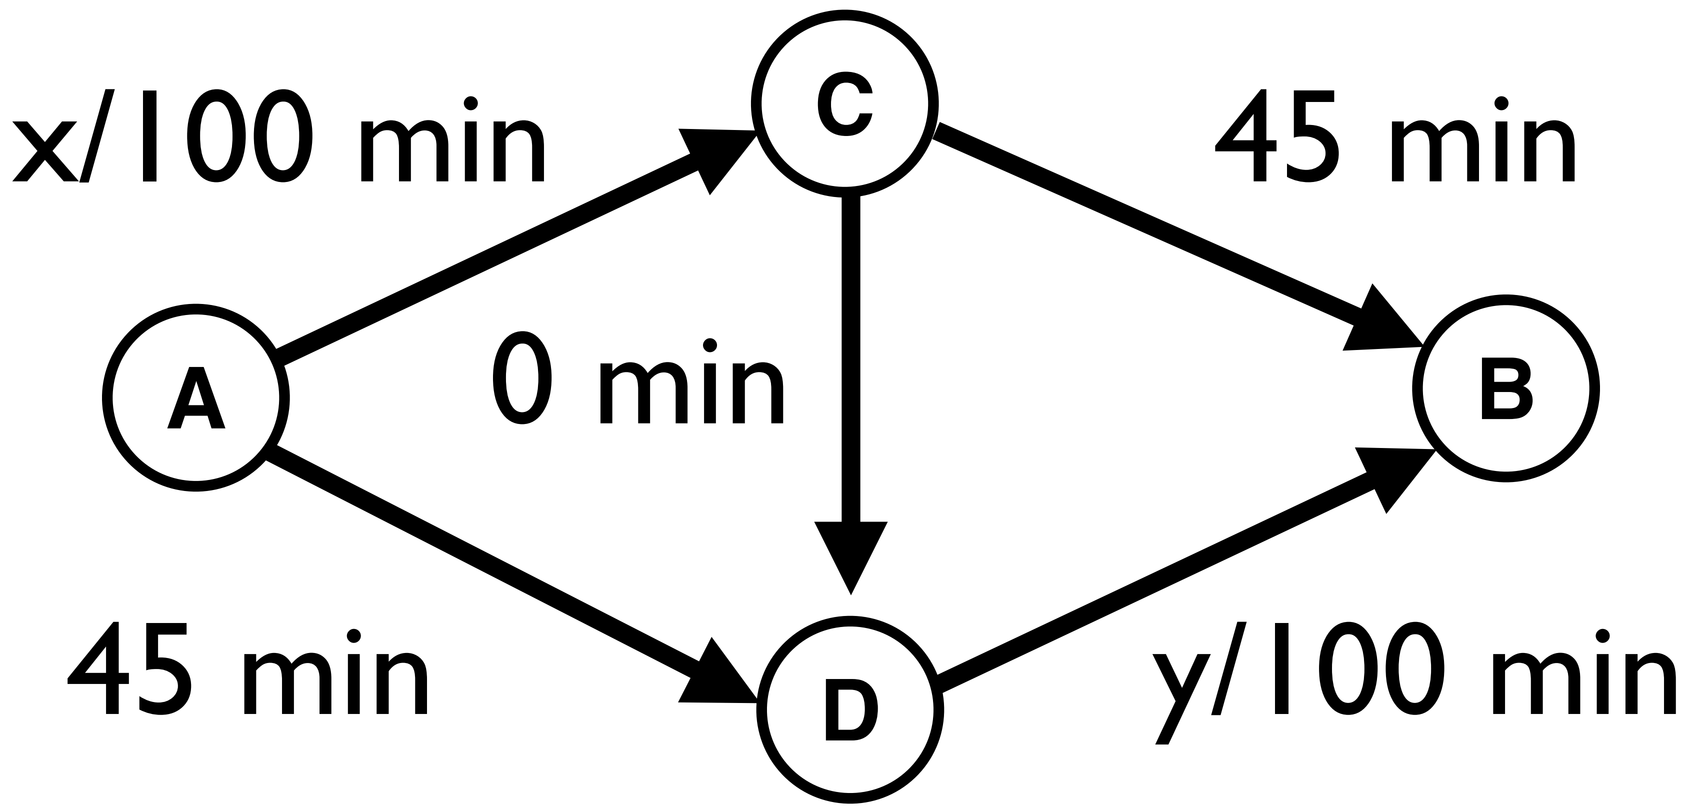

New equilibrium?

Payoffs when 0 ACB, 0 ADB, 4000 ACDB

A-C-B time: - (x/100 + 45)

A-D-B time: - (45 + y/100)

A-C-D-B time: - (x/100 + y/100)

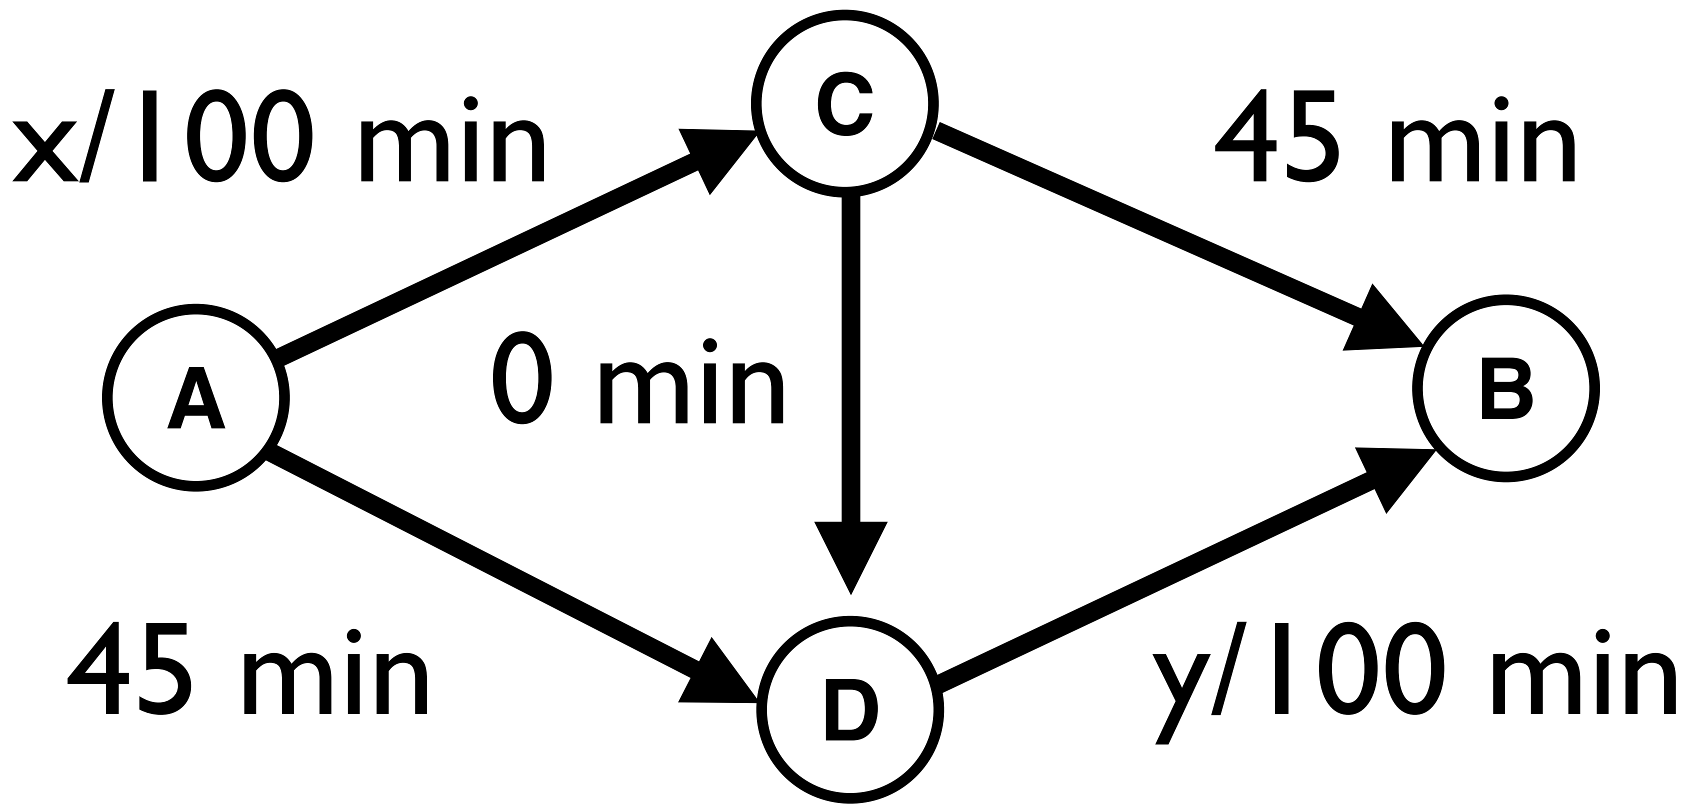

New equilibrium?

Payoffs when 0 ACB, 0 ADB, 4000 ACDB

A-C-B time: - (x/100 + 45)

4000/100 + 45 = 85 minutes

A-D-B time: - (45 + y/100)

45 + 4000/100 = 85 minutes

A-C-D-B time: - (x/100 + y/100)

4000/100 + 4000/100 = 80 minutes

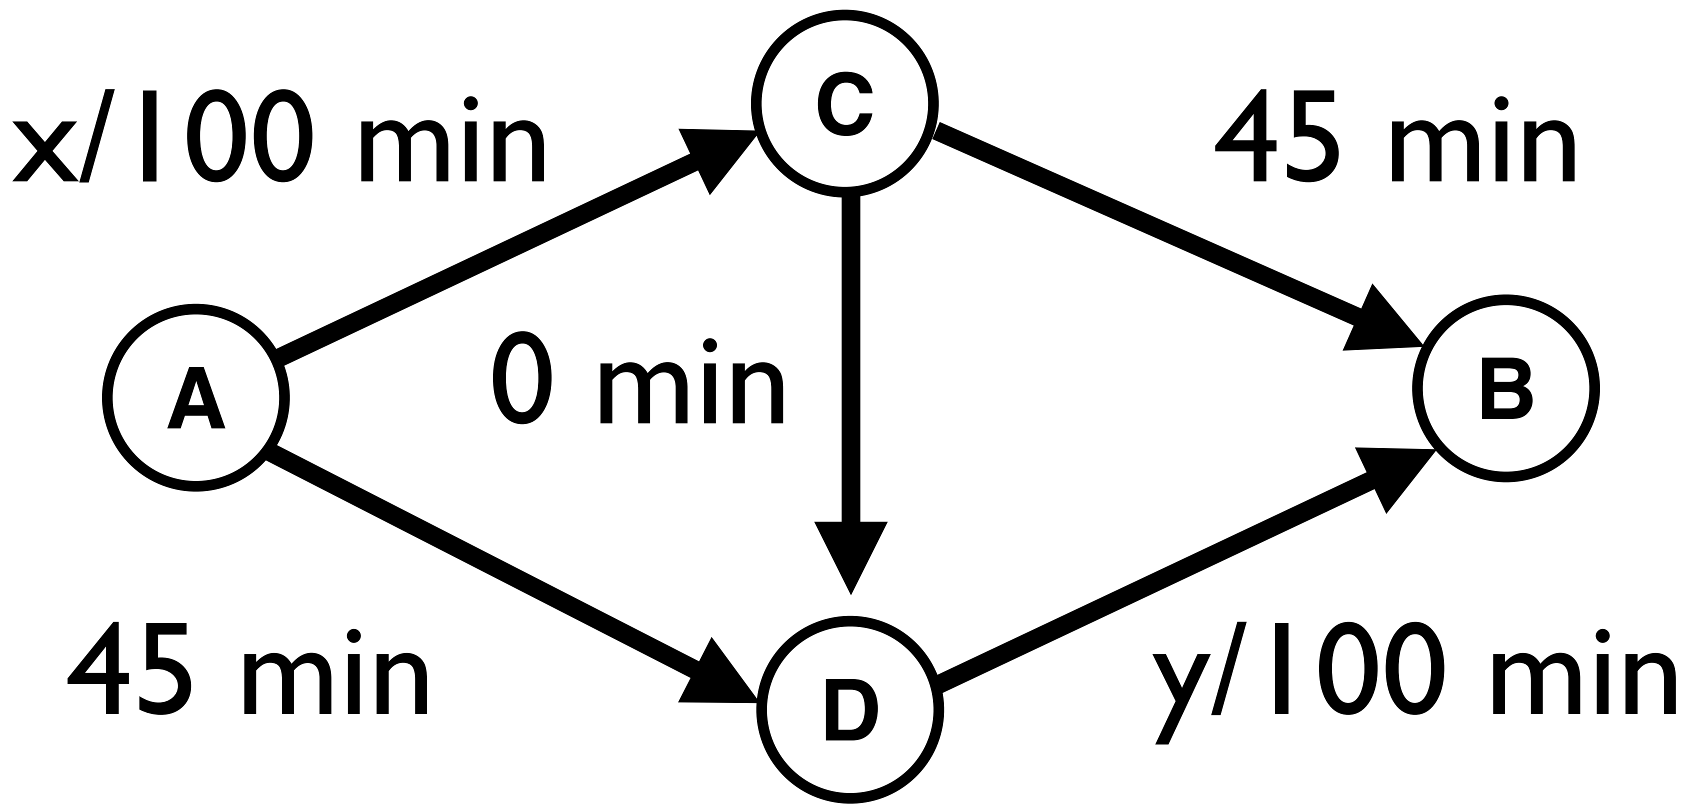

New equilibrium?

Payoffs when 0 ACB, 0 ADB, 4000 ACDB

A-C-B time: - (x/100 + 45) = 4000/100 + 45 = 85 minutes

A-D-B time: - (45 + y/100) = 45 + 4000/100 = 85 minutes

A-C-D-B time: - (x/100 + y/100) = 4000/100 + 4000/100 = 80 minutes

ACDB is a **strictly dominant strategy**

Everyone playing ACDB is the only equilibrium!

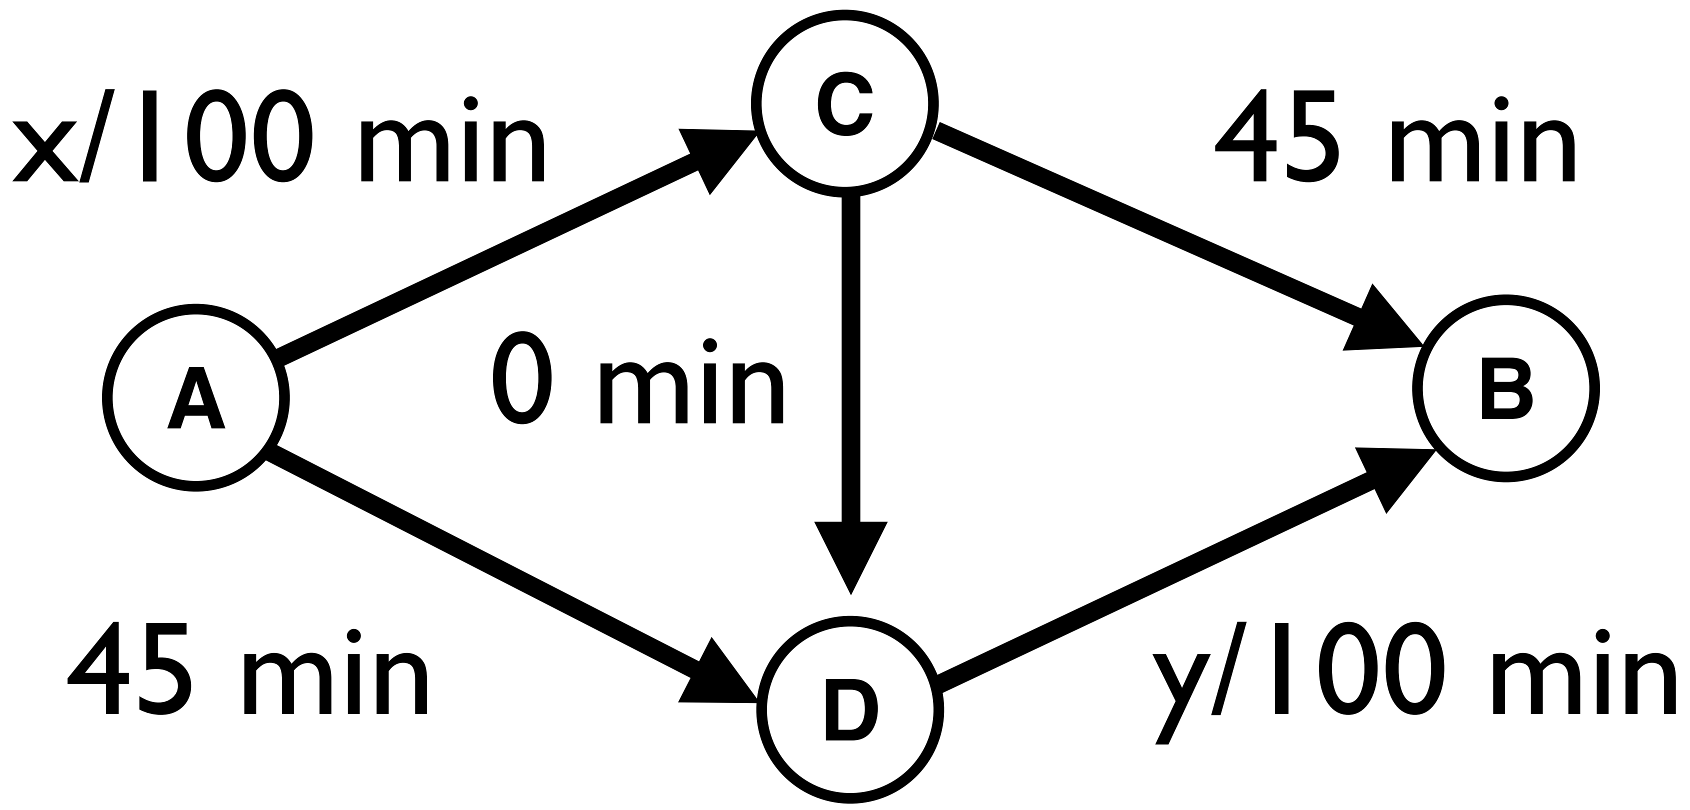

What just happened?

Equilibrium: 65 minutes for everyone

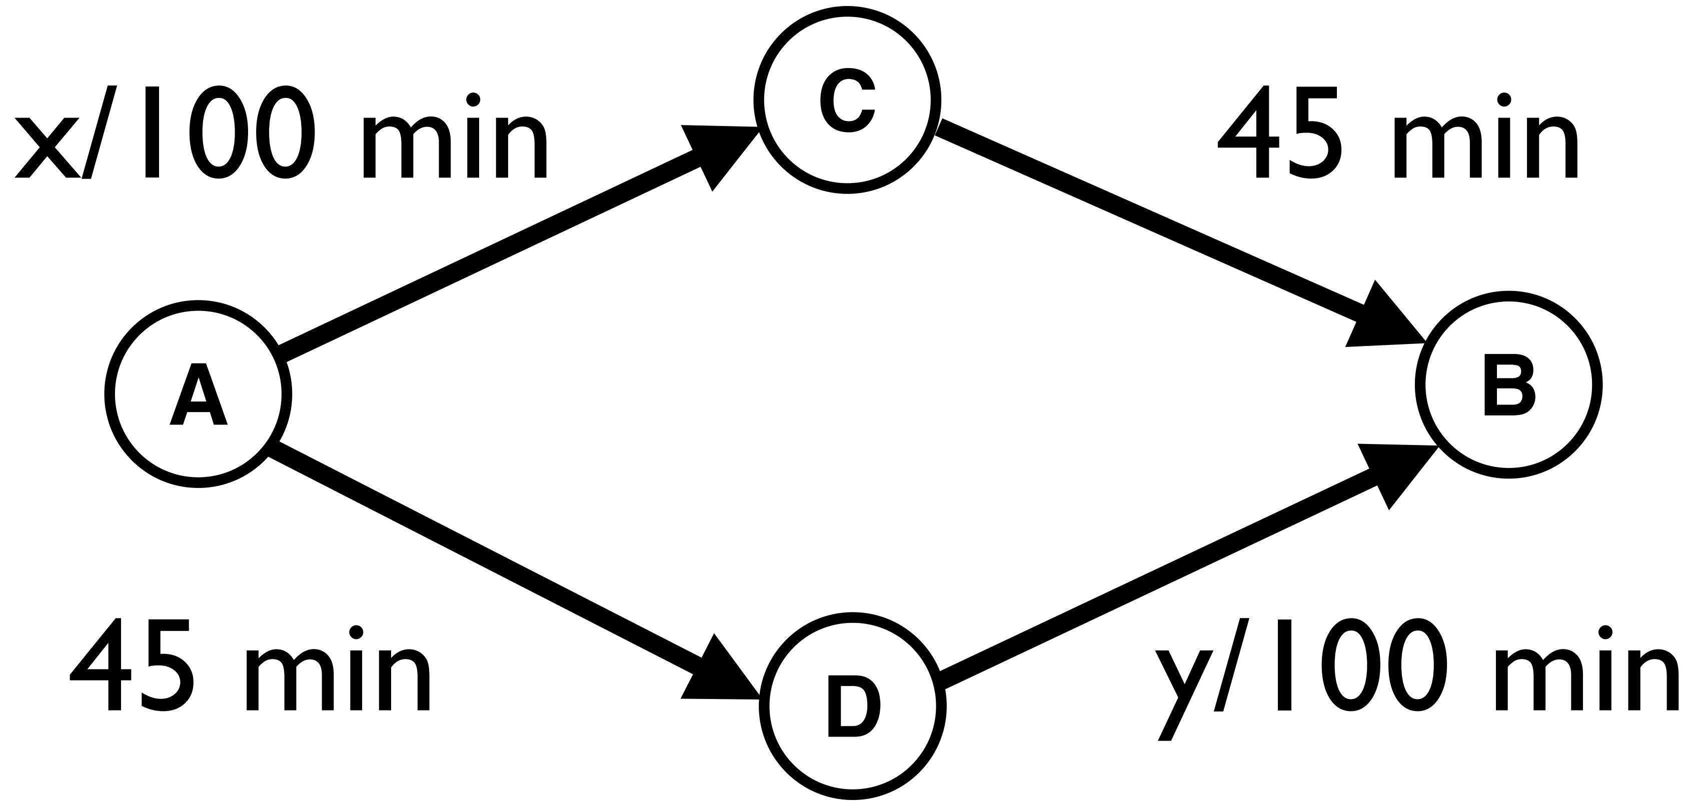

Equilibrium: 80 minutes for everyone

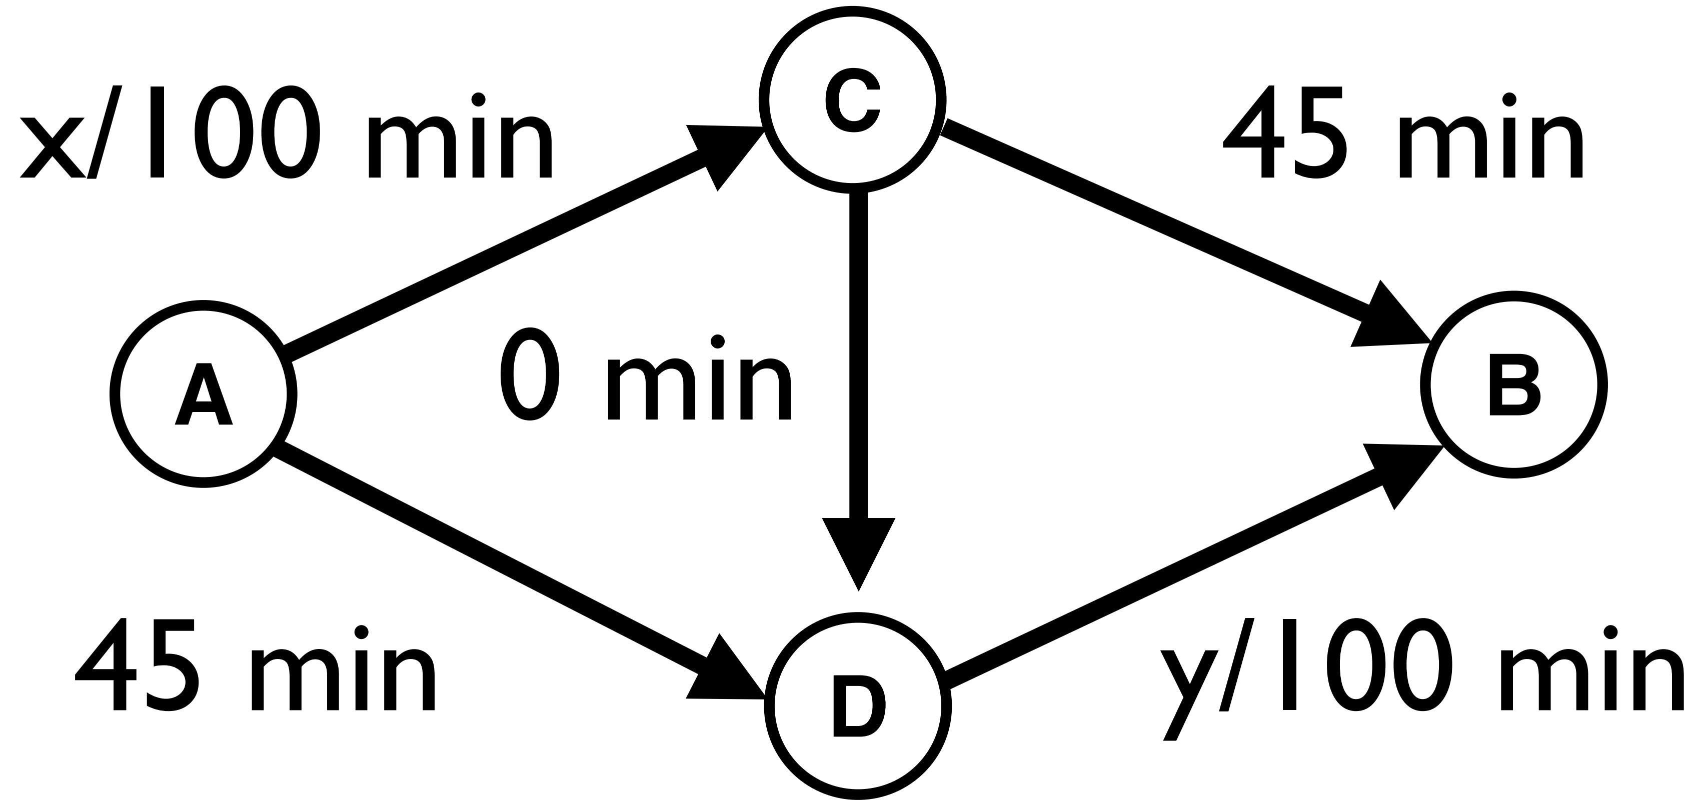

Same network but with an extra teleport

Braess's Paradox

Routing:

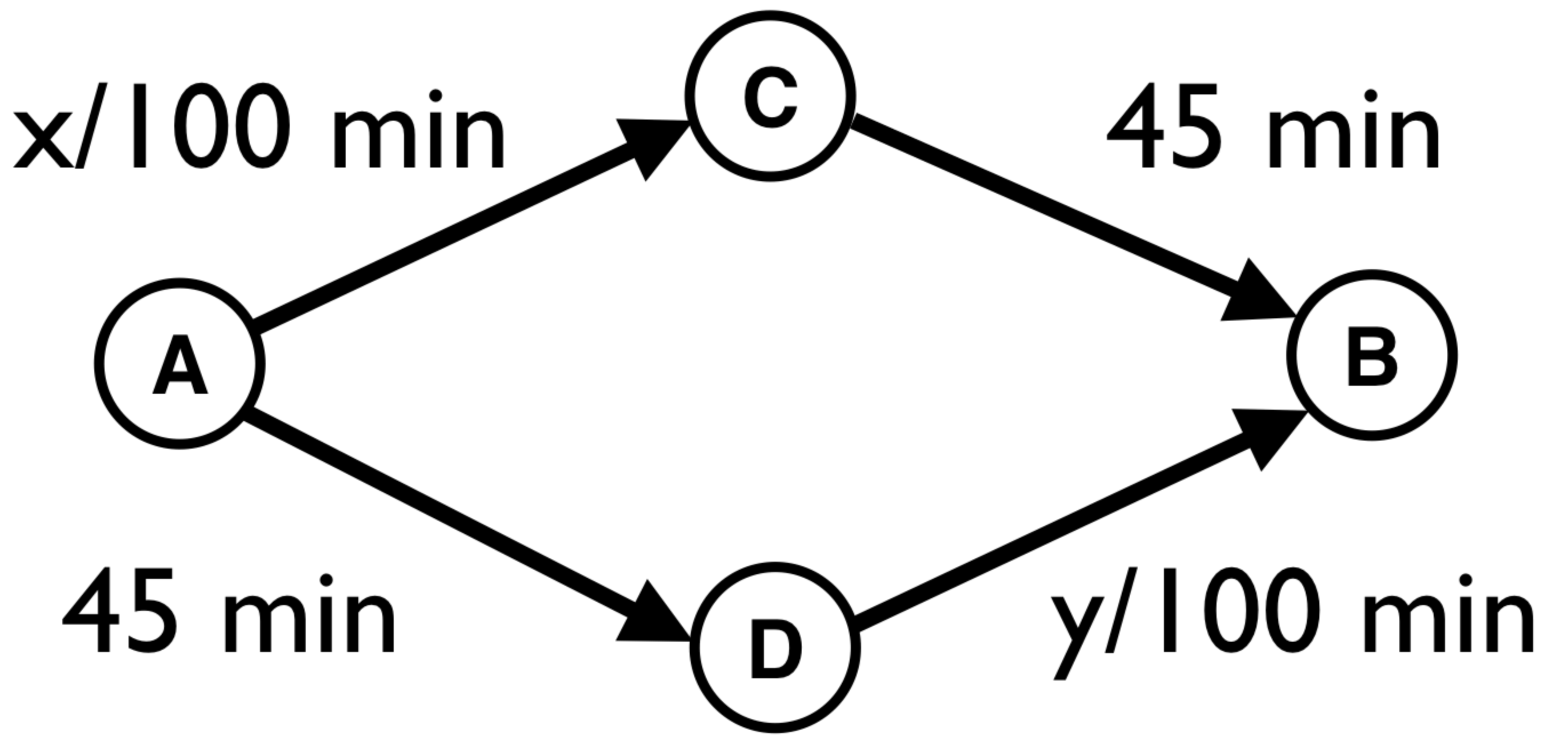

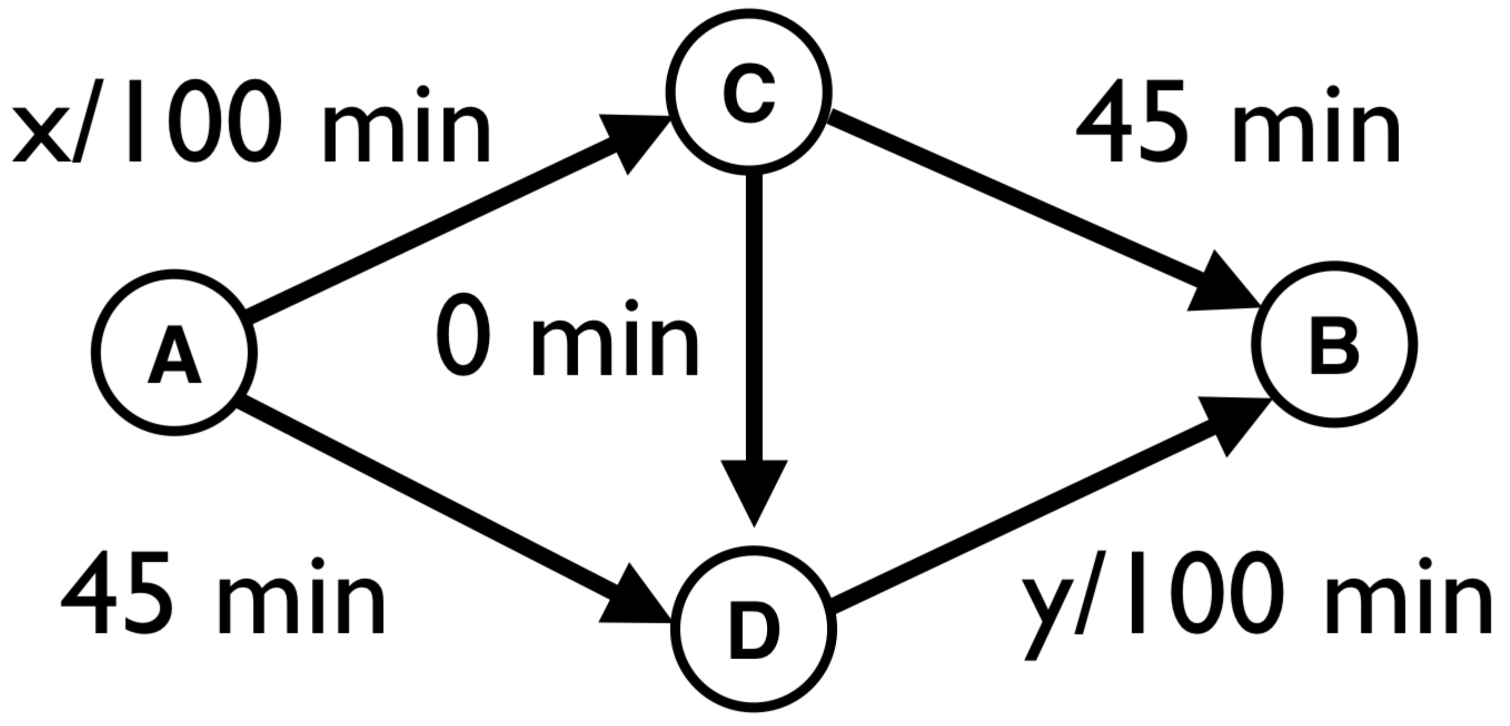

Prisoner's Dilemma:

		Suspect 2	
		NC	C
Suspect 1	NC	$-1, -1$	$-10, 0$
	C	$0, -10$	$-4, -4$

Sometimes strategies can hurt you

Routing:

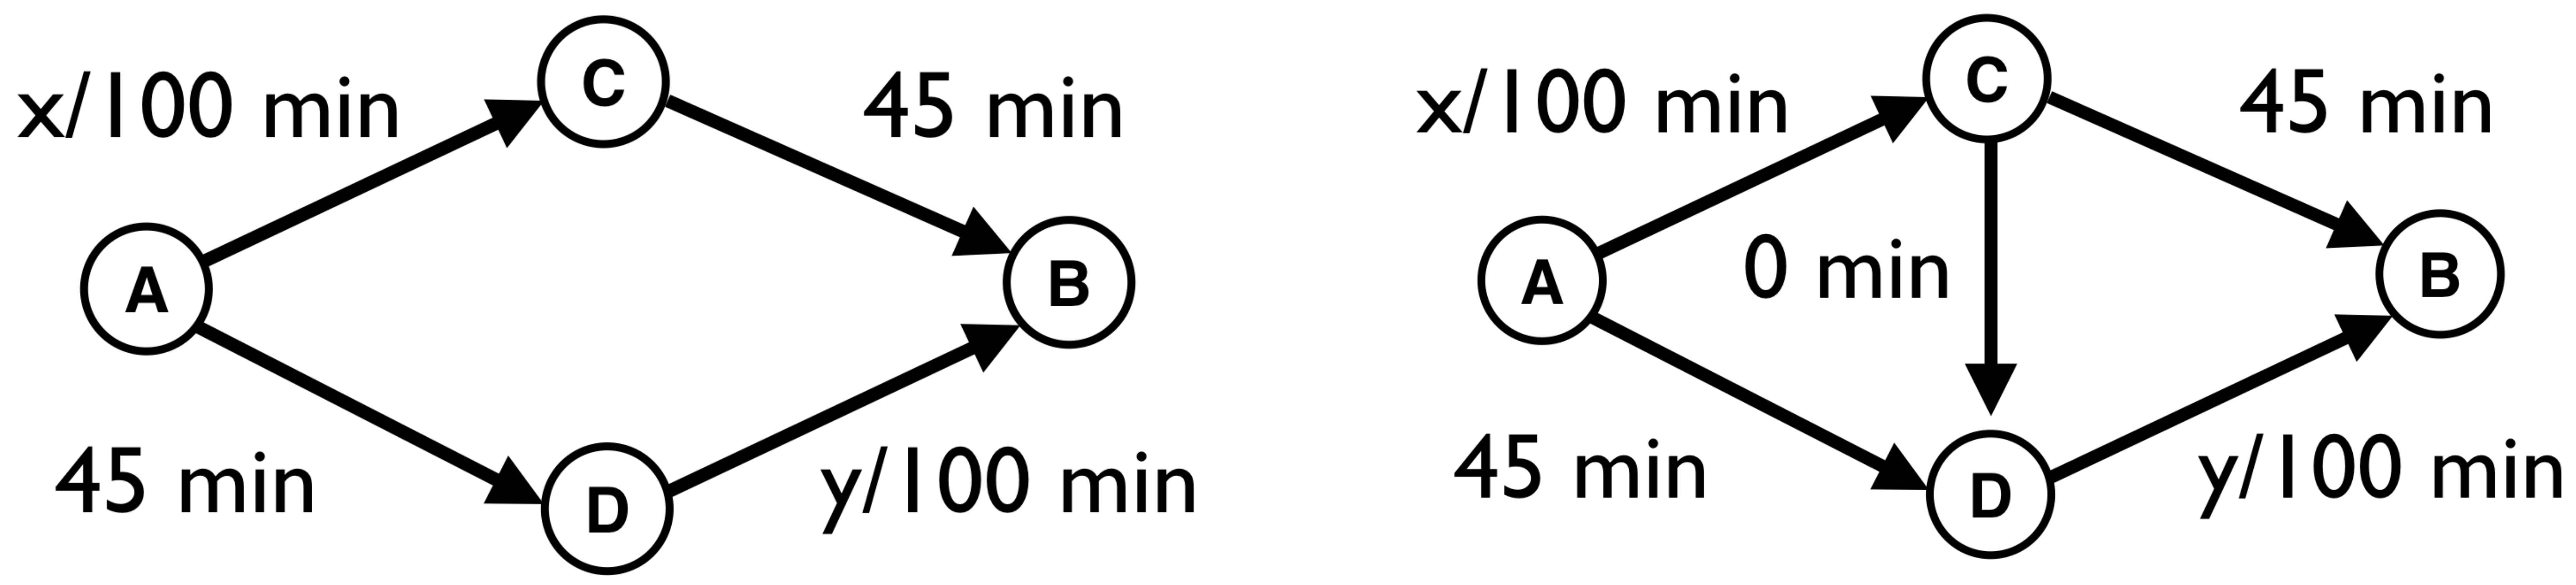

Prisoner's Dilemma:

		Suspect 2	
		NC	C
Suspect 1	NC	$-1, -1$	$-10, 0$
	C	$0, -10$	$-4, -4$

How bad can it get?

Routing:

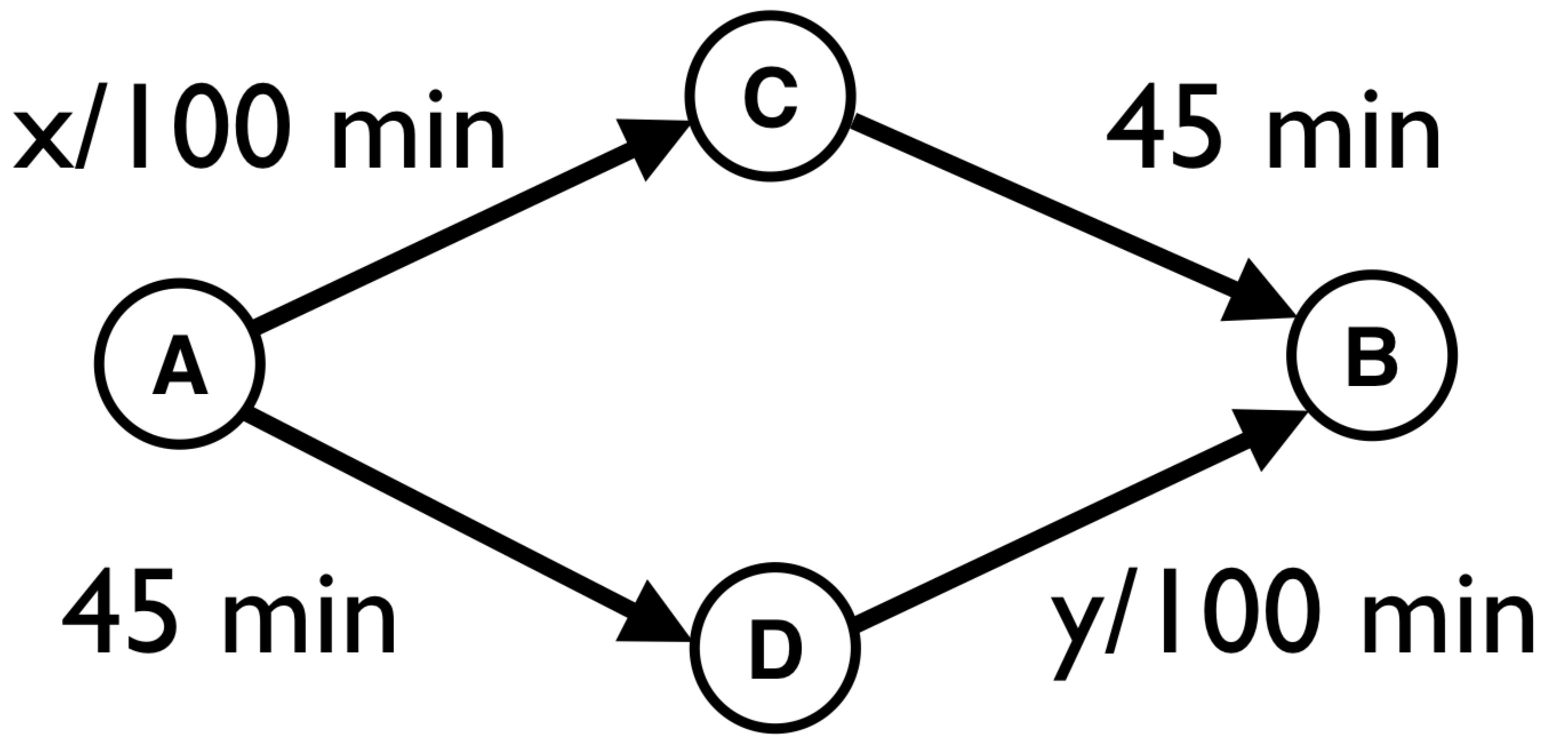

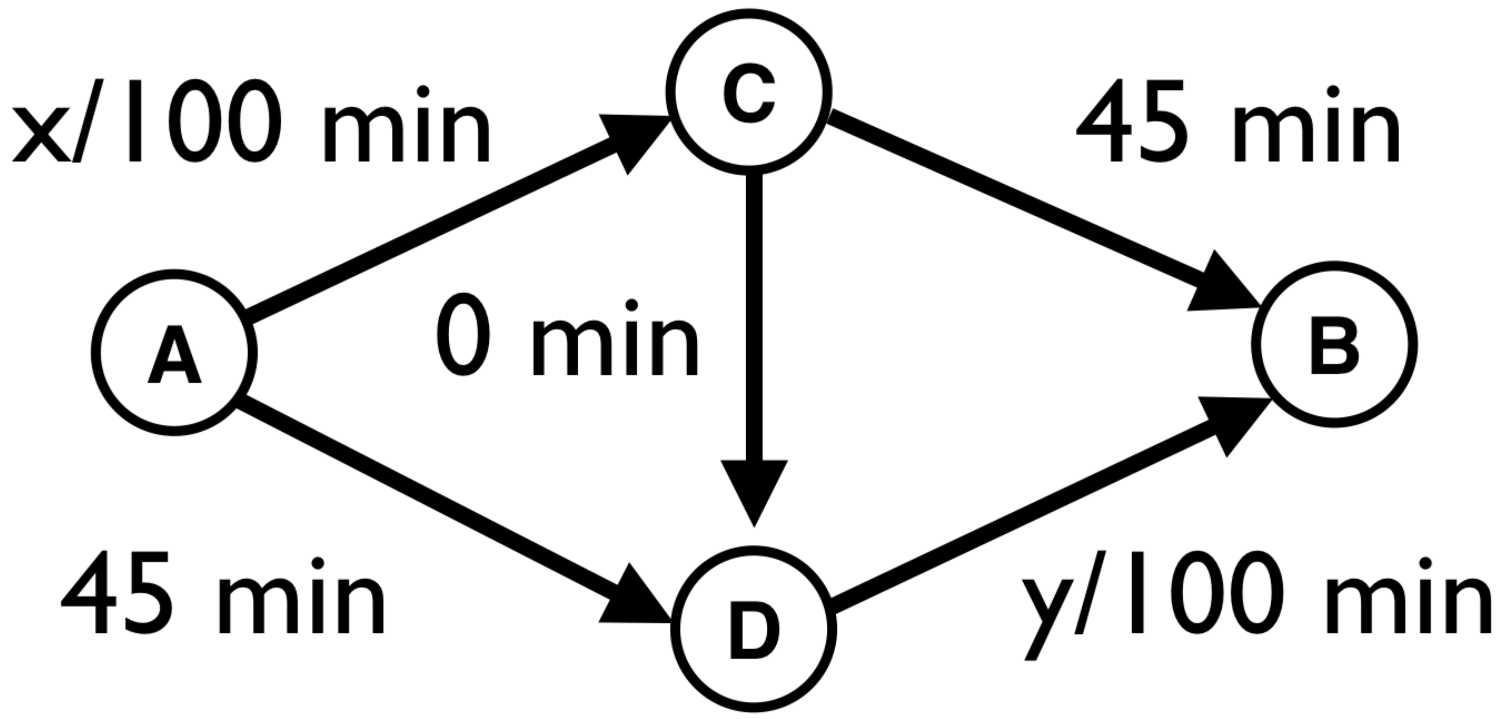

Ratio between socially optimal and selfish routing (called the "Price of Anarchy")?

This example: 80/65 = 1.23x worse

Worst case: How bad can it get?

For selfish routing, "Price of Anarchy" = 4/3

Diffusion of Decisions

Social Decisions

Lots of decisions you make depend on what your friends are doing

Where to go?

What game to play?

What software to use?

What OS to use?

iPhone vs. Android

BluRay vs. HD DVD

Blu-ray Disc

HD DVD

Electric Car vs. Diesel Truck

How to Reason About Social Decisions?

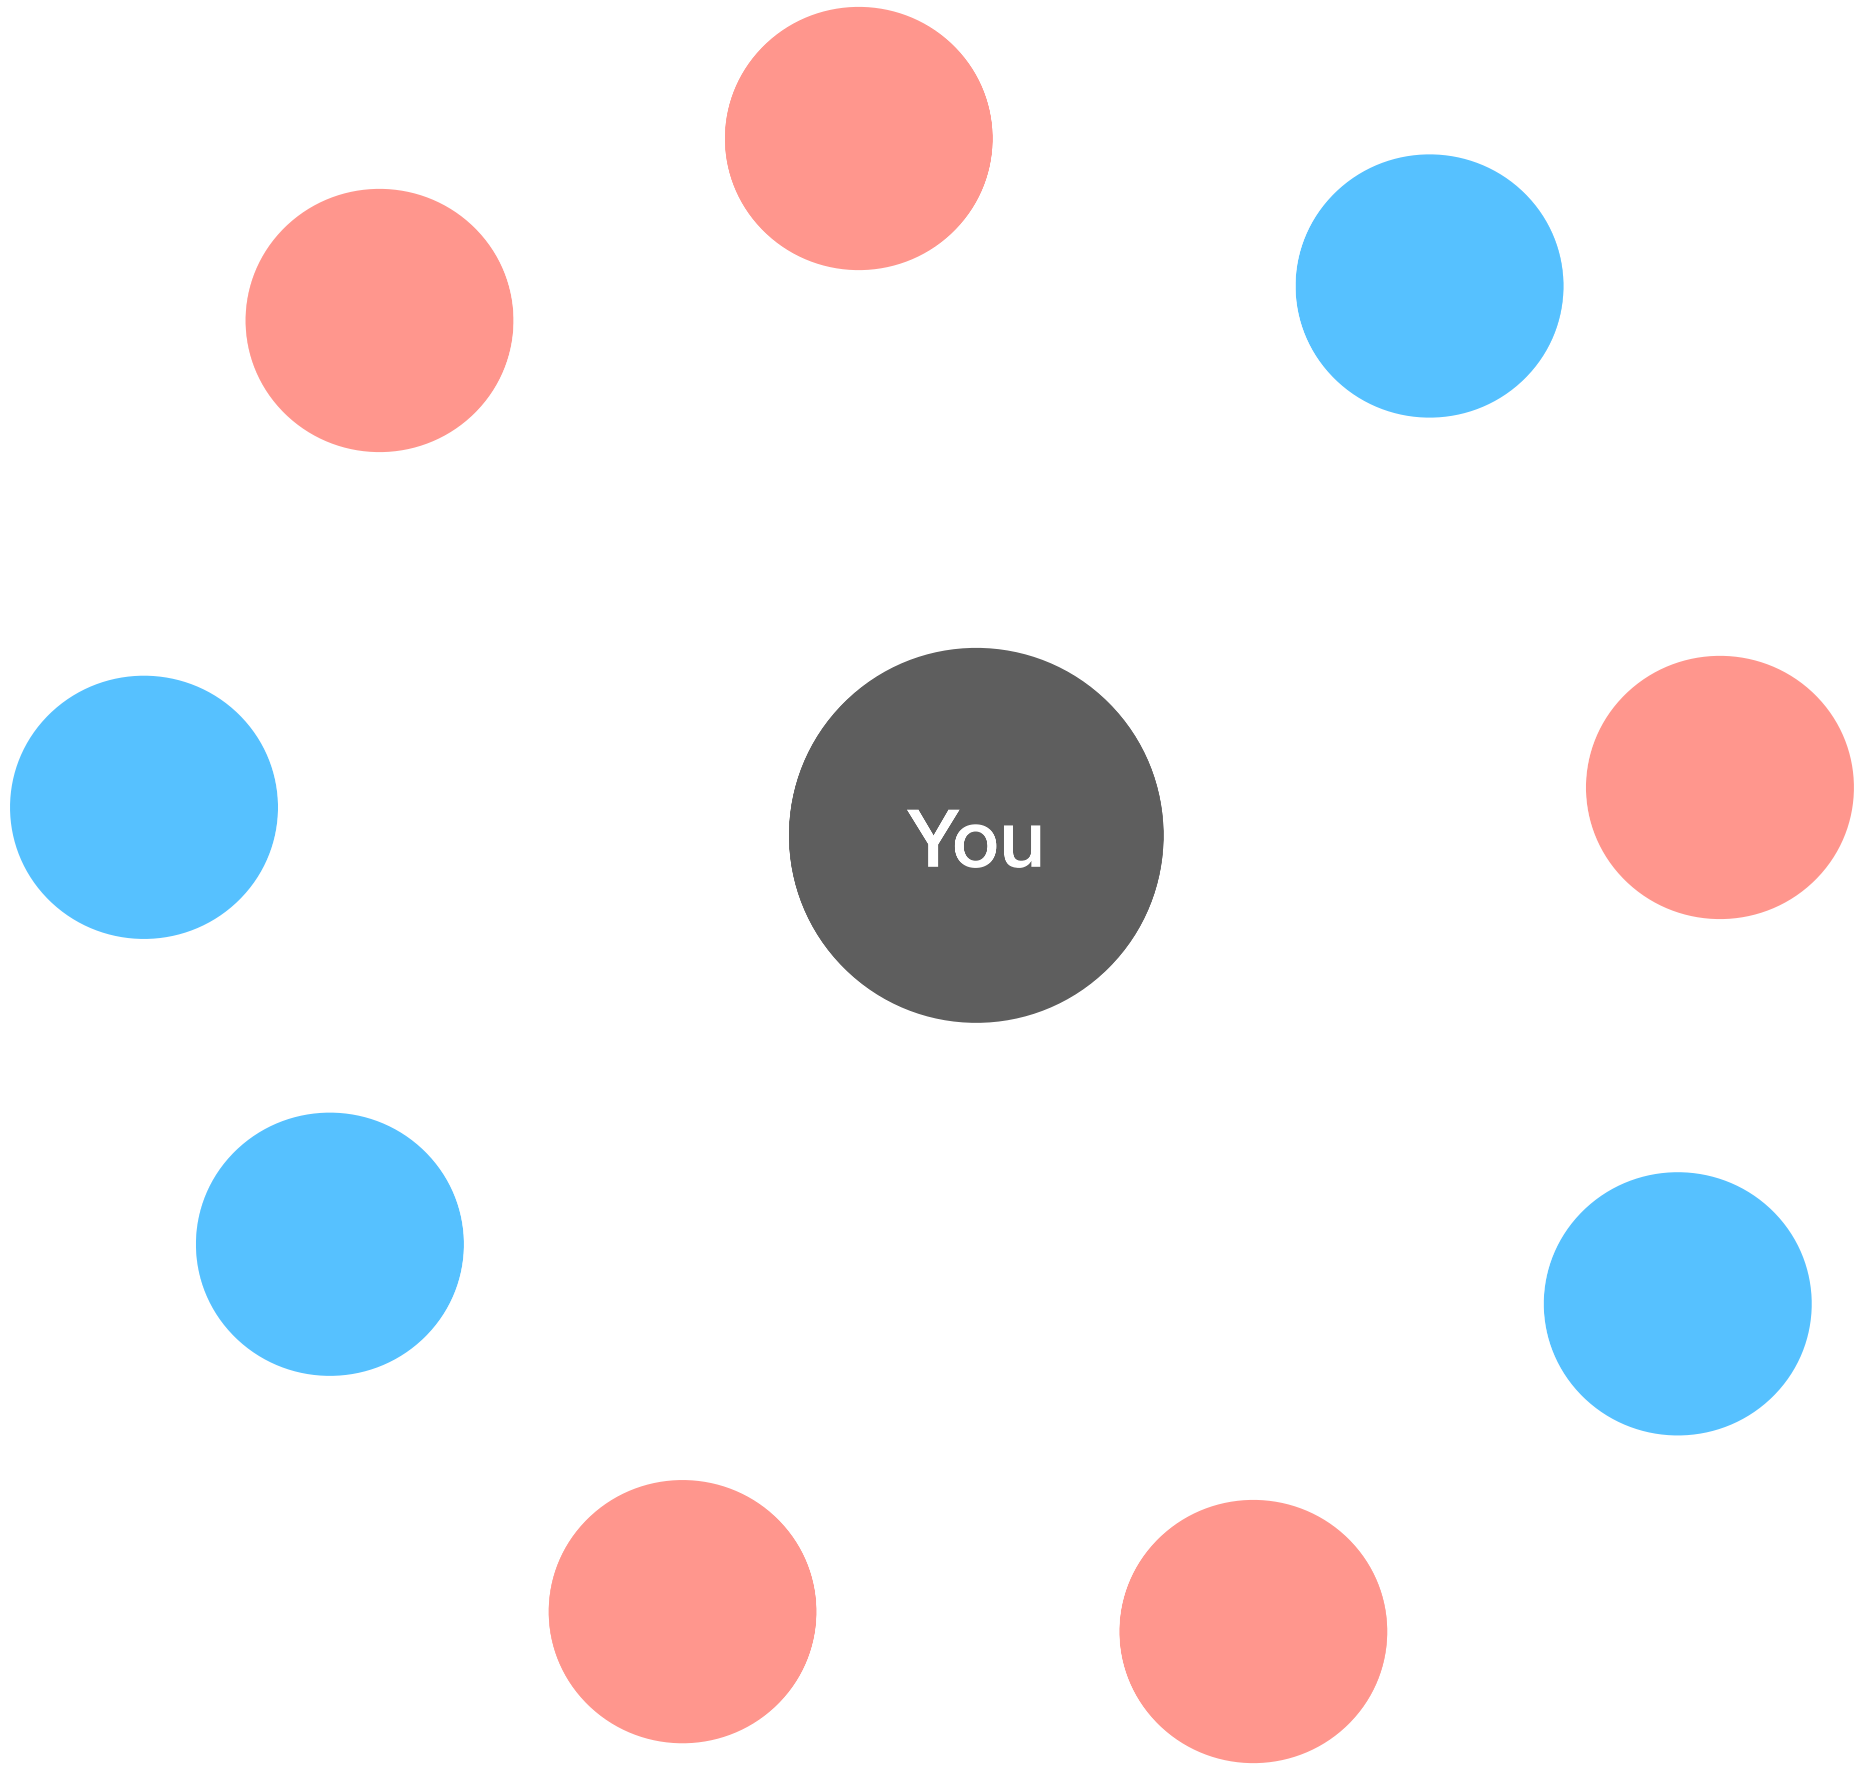

Given that your friends have all chosen one way or another, what should you choose?

How to Reason About Social Decisions?

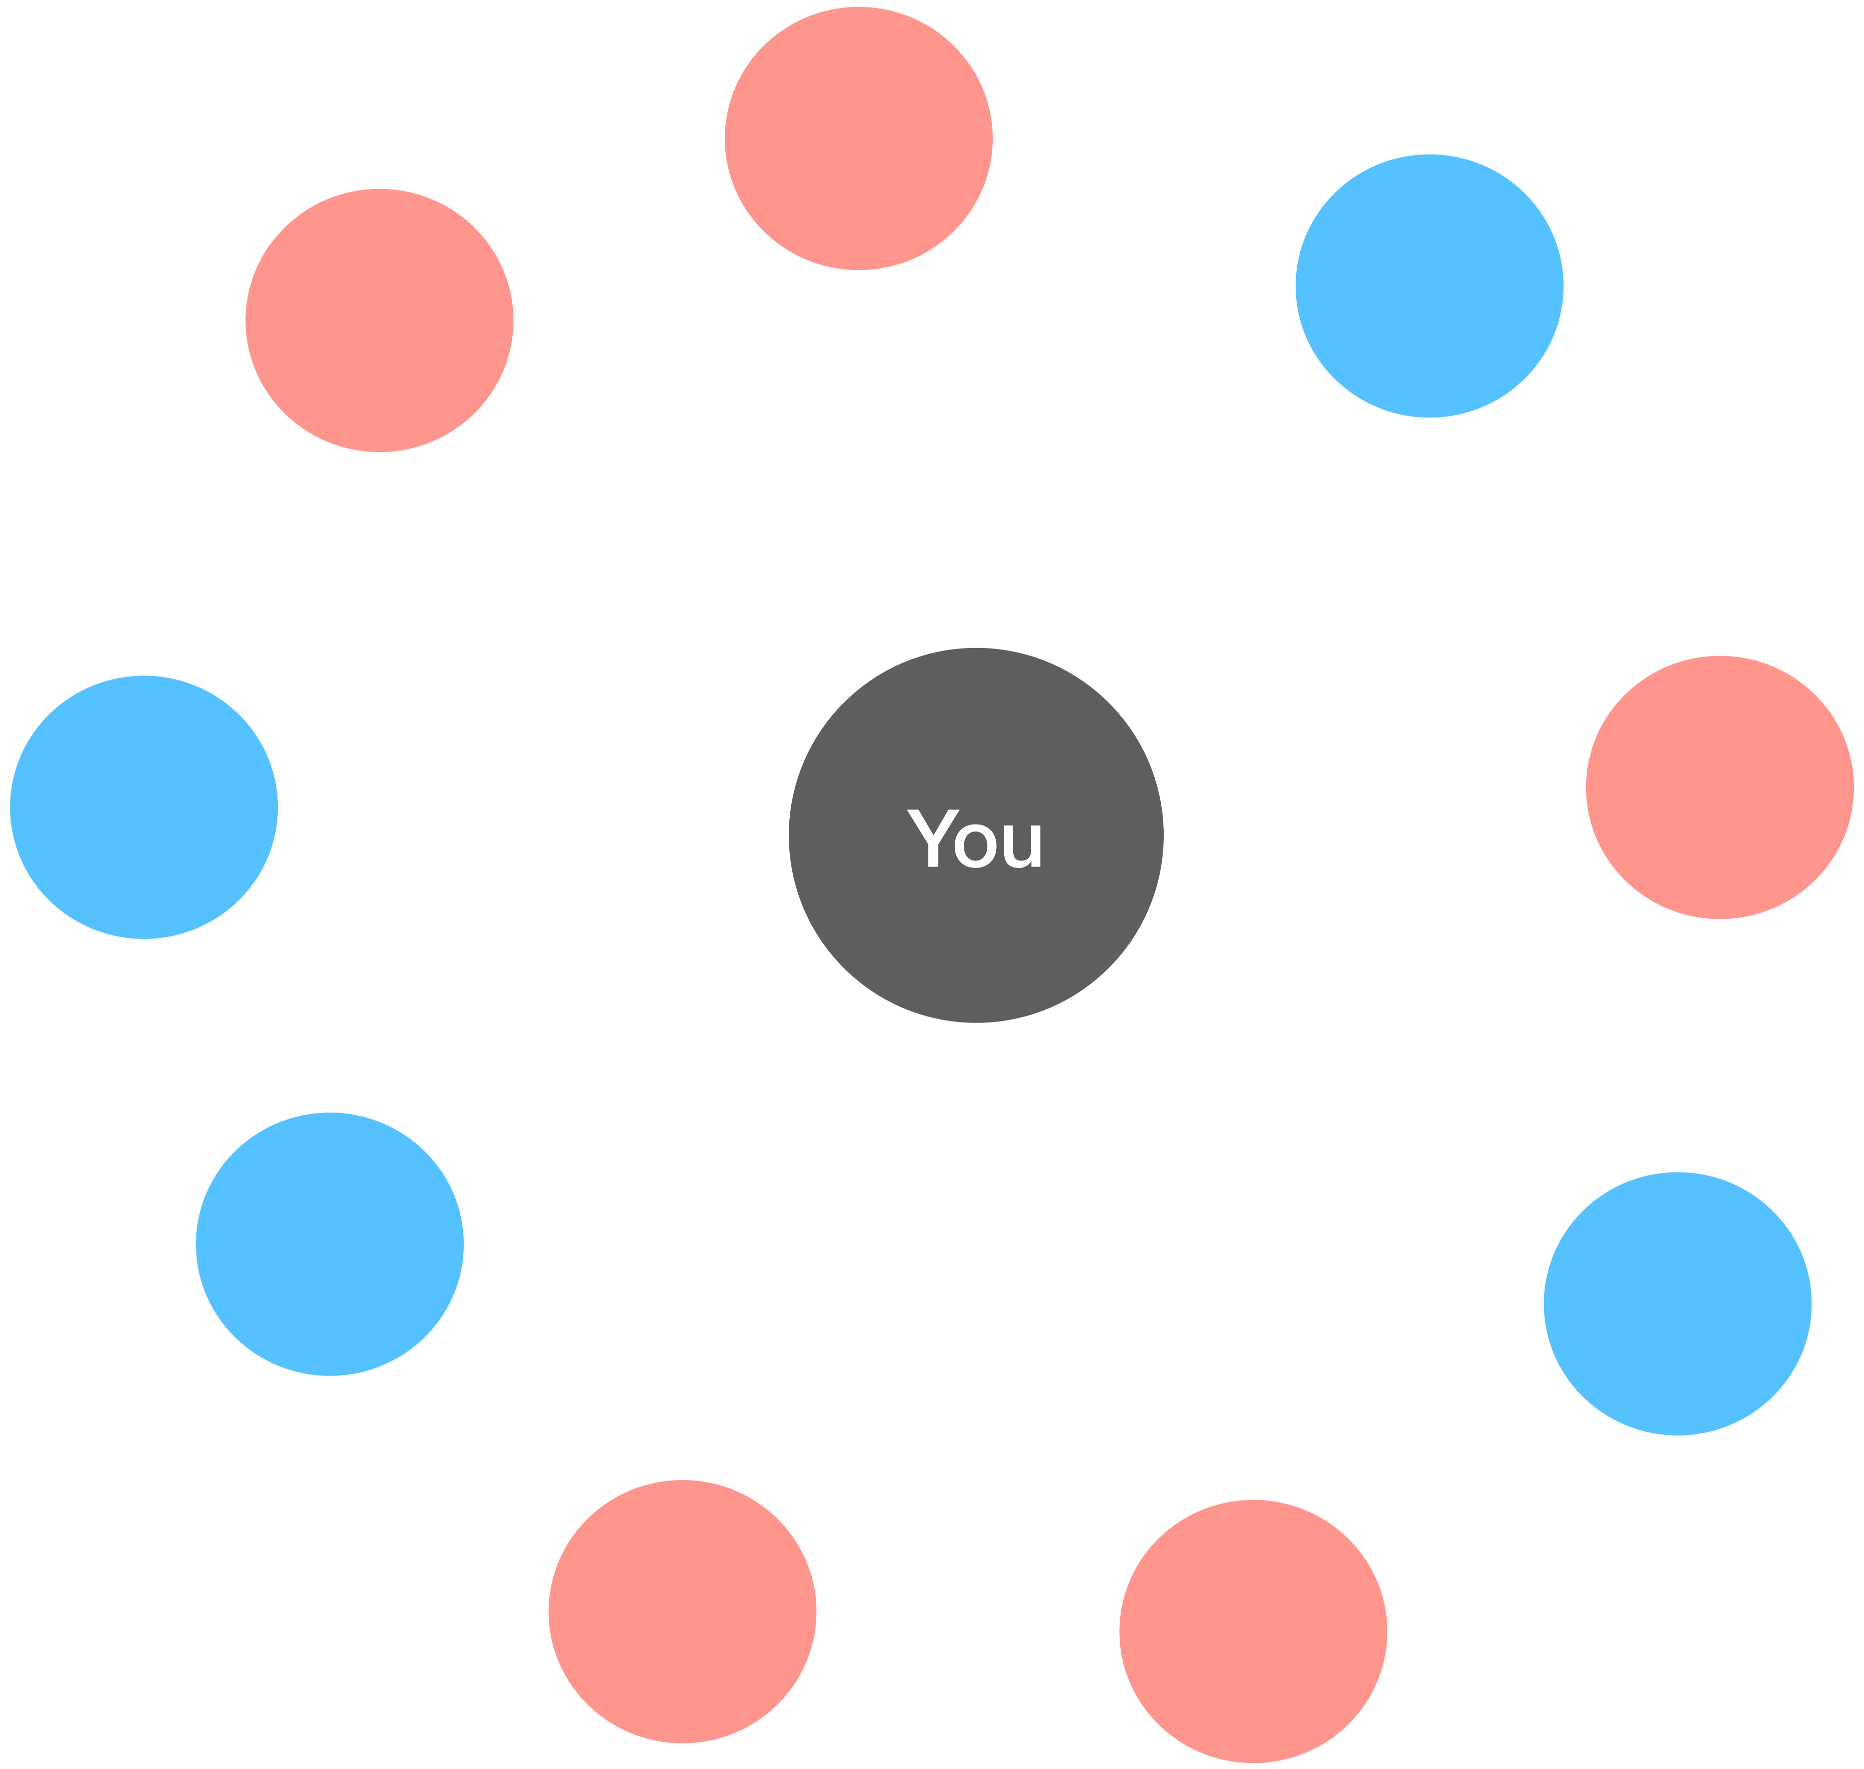

“Network Effects”

Game Theoretic Model of Cascades

Social Networks + Game Theory can help us think about this question!

Model every friendship edge as a 2 player coordination game

2 players – each chooses technology A or B

Each person can only adopt **one** "behavior", **A** or **B**

You gain more payoff if your friend has adopted the **same** behavior as you

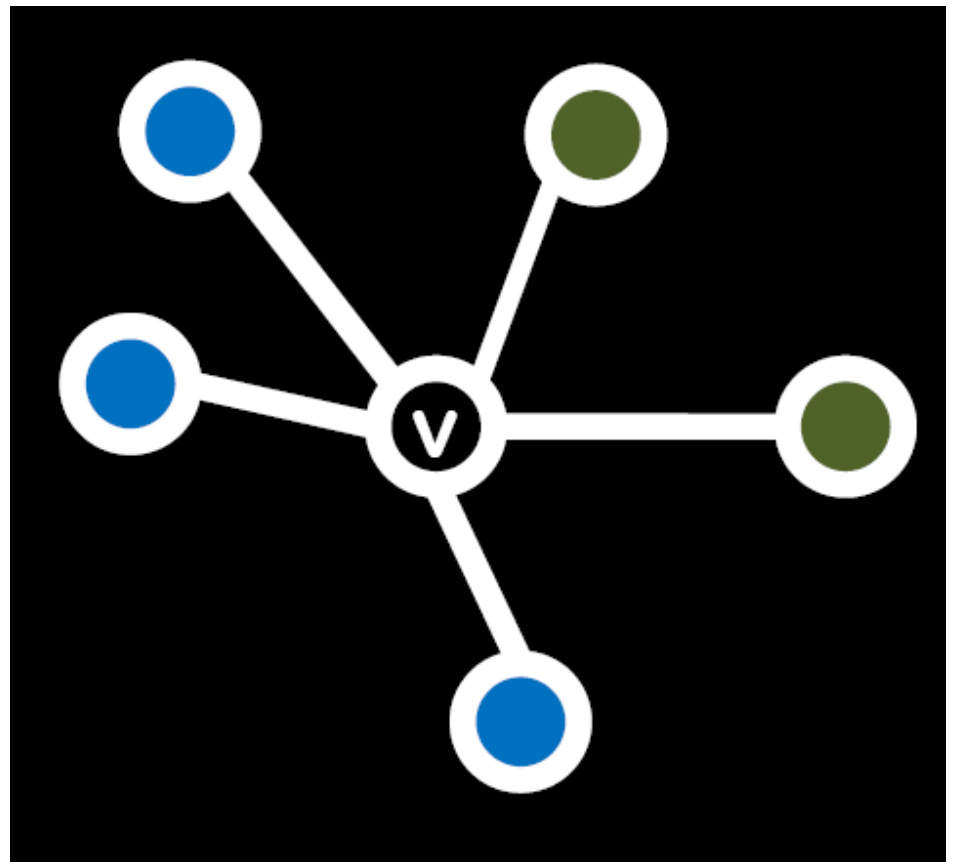

Local view of the network of node **v**

w

v		A	B
	A	a, a	$0, 0$
	B	$0, 0$	b, b

The Model for Two Nodes

Payoff matrix:

If both v and w adopt behaviour A, they each get payoff $a > 0$

If v and w adopt behaviour B, they each get payoff $b > 0$

If v and w adopt the opposite behaviours, they each get 0

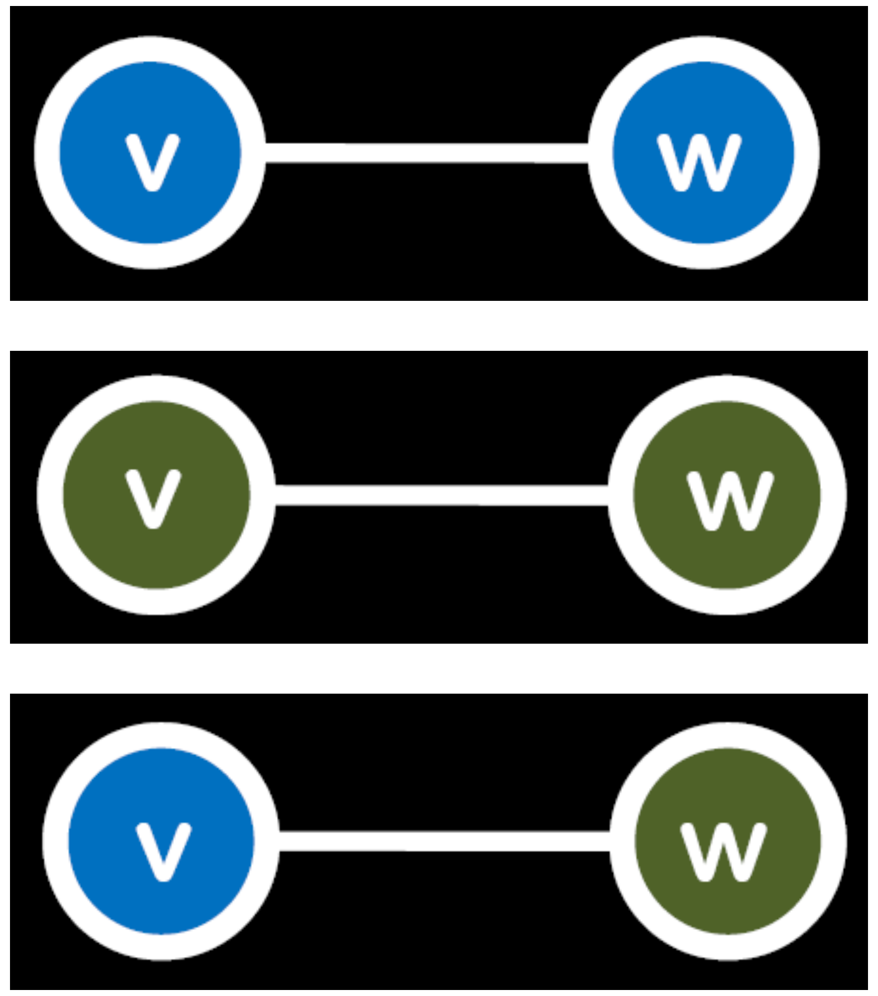

In some large network:

Each node v is playing a copy of the coordination game with each of its neighbours

Payoff: sum of node payoffs per game

		w: A	w: B
v	A	a, a	$0, 0$
	B	$0, 0$	b, b

Calculation of Node *v*

Let **v** have ***d*** neighbours — some adopt **A** and some adopt **B**

Say fraction ***p*** of ***v***'s neighbours adopt ***A*** and **1-p** adopt **B**

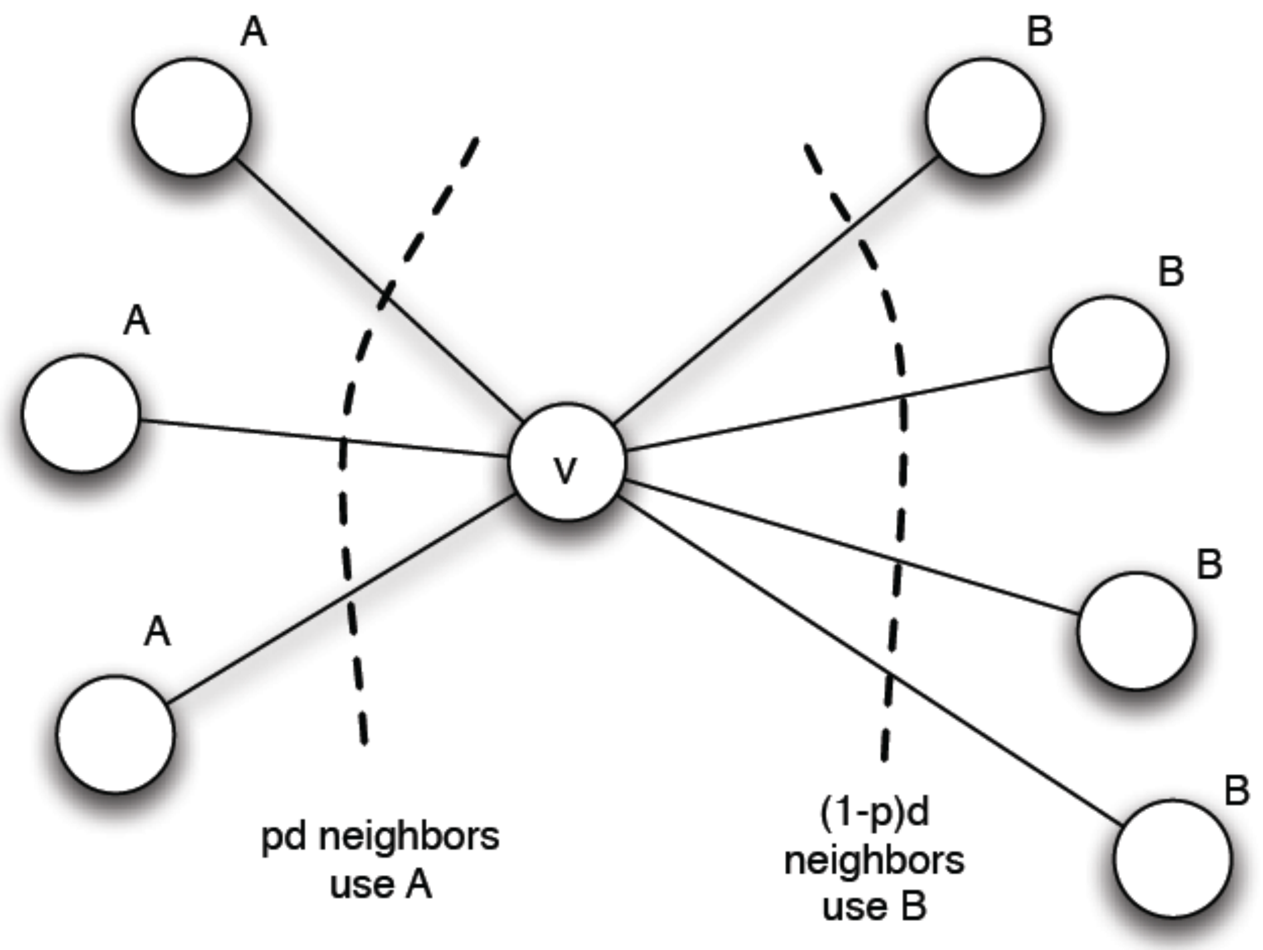

$Payoff_v = a \cdot p \cdot d$ if v chooses A

$= b \cdot (1-p) \cdot d$ if v chooses B

Thus: *v* chooses A if:

$a \cdot p \cdot d > b \cdot (1-p) \cdot d$

Threshold:

v chooses ***A*** if $p > \frac{b}{a+b} = q$

p… frac. v's neighbours choosing A
q… payoff threshold

Example Scenario

Scenario:
Graph where everyone starts with *B*
Small set *S* of early adopters of *A*

Hard-wire *S* – they keep using *A* no matter what payoffs tell them to do

Assume payoffs are set in such a way that nodes say:
If **more than 50%** of my friends take A
I'll also take A

(this means: $a = b-\varepsilon$ and $q>1/2$)

Example Scenario

$S = \{u, v\}$

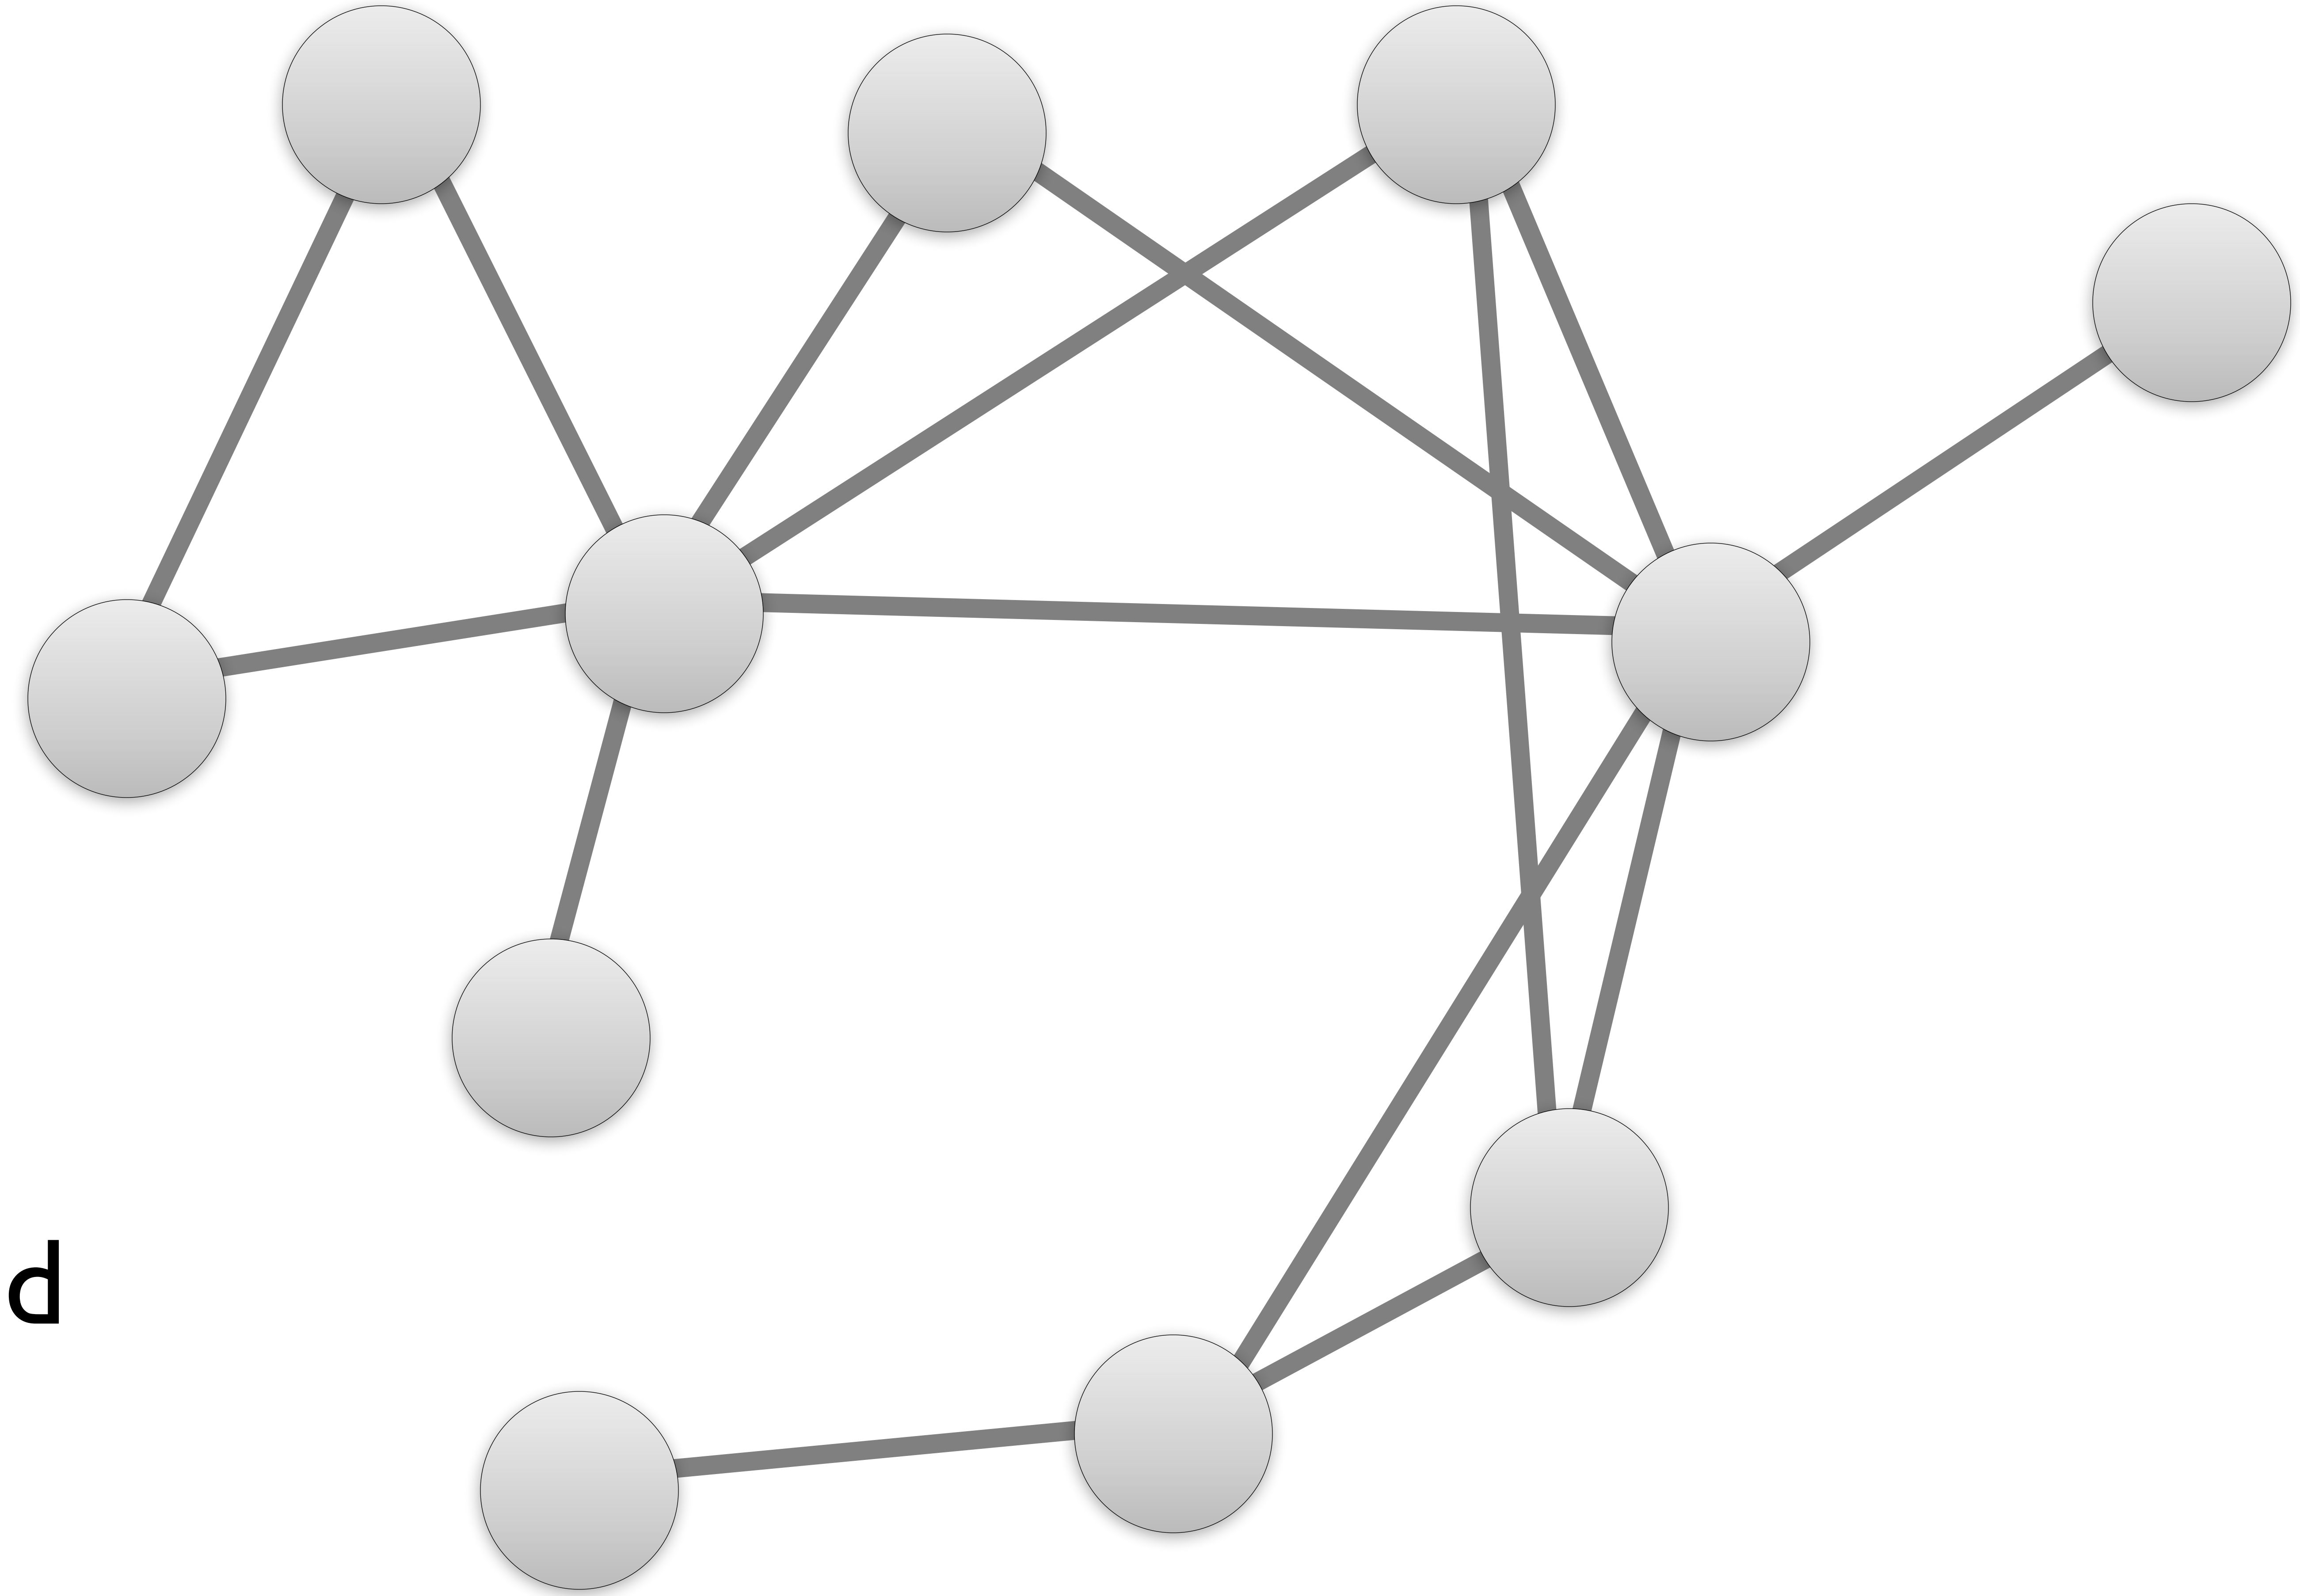

If more than q=50% of my friends are red I'll be red

Example Scenario

$$S = \{u, v\}$$

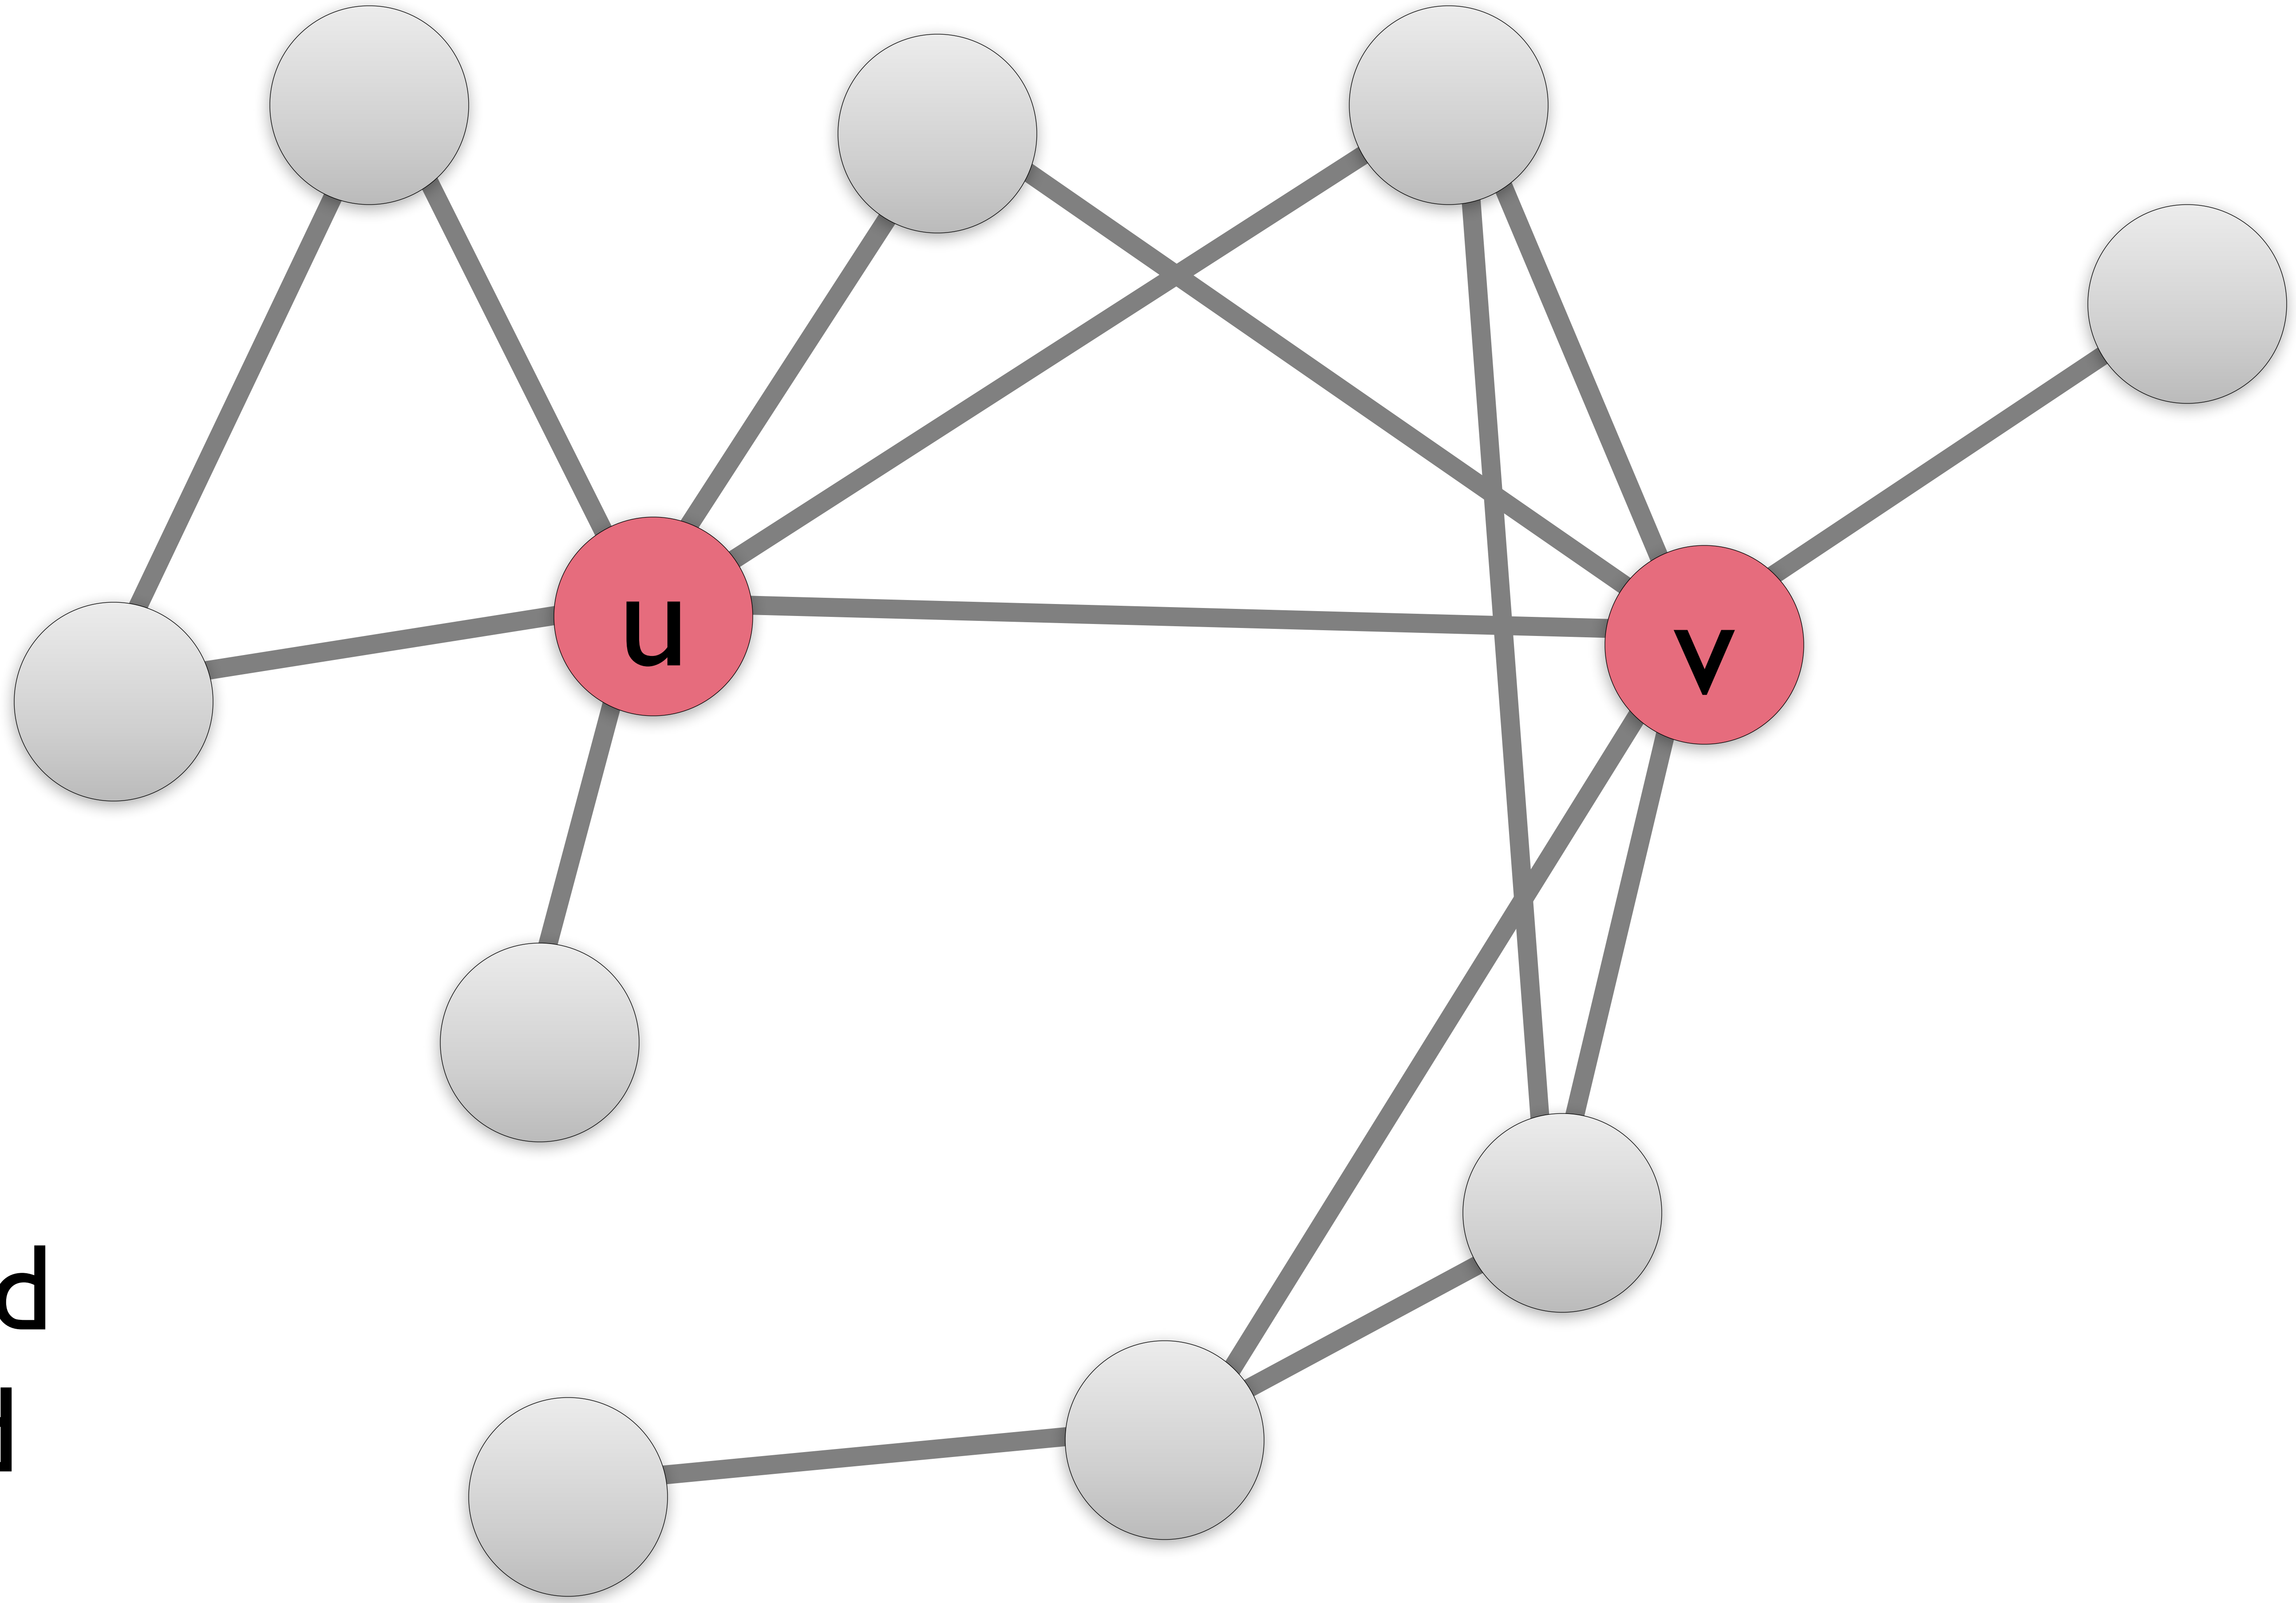

If more than q=50% of my friends are red I'll also be red

Example Scenario

$S = \{u, v\}$

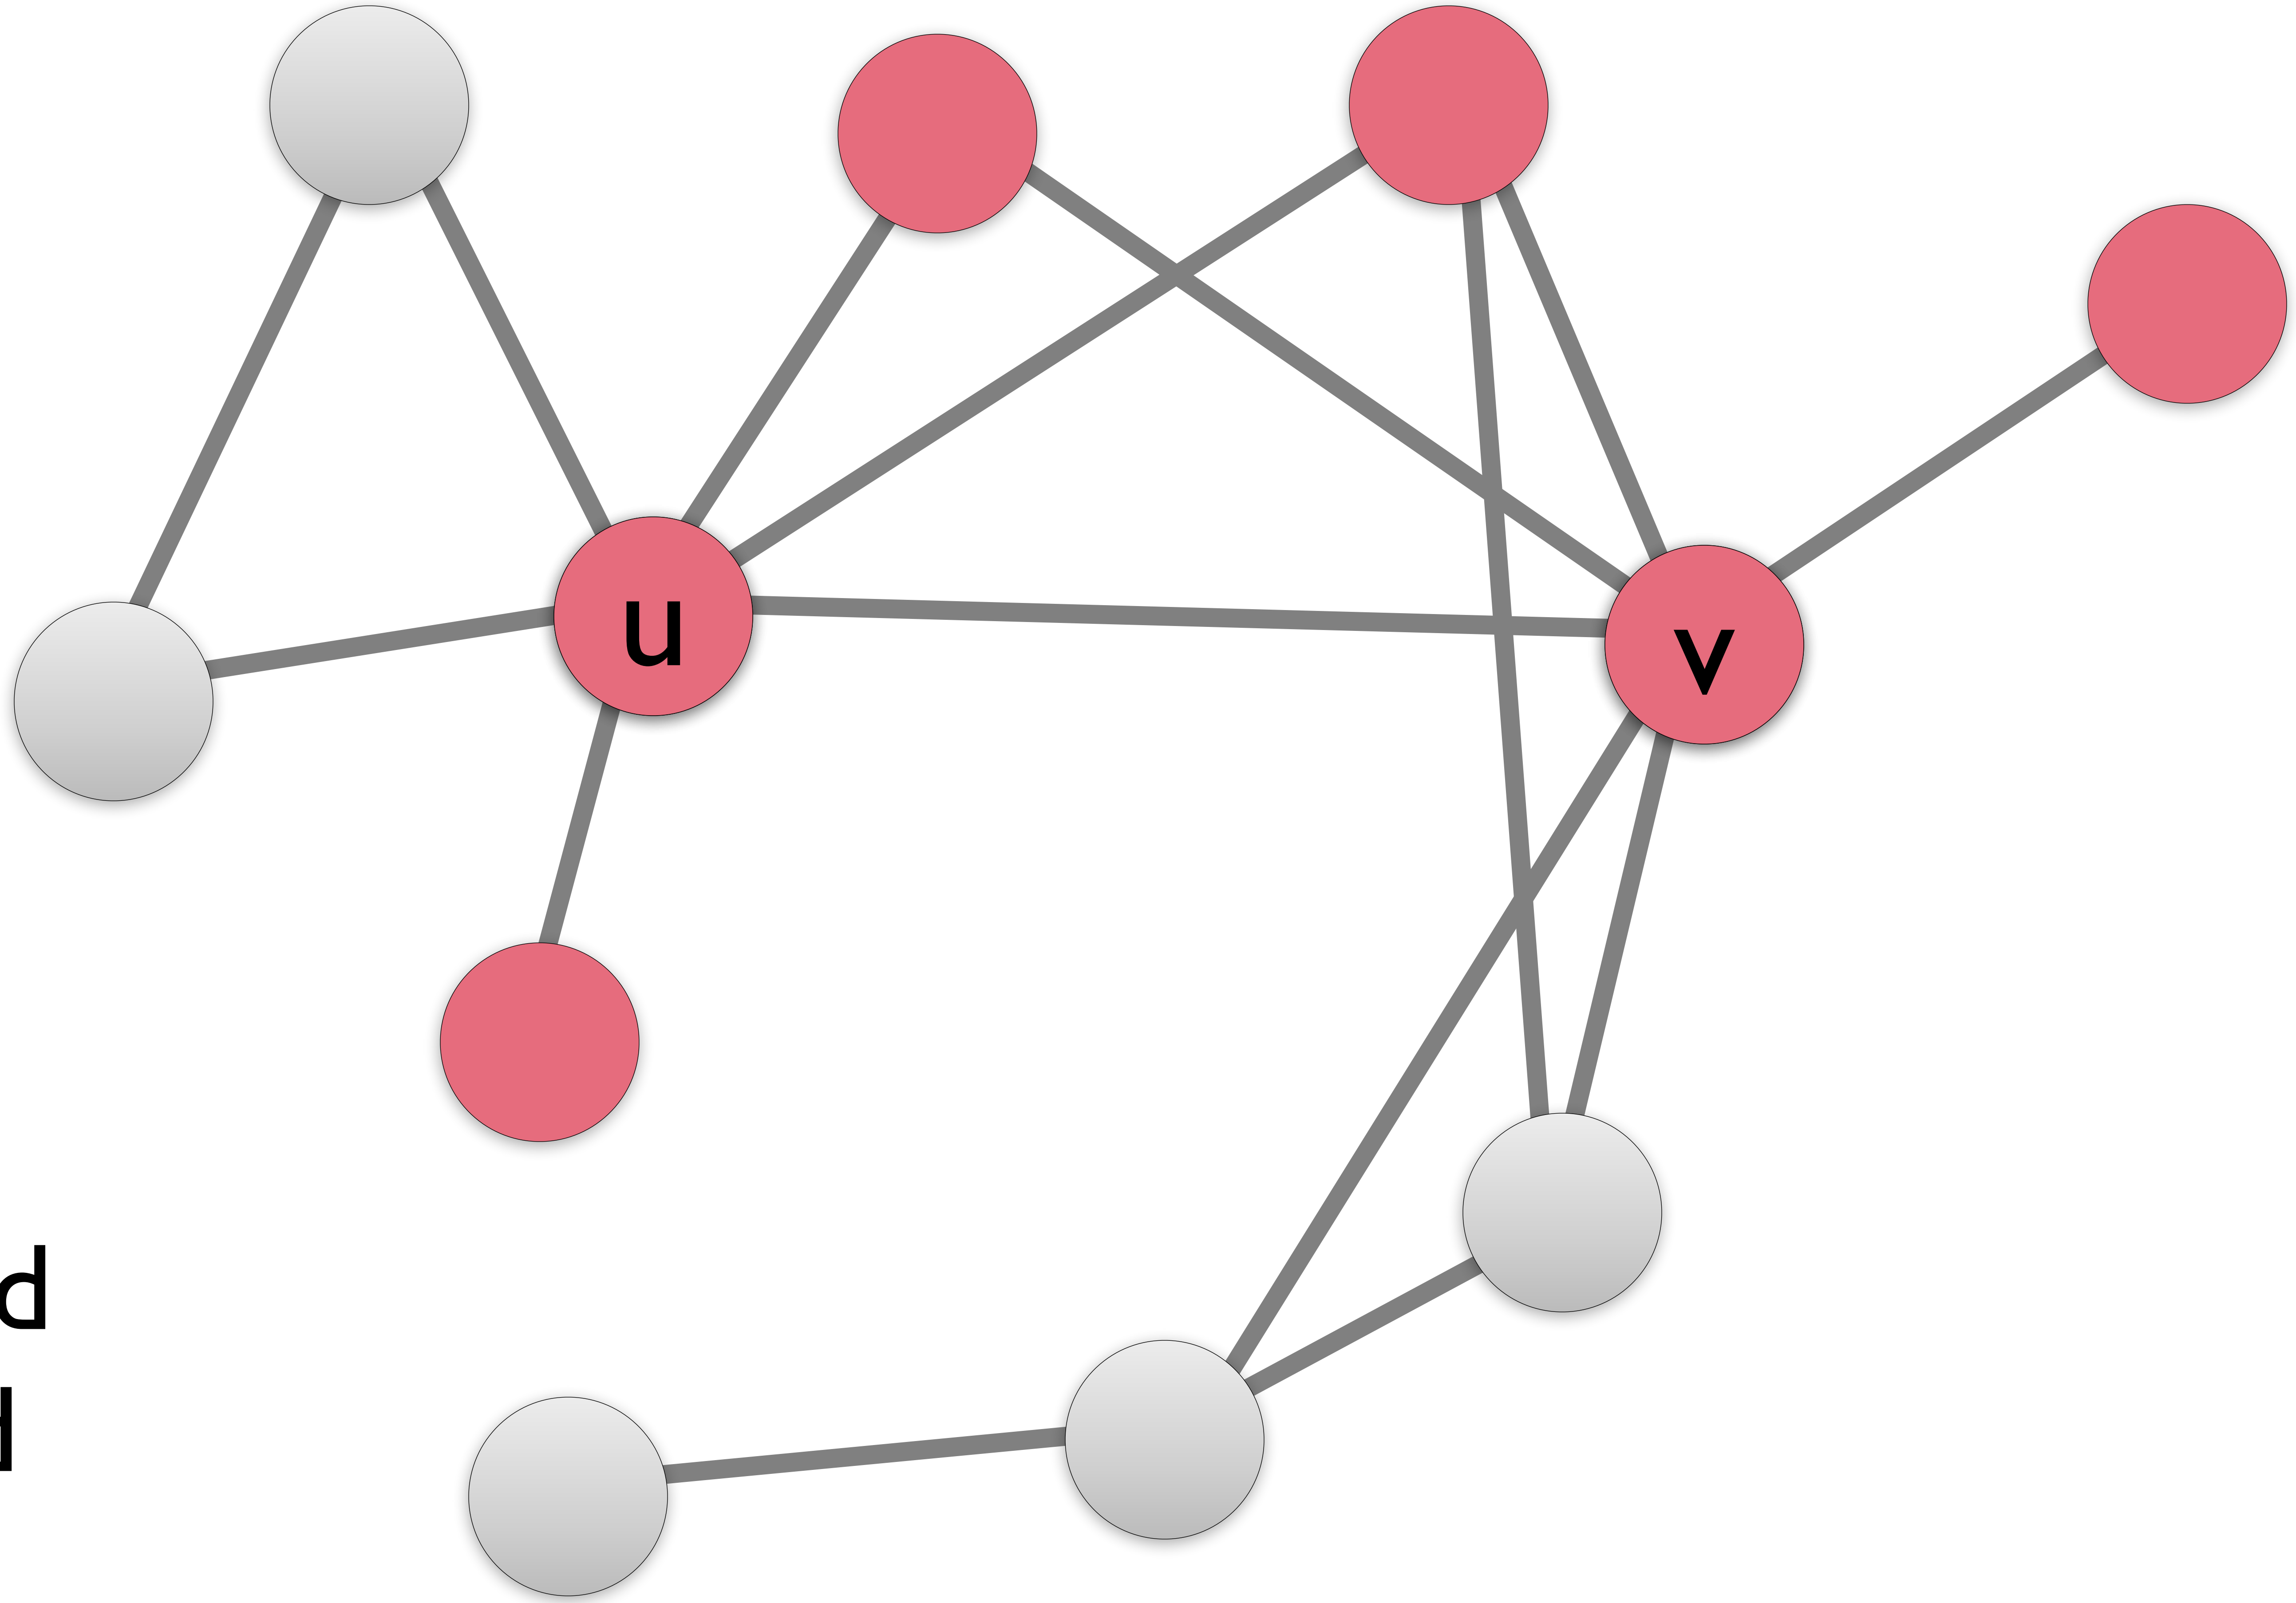

If more than q=50% of my friends are red I'll also be red

Example Scenario

$S = \{u, v\}$

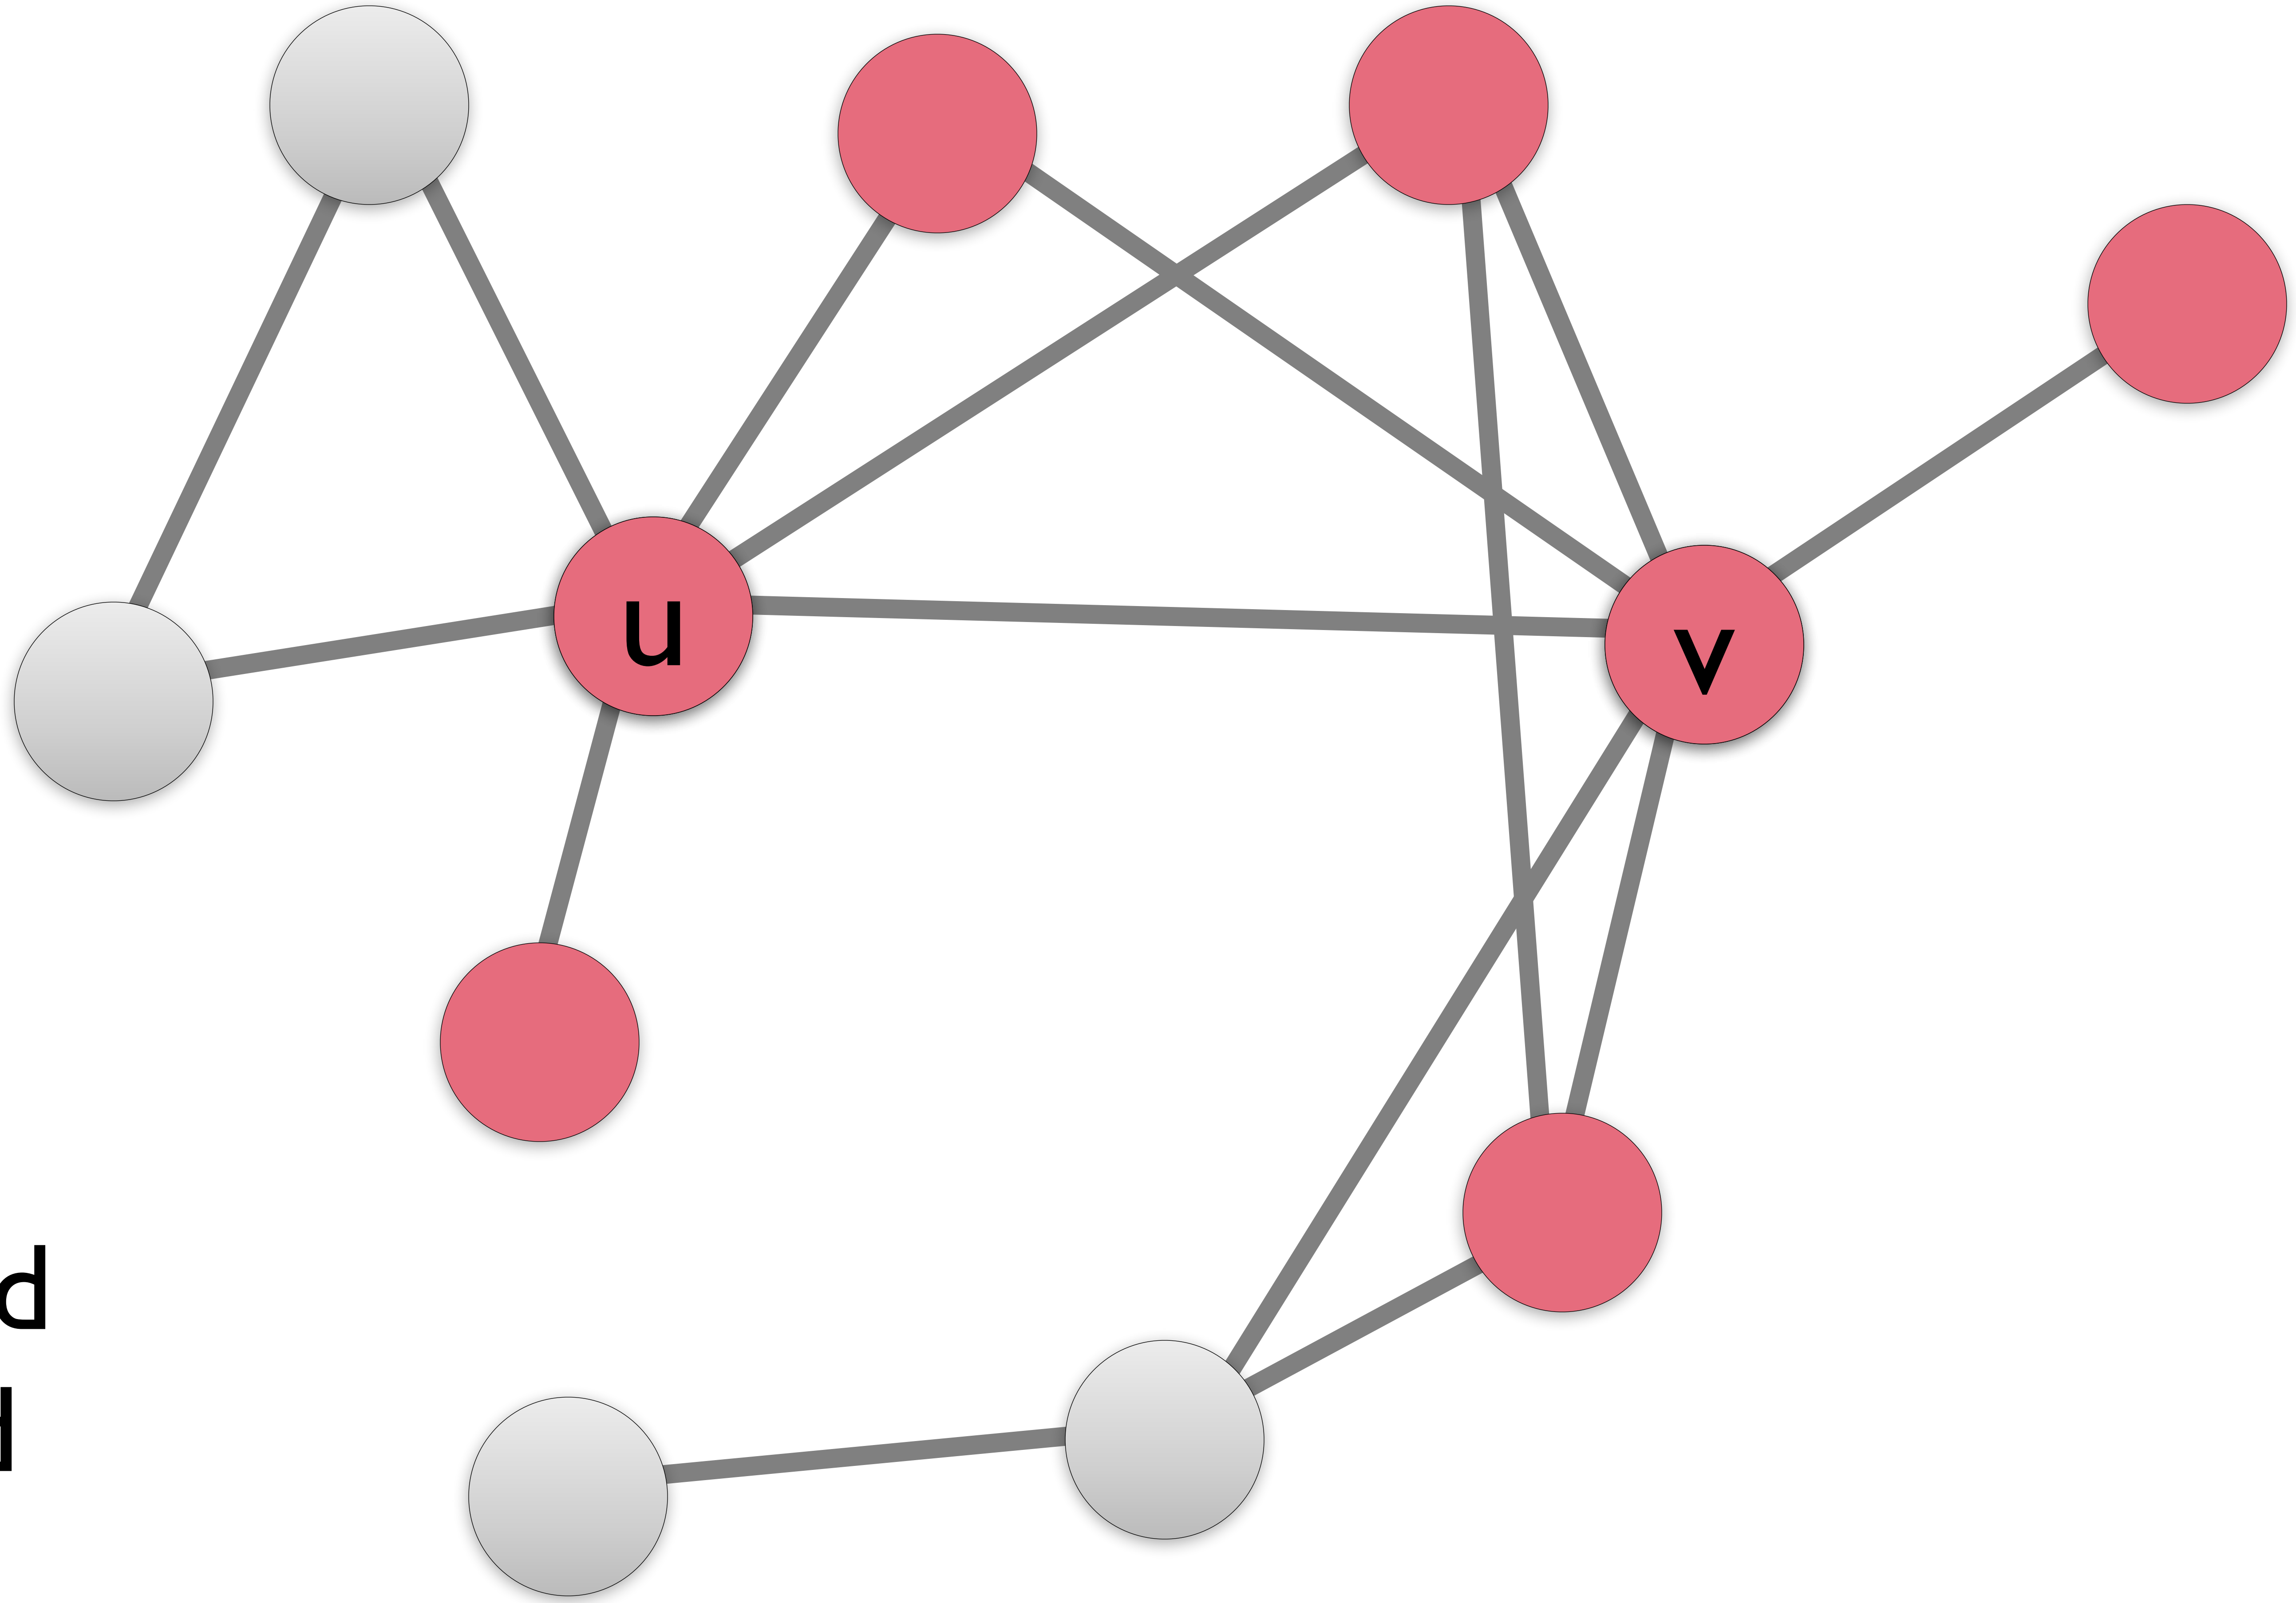

If more than q=50% of my friends are red I'll also be red

Example Scenario

$S = \{u, v\}$

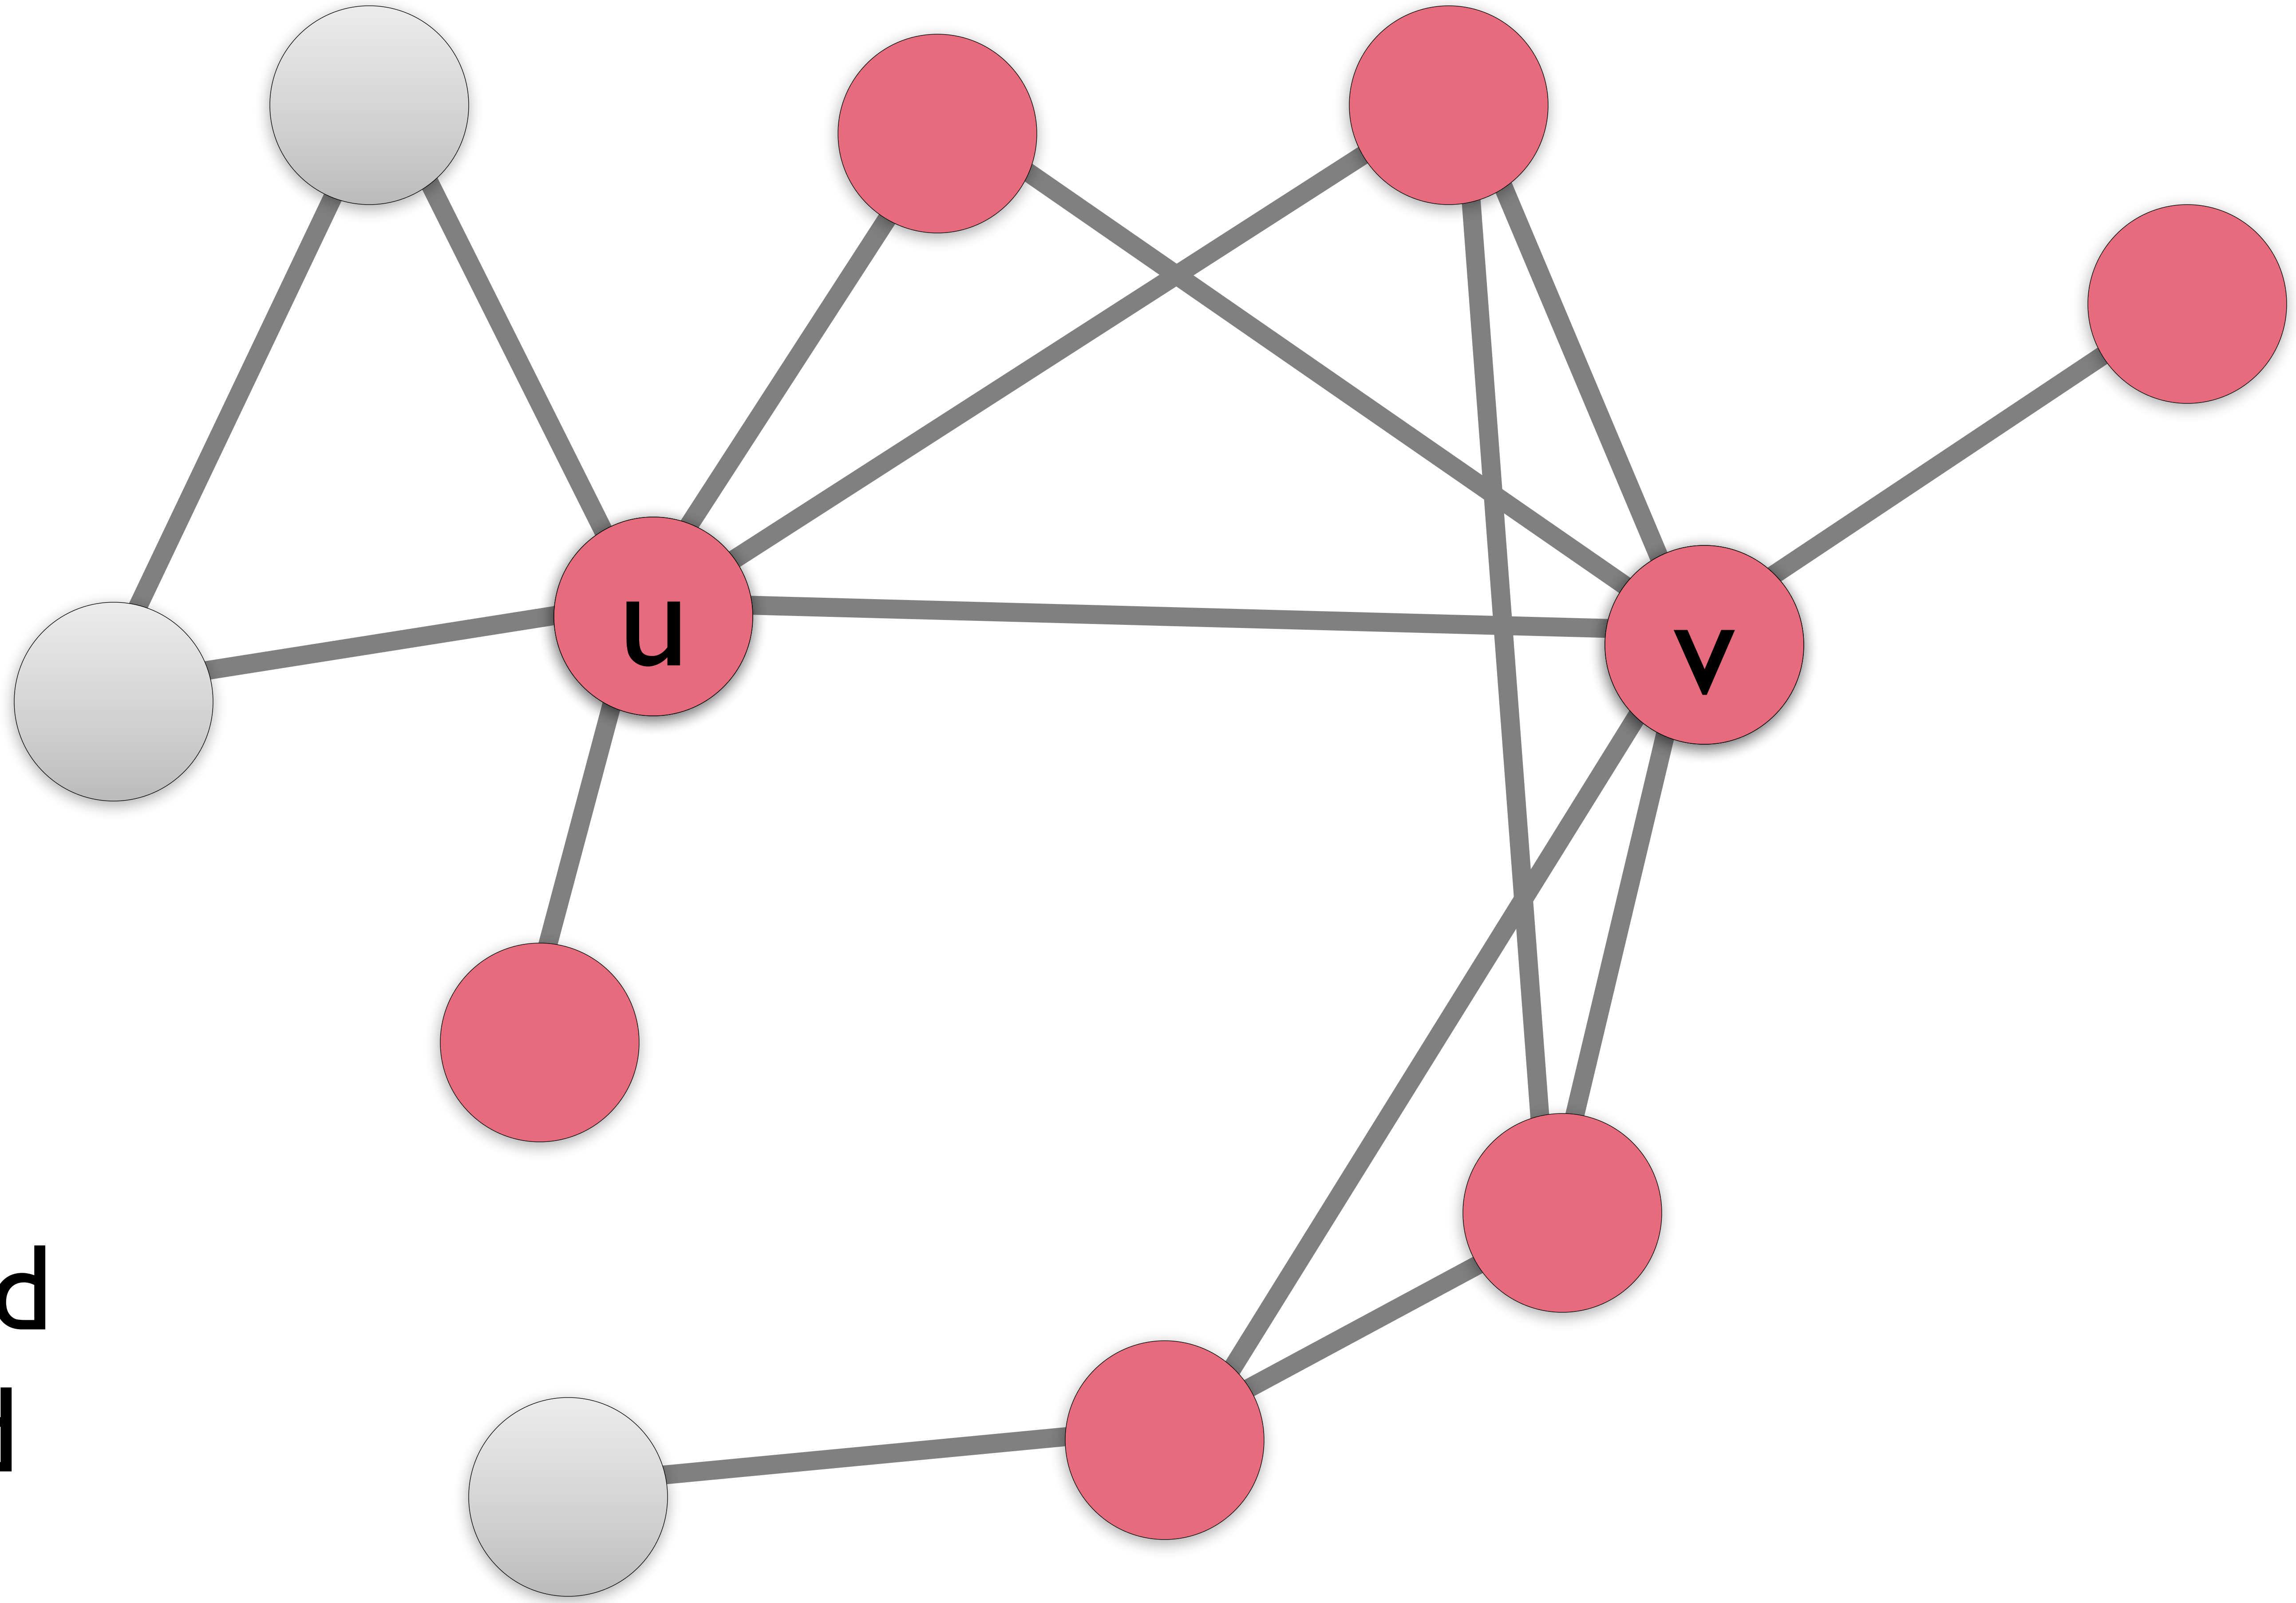

If more than q=50% of my friends are red I'll also be red

Another example with a=3 and b=2

$$p > \frac{b}{a+b} = q$$

q = 2/5

(new technology better,
so q<1/2)

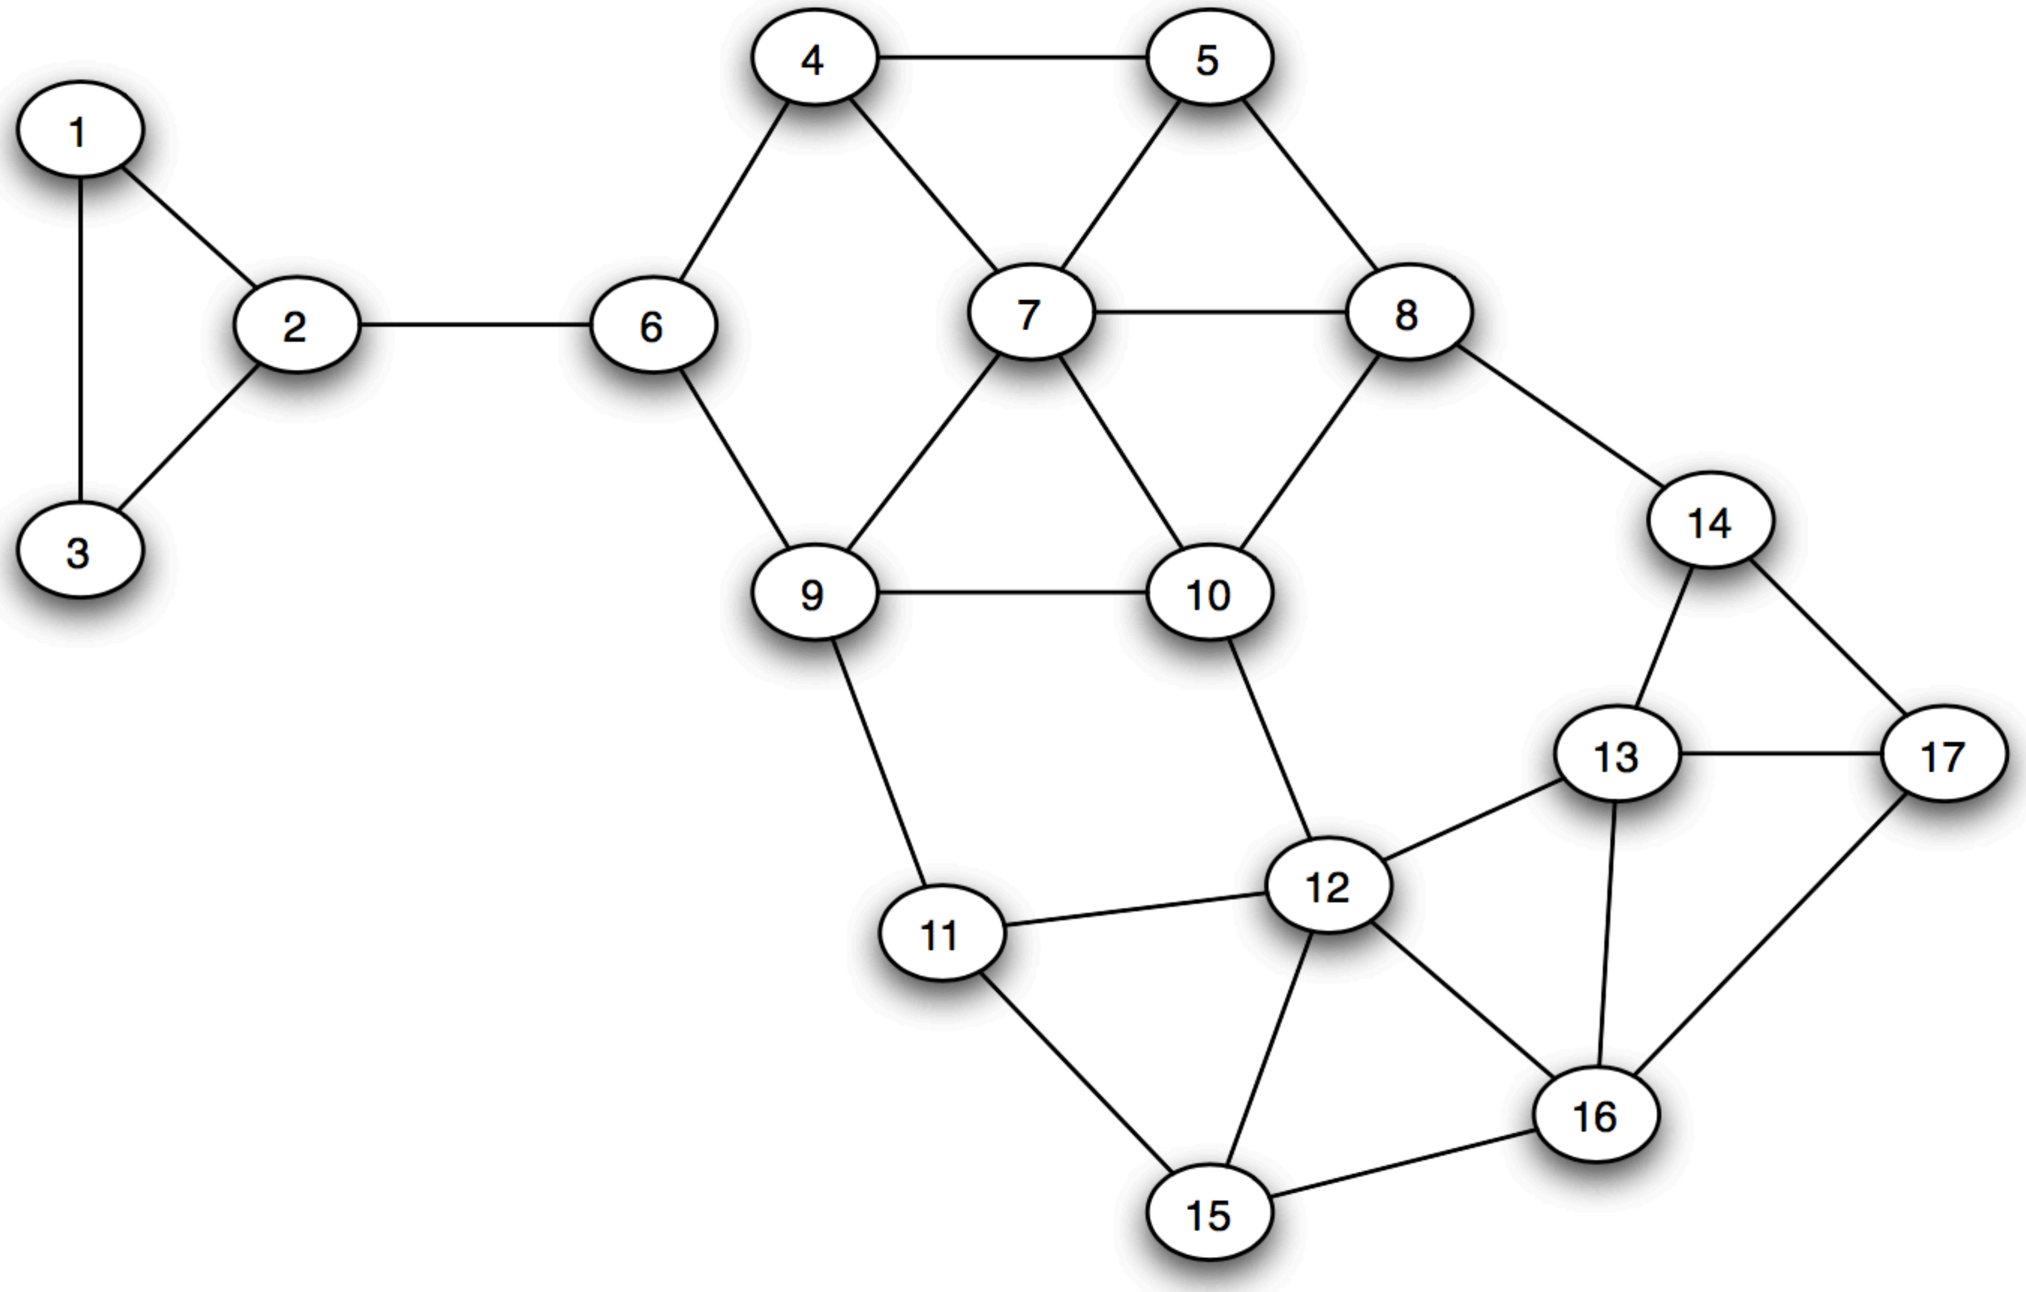

Another example with a=3 and b=2

$$p > \frac{b}{a+b} = q$$

q = 2/5

(new technology better, so q<1/2)

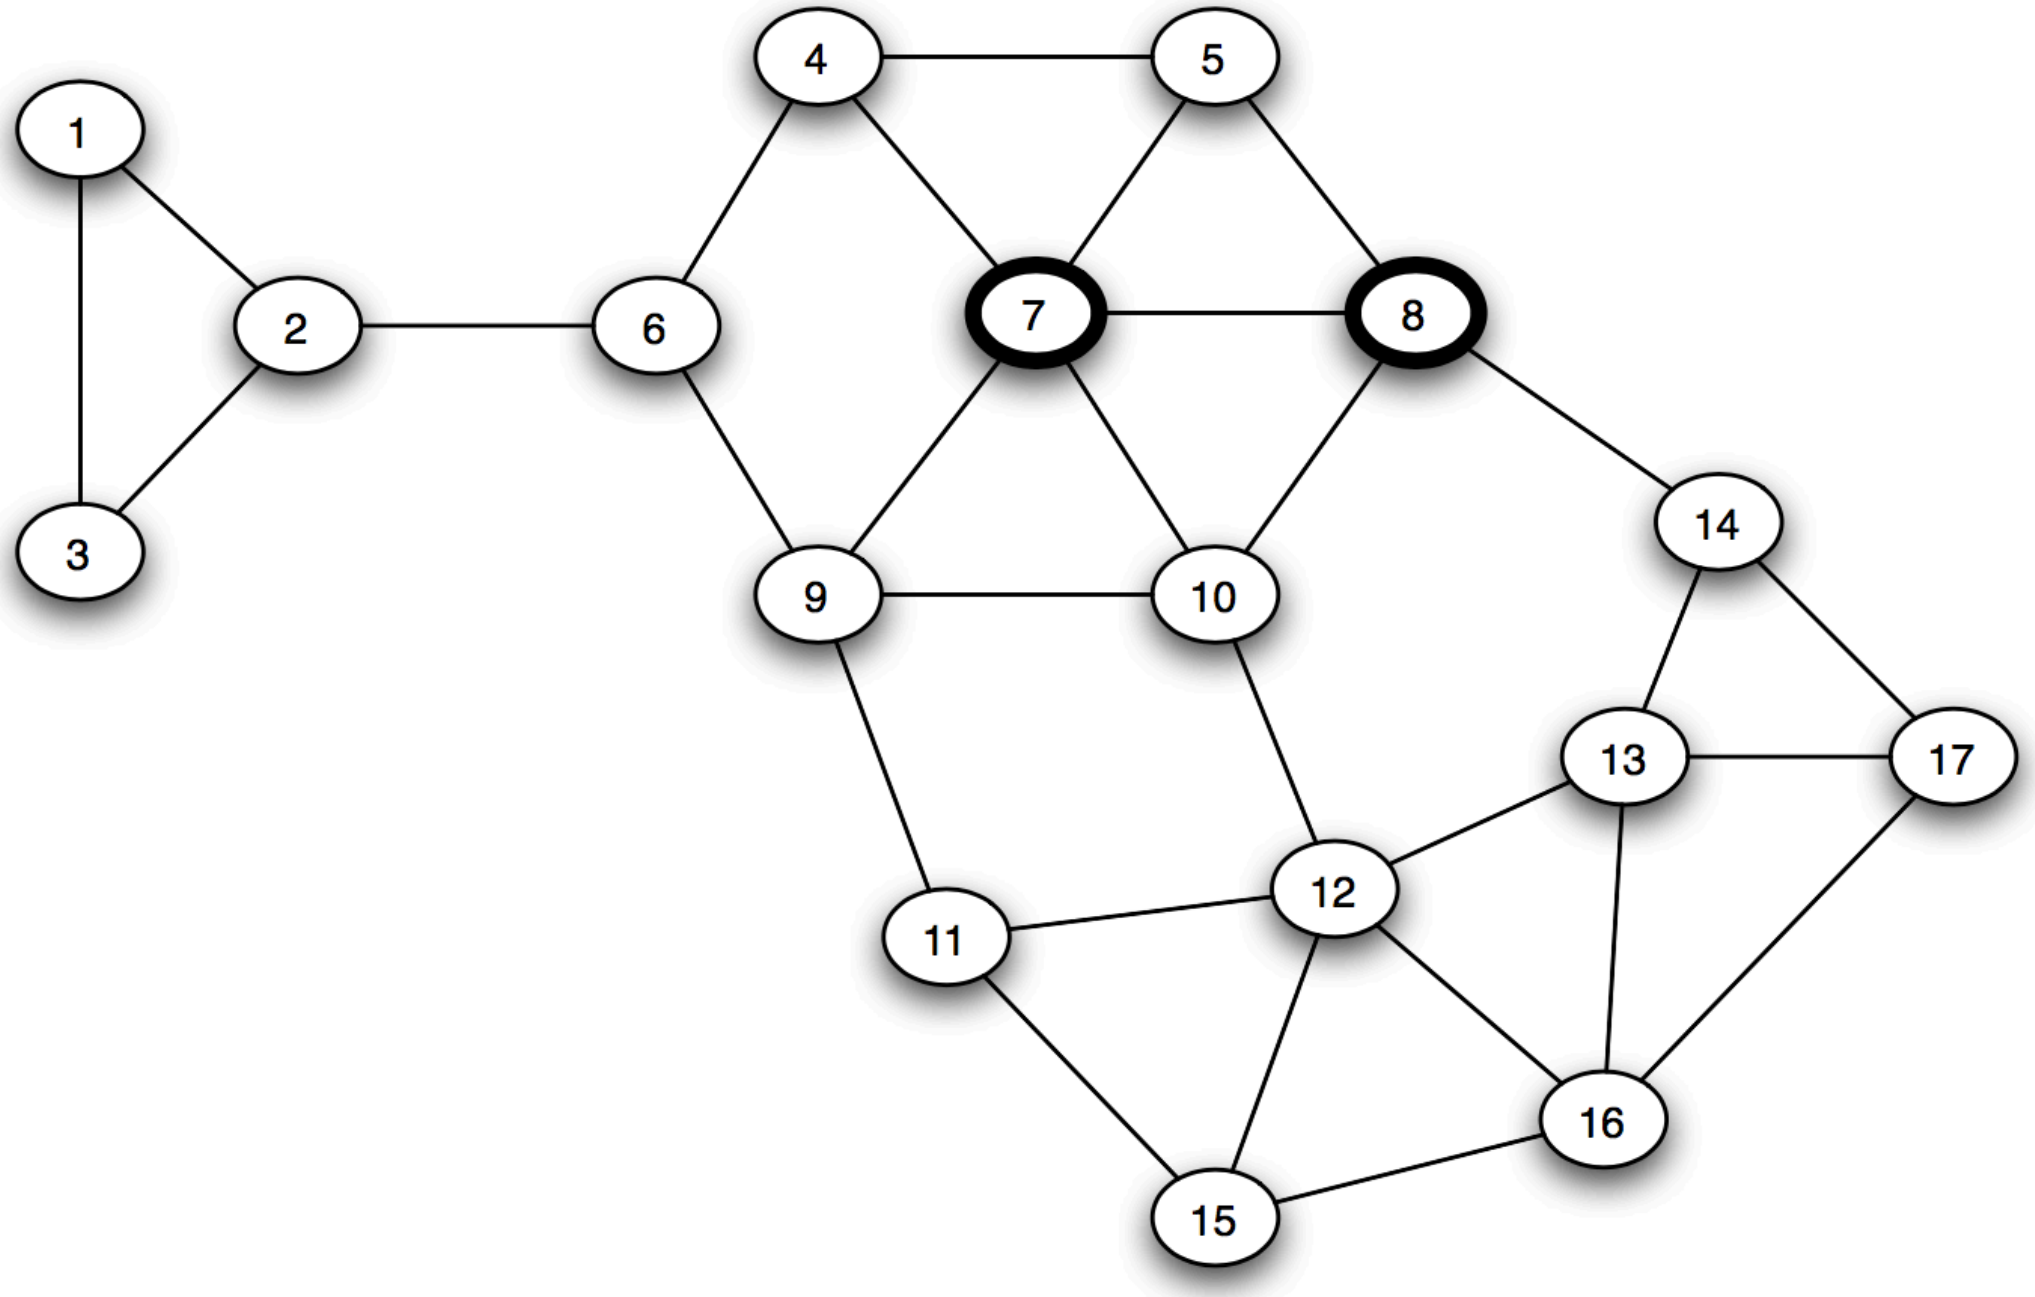

Another example with a=3 and b=2

$$p > \frac{b}{a+b} = q$$

q = 2/5

(new technology better, so q<1/2)

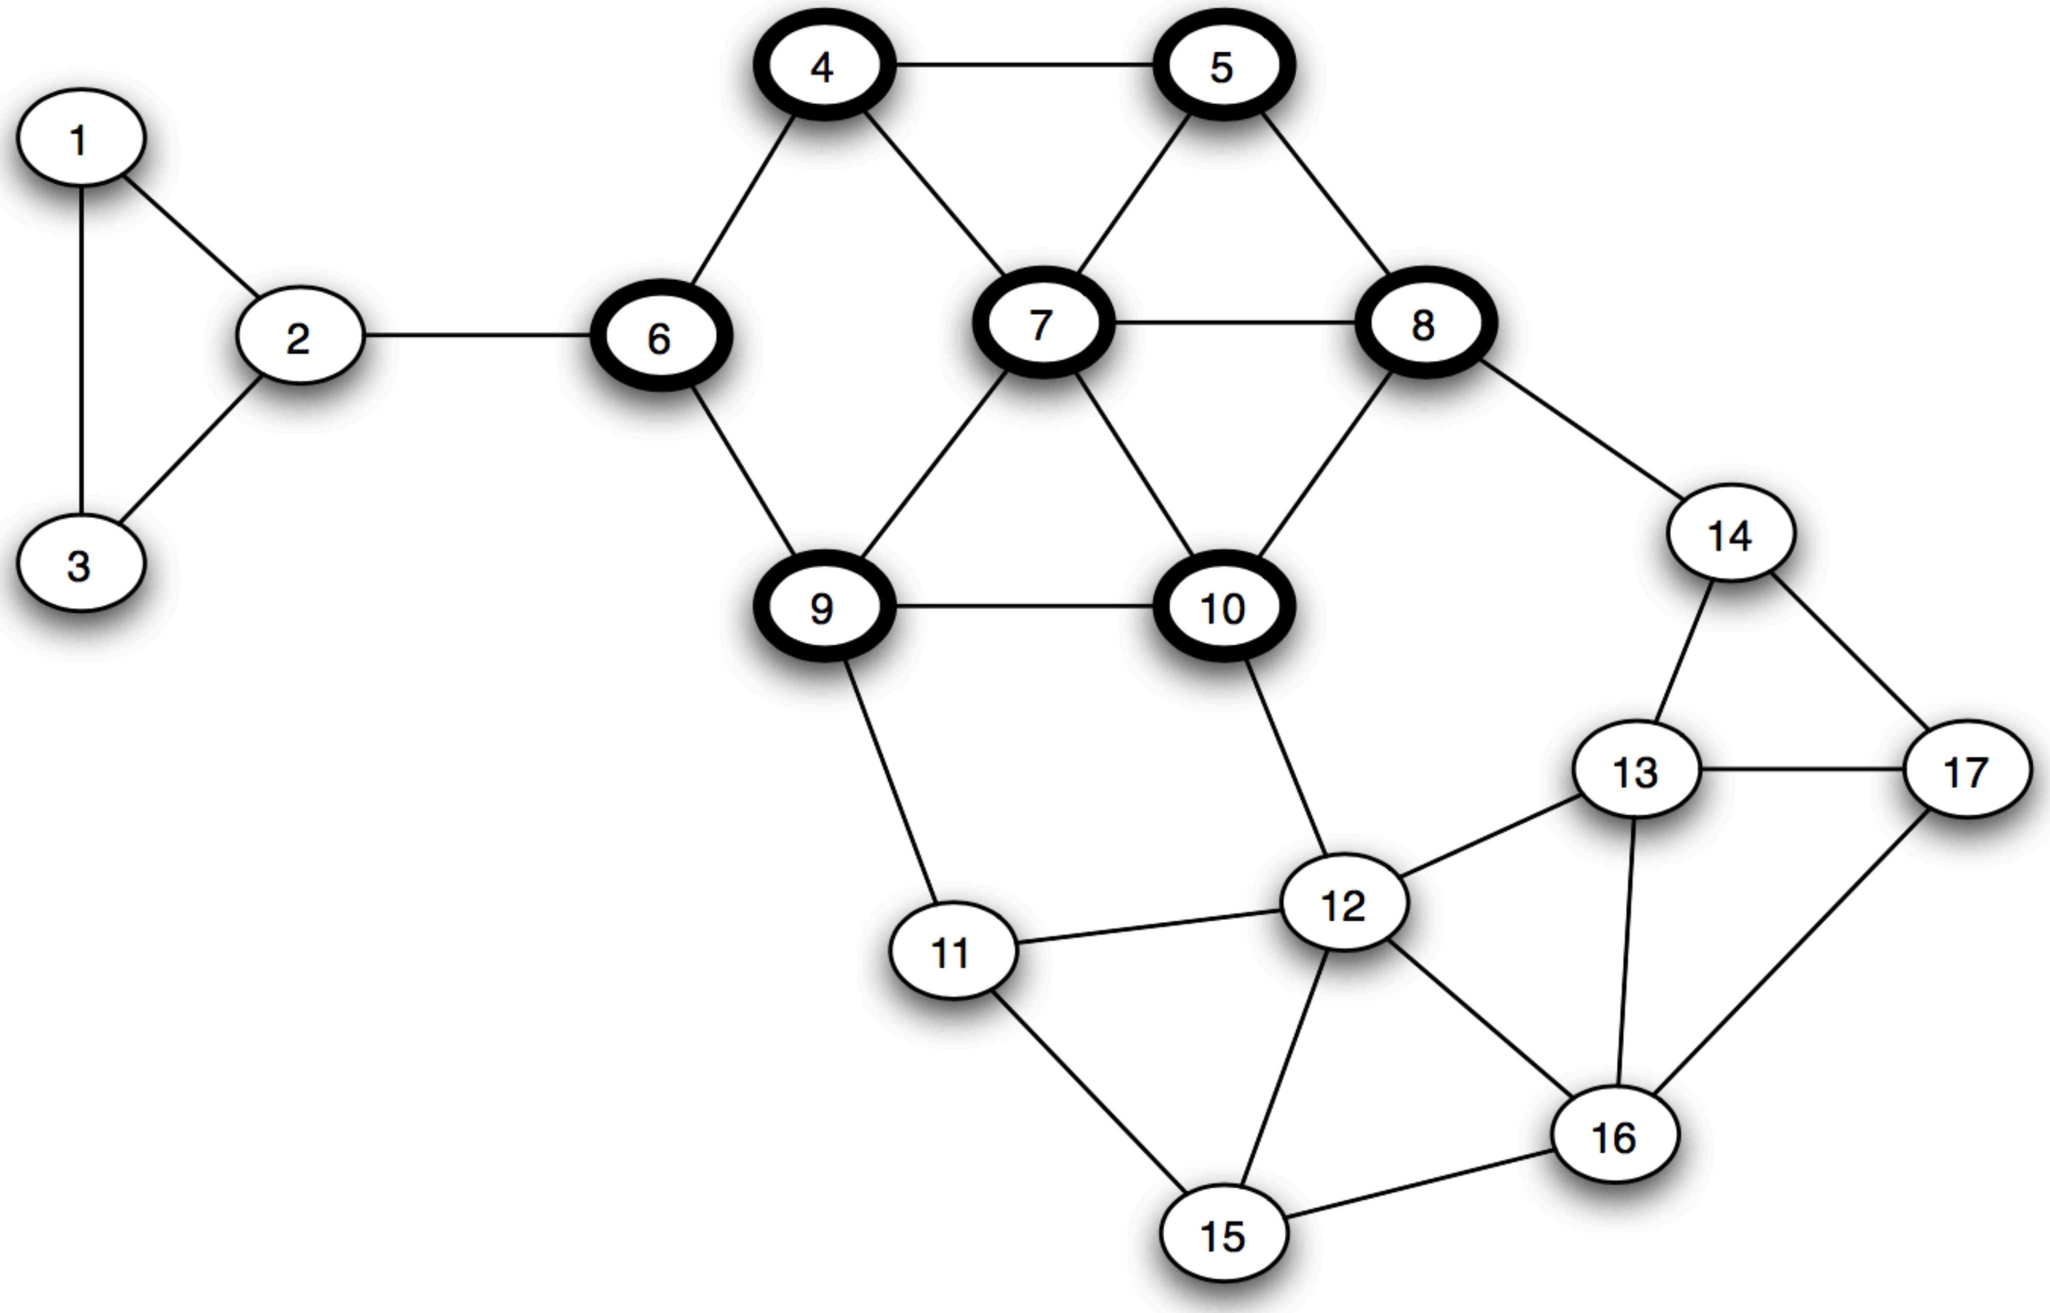

After three steps it stops

Another example with a=3 and b=2

A spread to nodes with sufficiently dense internal connectivity

But it could never bridge the "gaps" that separate nodes 8–10 and 11–14, and node 6 and node 2

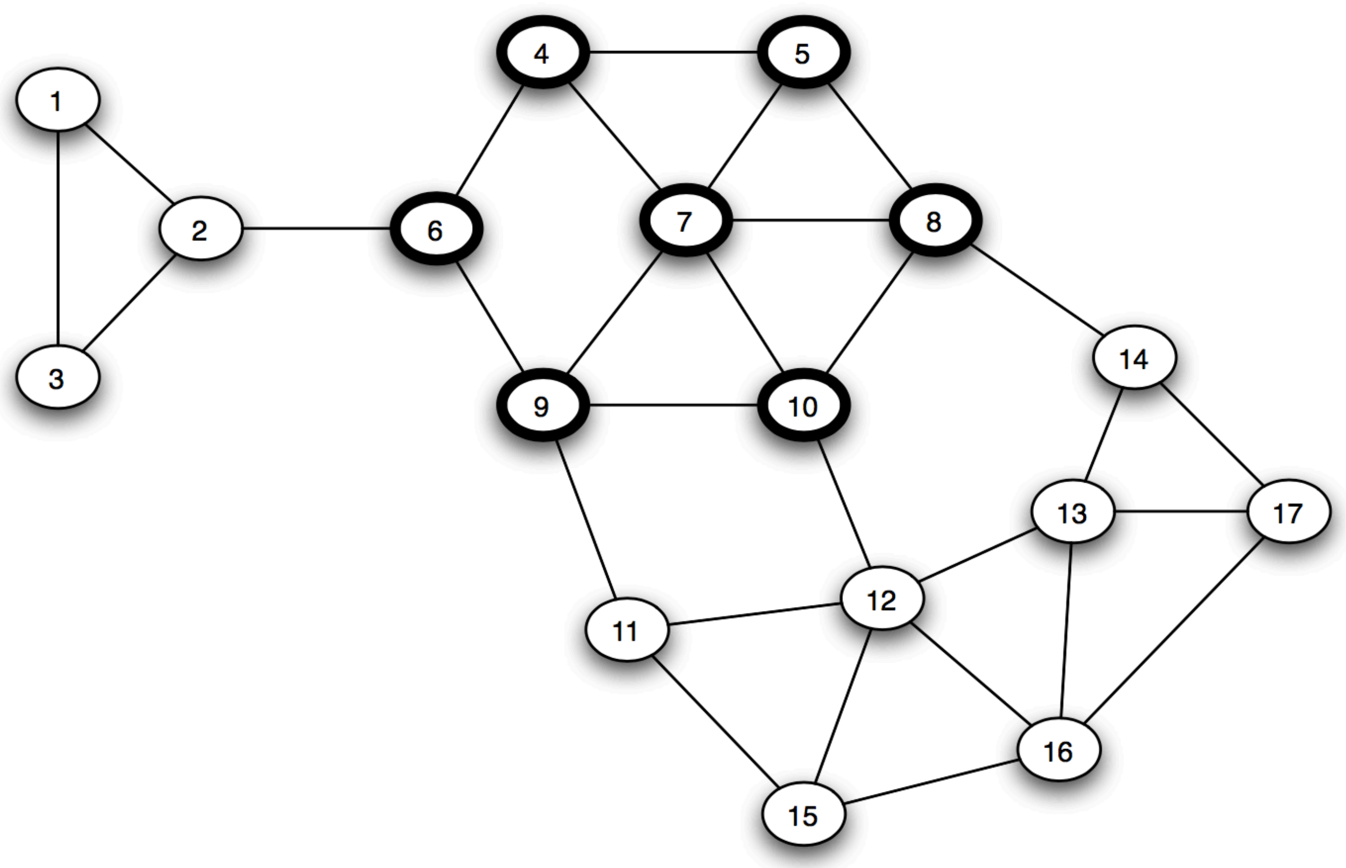

Result: coexistence of **A** and **B**, boundaries in the network where the two meet

- Different dominant political/religious views between adjacent communities
- Different social networking sites dominated by different age groups and lifestyles
- Windows vs. Mac (some industries heavily use Mac, even though Windows generally dominates)

Another example with a=3 and b=2

What could **A** do to improve its reach?

Raise quality of the product:

- If payoff in underlying coordination game improves from a=3 to a=4
- Threshold to switch drops from q=2/5 to q=1/3
- All nodes eventually switch to **A**

Slightly increasing the quality of innovations can dramatically alter their reach

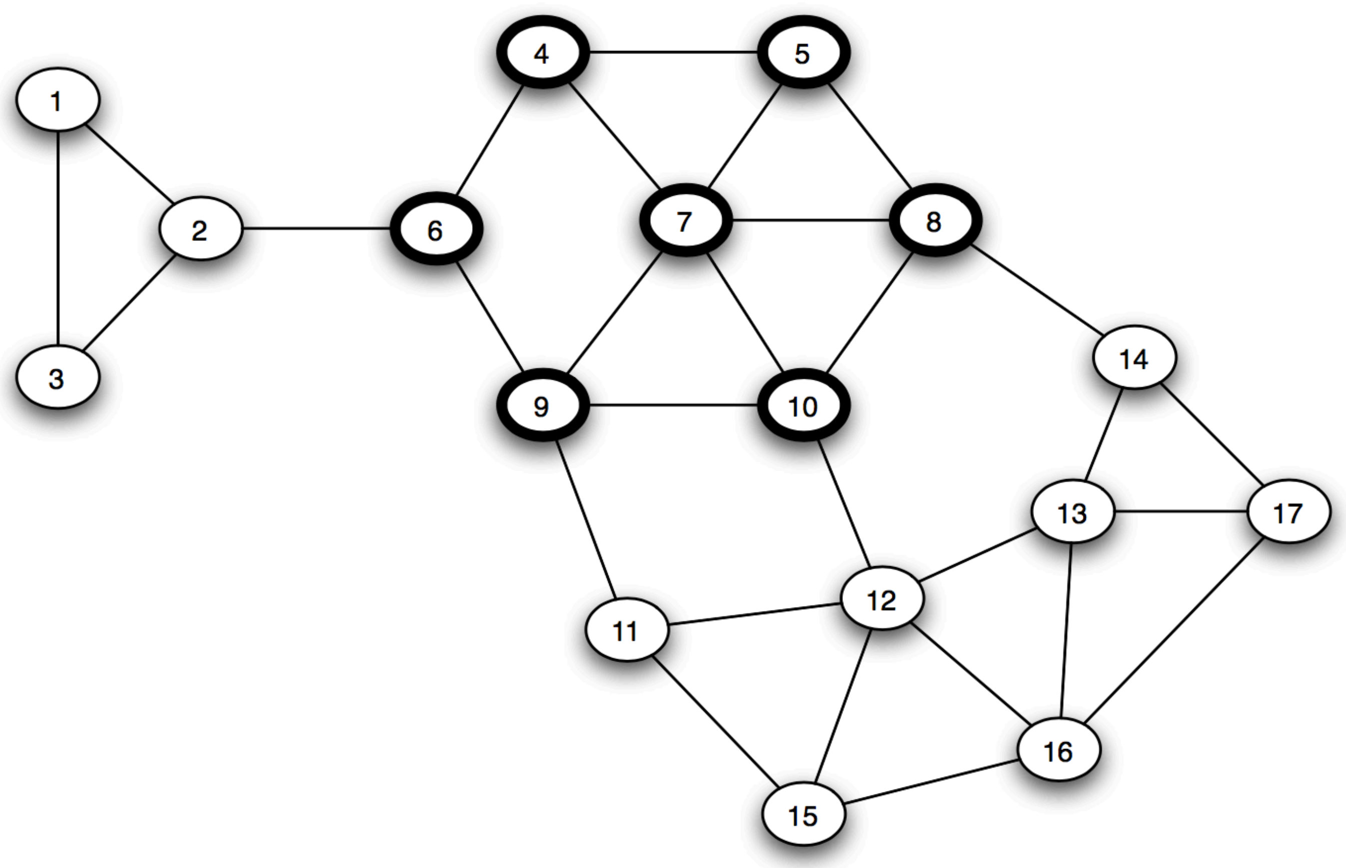

Another example with a=3 and b=2

What could **A** do to improve its reach?

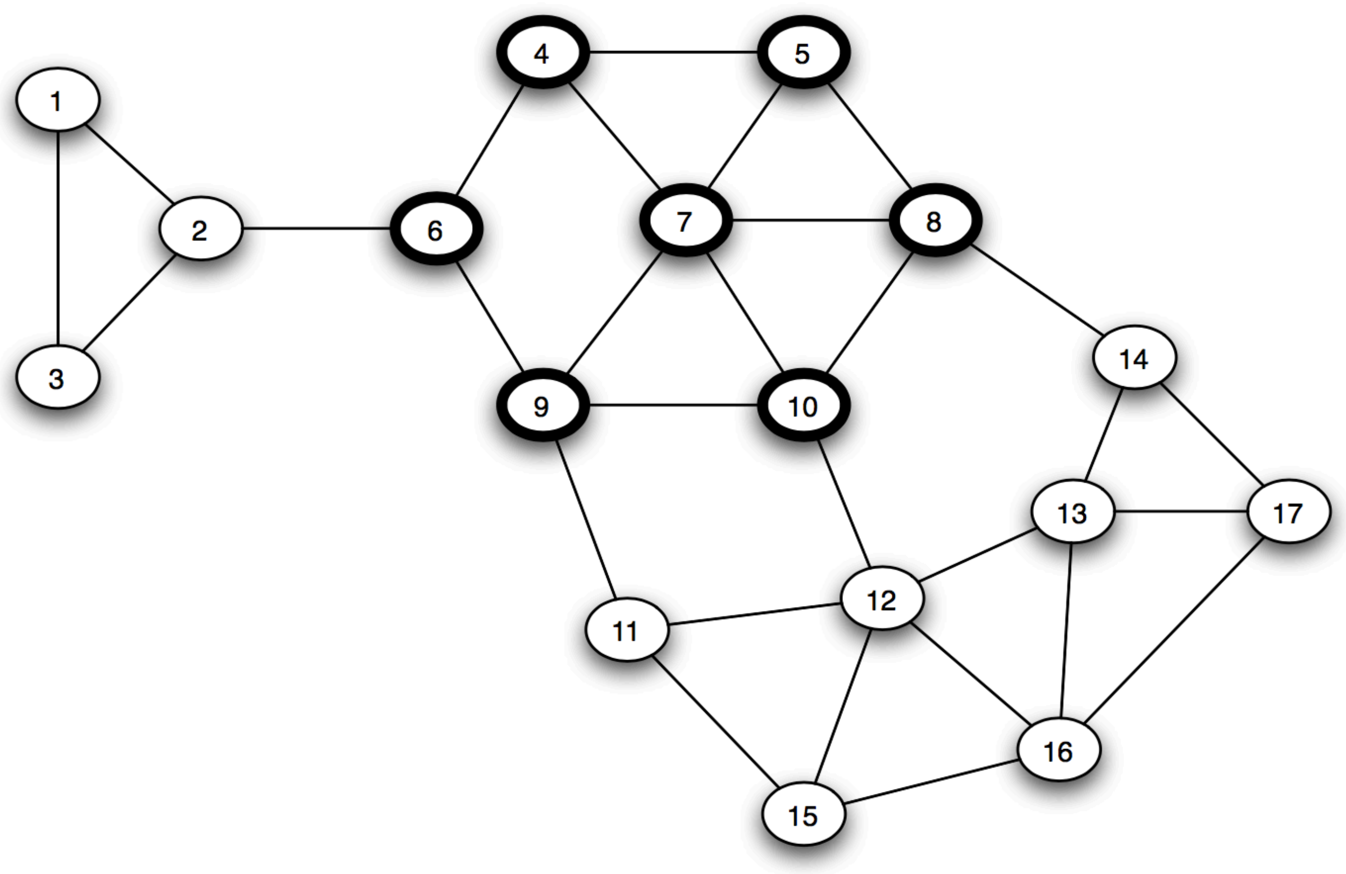

Convince key people to be early adopters

- Sometimes it's impossible to raise the quality any higher than it already is
- Threshold stays the same (here q=2/5)
- If 12 or 13 switch, then all nodes 11–17 switch
- If 11 or 14 switch, nothing else happens

Certain people occupy **structurally important positions**

Another example with a=3 and b=2

What are the impediments to spread?

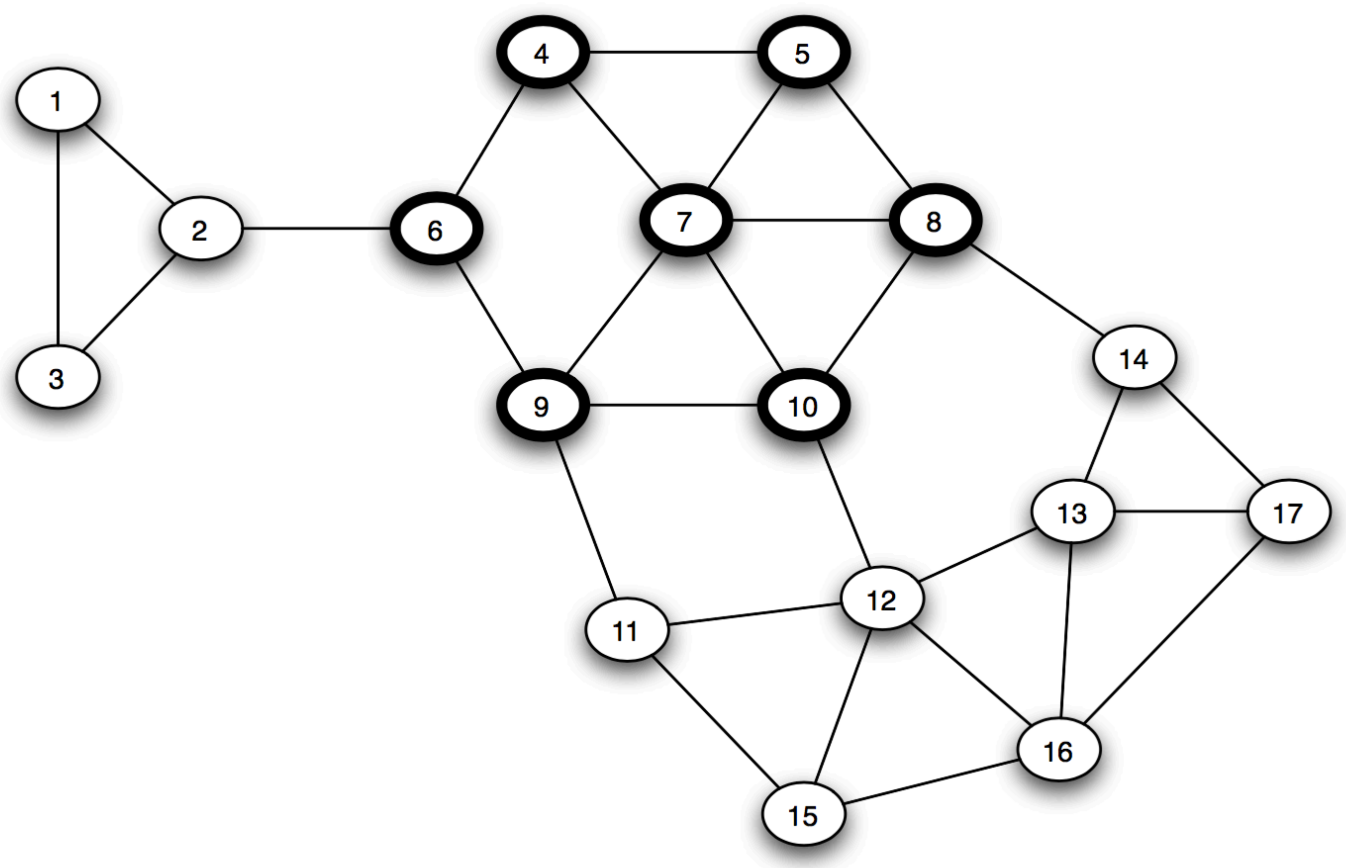

Densely connected communities

- 1–3 are well-connected with each other but poorly connected to the rest of the network
- Similar story for 11–17
- **Homophily impedes diffusion**

A **cluster of density *p*** is a set of nodes such that every node in the set has at least a p fraction of its neighbours in the set

Nodes {1,2,3} are a cluster of density p = ?

Nodes {11,12,13,14,15,16,17} are a cluster of density p = ?

Another example with a=3 and b=2

What are the impediments to spread?

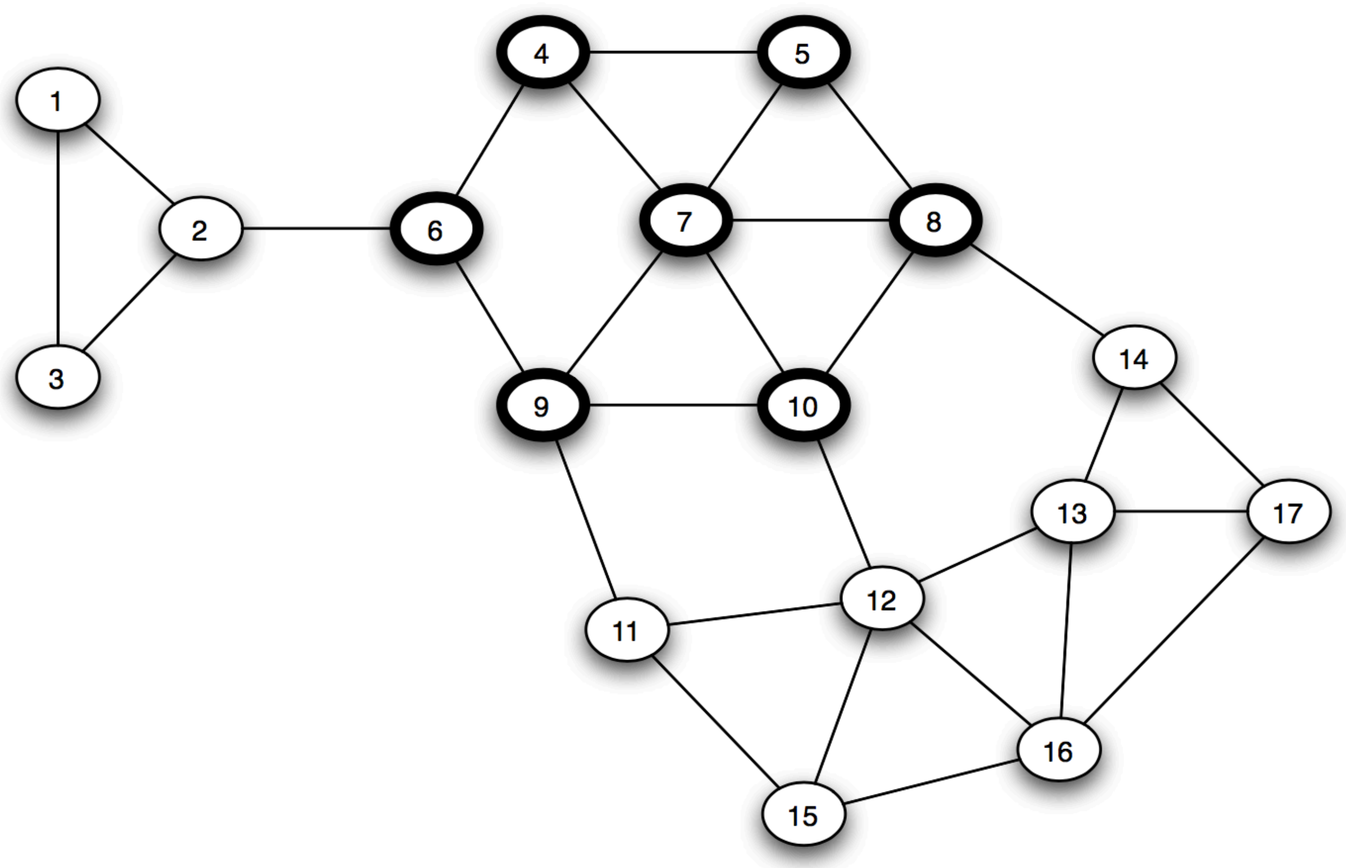

Densely connected communities

- 1–3 are well-connected with each other but poorly connected to the rest of the network
- Similar story for 11–17
- **Homophily impedes diffusion**

A **cluster of density *p*** is a set of nodes such that every node in the set has at least a p fraction of its neighbours in the set

Nodes {1,2,3} are a cluster of density p = *2/3*

Nodes {11,12,13,14,15,16,17} are a cluster of density p = *2/3*

Another example with a=3 and b=2

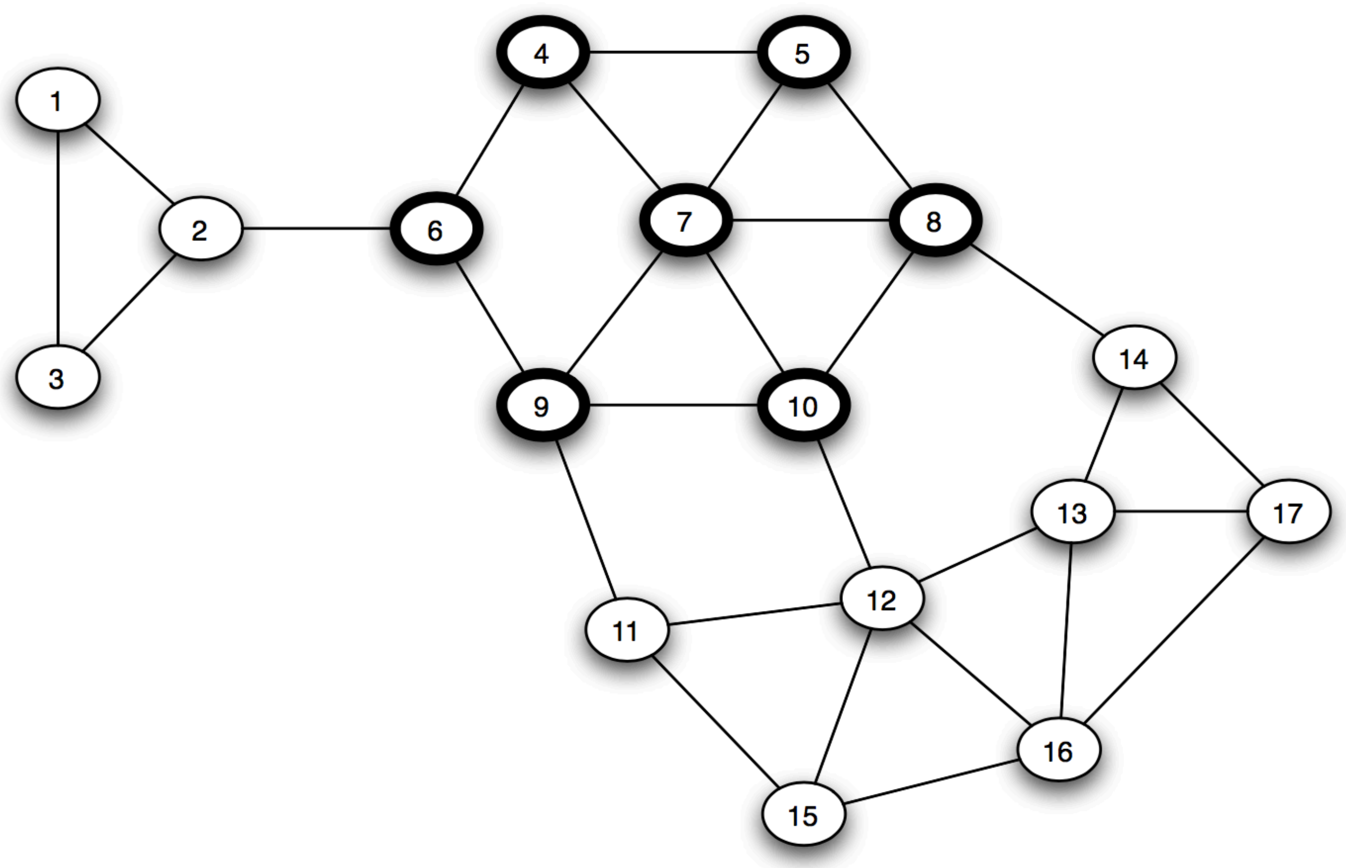

<u>Fact</u>: Consider a set of initial adopters of behavior A, with a threshold of q for nodes in the remaining network to adopt behavior A.

- If the remaining network contains a cluster of density greater than 1−q, then the set of initial adopters will not cause a complete cascade.
- Moreover, whenever a set of initial adopters does not cause a complete cascade with threshold q, the remaining network must contain a cluster of density greater than 1−q

In this model, densely connected communities are impediments to diffusion — and they are the only impediments to diffusion

Monotonic Spreading

Observation: Use of A spreads monotonically
(Nodes only switch B→A, but never back to B)

Why? Proof sketch:

Nodes keep switching from B to A: *B*→*A*

Now, suppose some node switched back from *A*→*B*, consider the first node *u* to do so (say at time *t*)

Earlier at some time *t'* (*t'*<*t*) the same node *u* switched *B*→*A*

So at time *t'* *u* was above threshold for *A*

But up to time *t* no node switched back to B, so node *u* could only have more neighbors who used *A* at time *t* compared to *t'*.
There was no reason for *u* to switch at the first place!

‼ Contradiction ‼

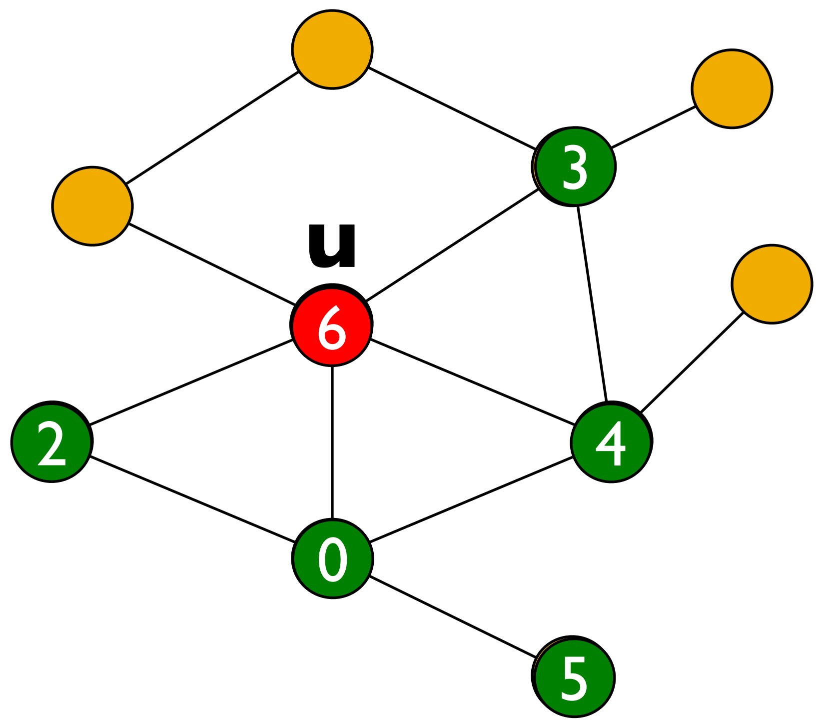

Infinite Graphs

Consider <u>infinite</u> graph *G*
(but each node has finite number of neighbors!)

v chooses ***A*** if ***p>q***

$$q = \frac{b}{a+b}$$

We say that a finite set ***S*** causes a **complete cascade** in *G* with **threshold *q*** if, when ***S*** adopts ***A***, eventually **every node in G adopts *A***

Example: **Path**

If ***q<1/2*** then cascade occurs

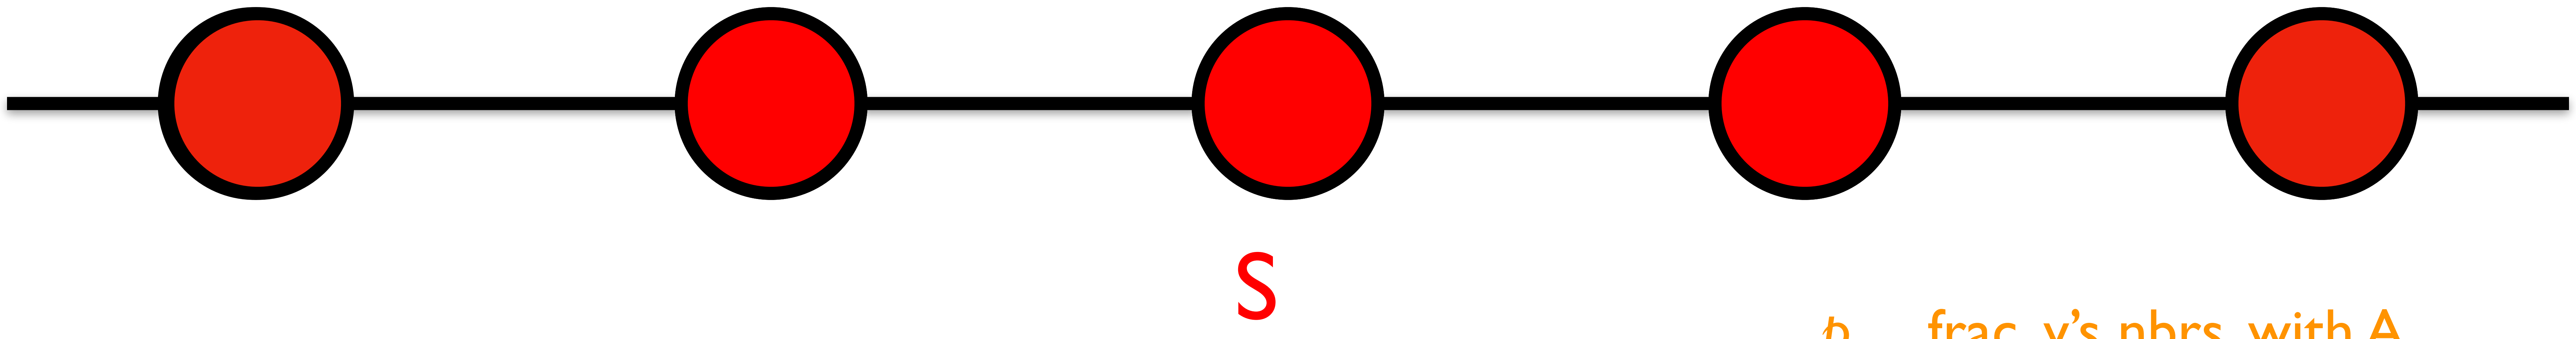

p… frac. v's nbrs. with A
q… payoff threshold

Infinite Graphs

Infinite Tree:

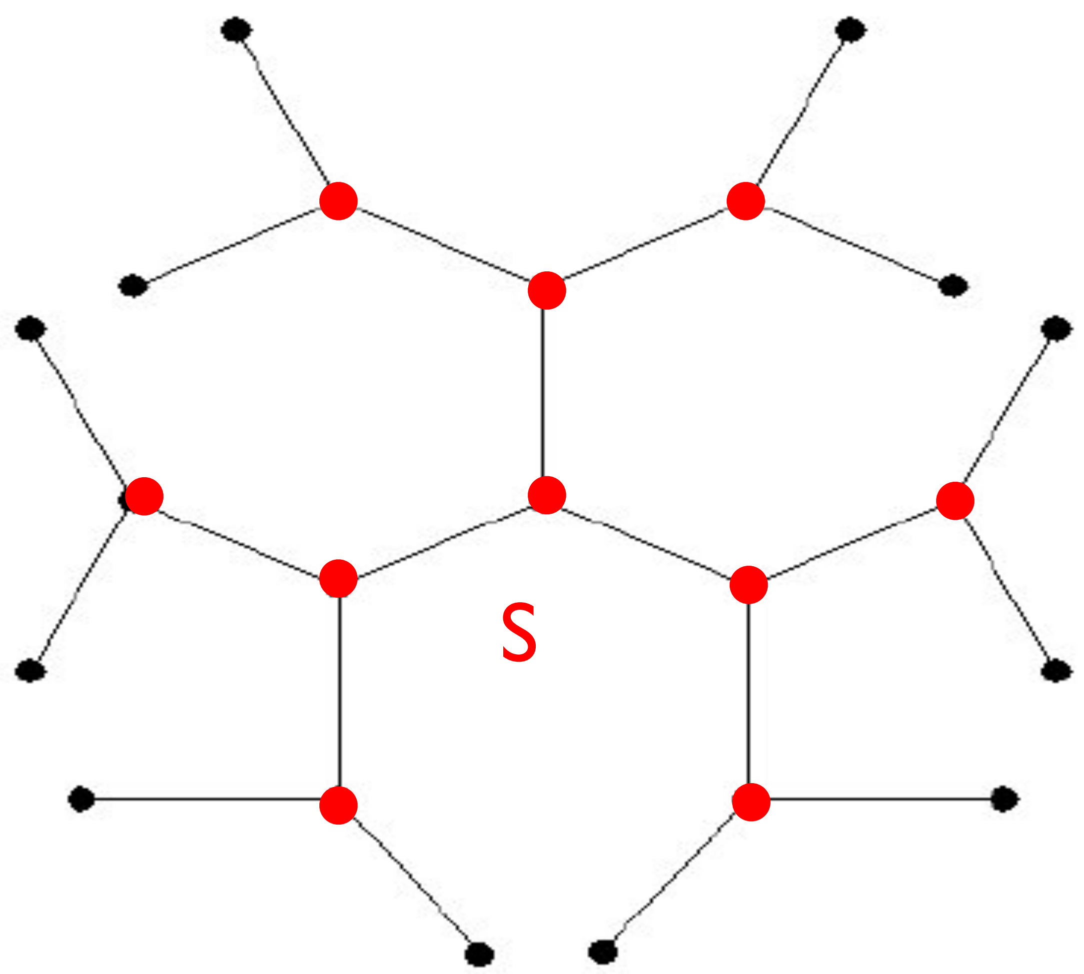

If $q<1/3$ then cascade occurs

Infinite Grid:

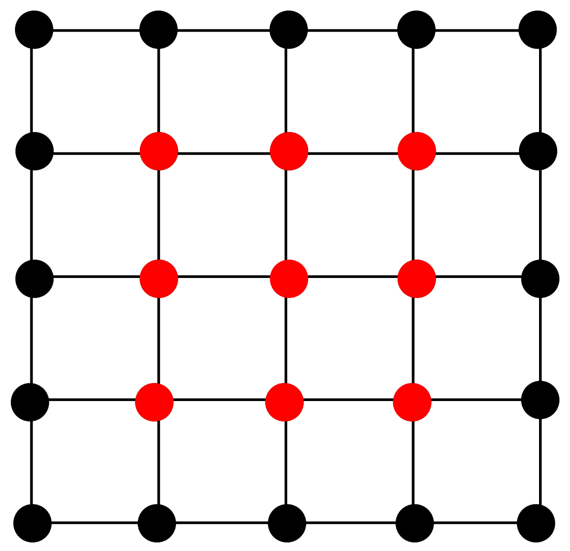

If $q<1/4$ then cascade occurs

Information Diffusion

Influence Through Networks

- If people are connected through a network, it's possible for them to influence each other's behaviour and actions
- Today: why?
  - Direct benefit
  - Informational
  - Social conformity

Information Diffusion: Media

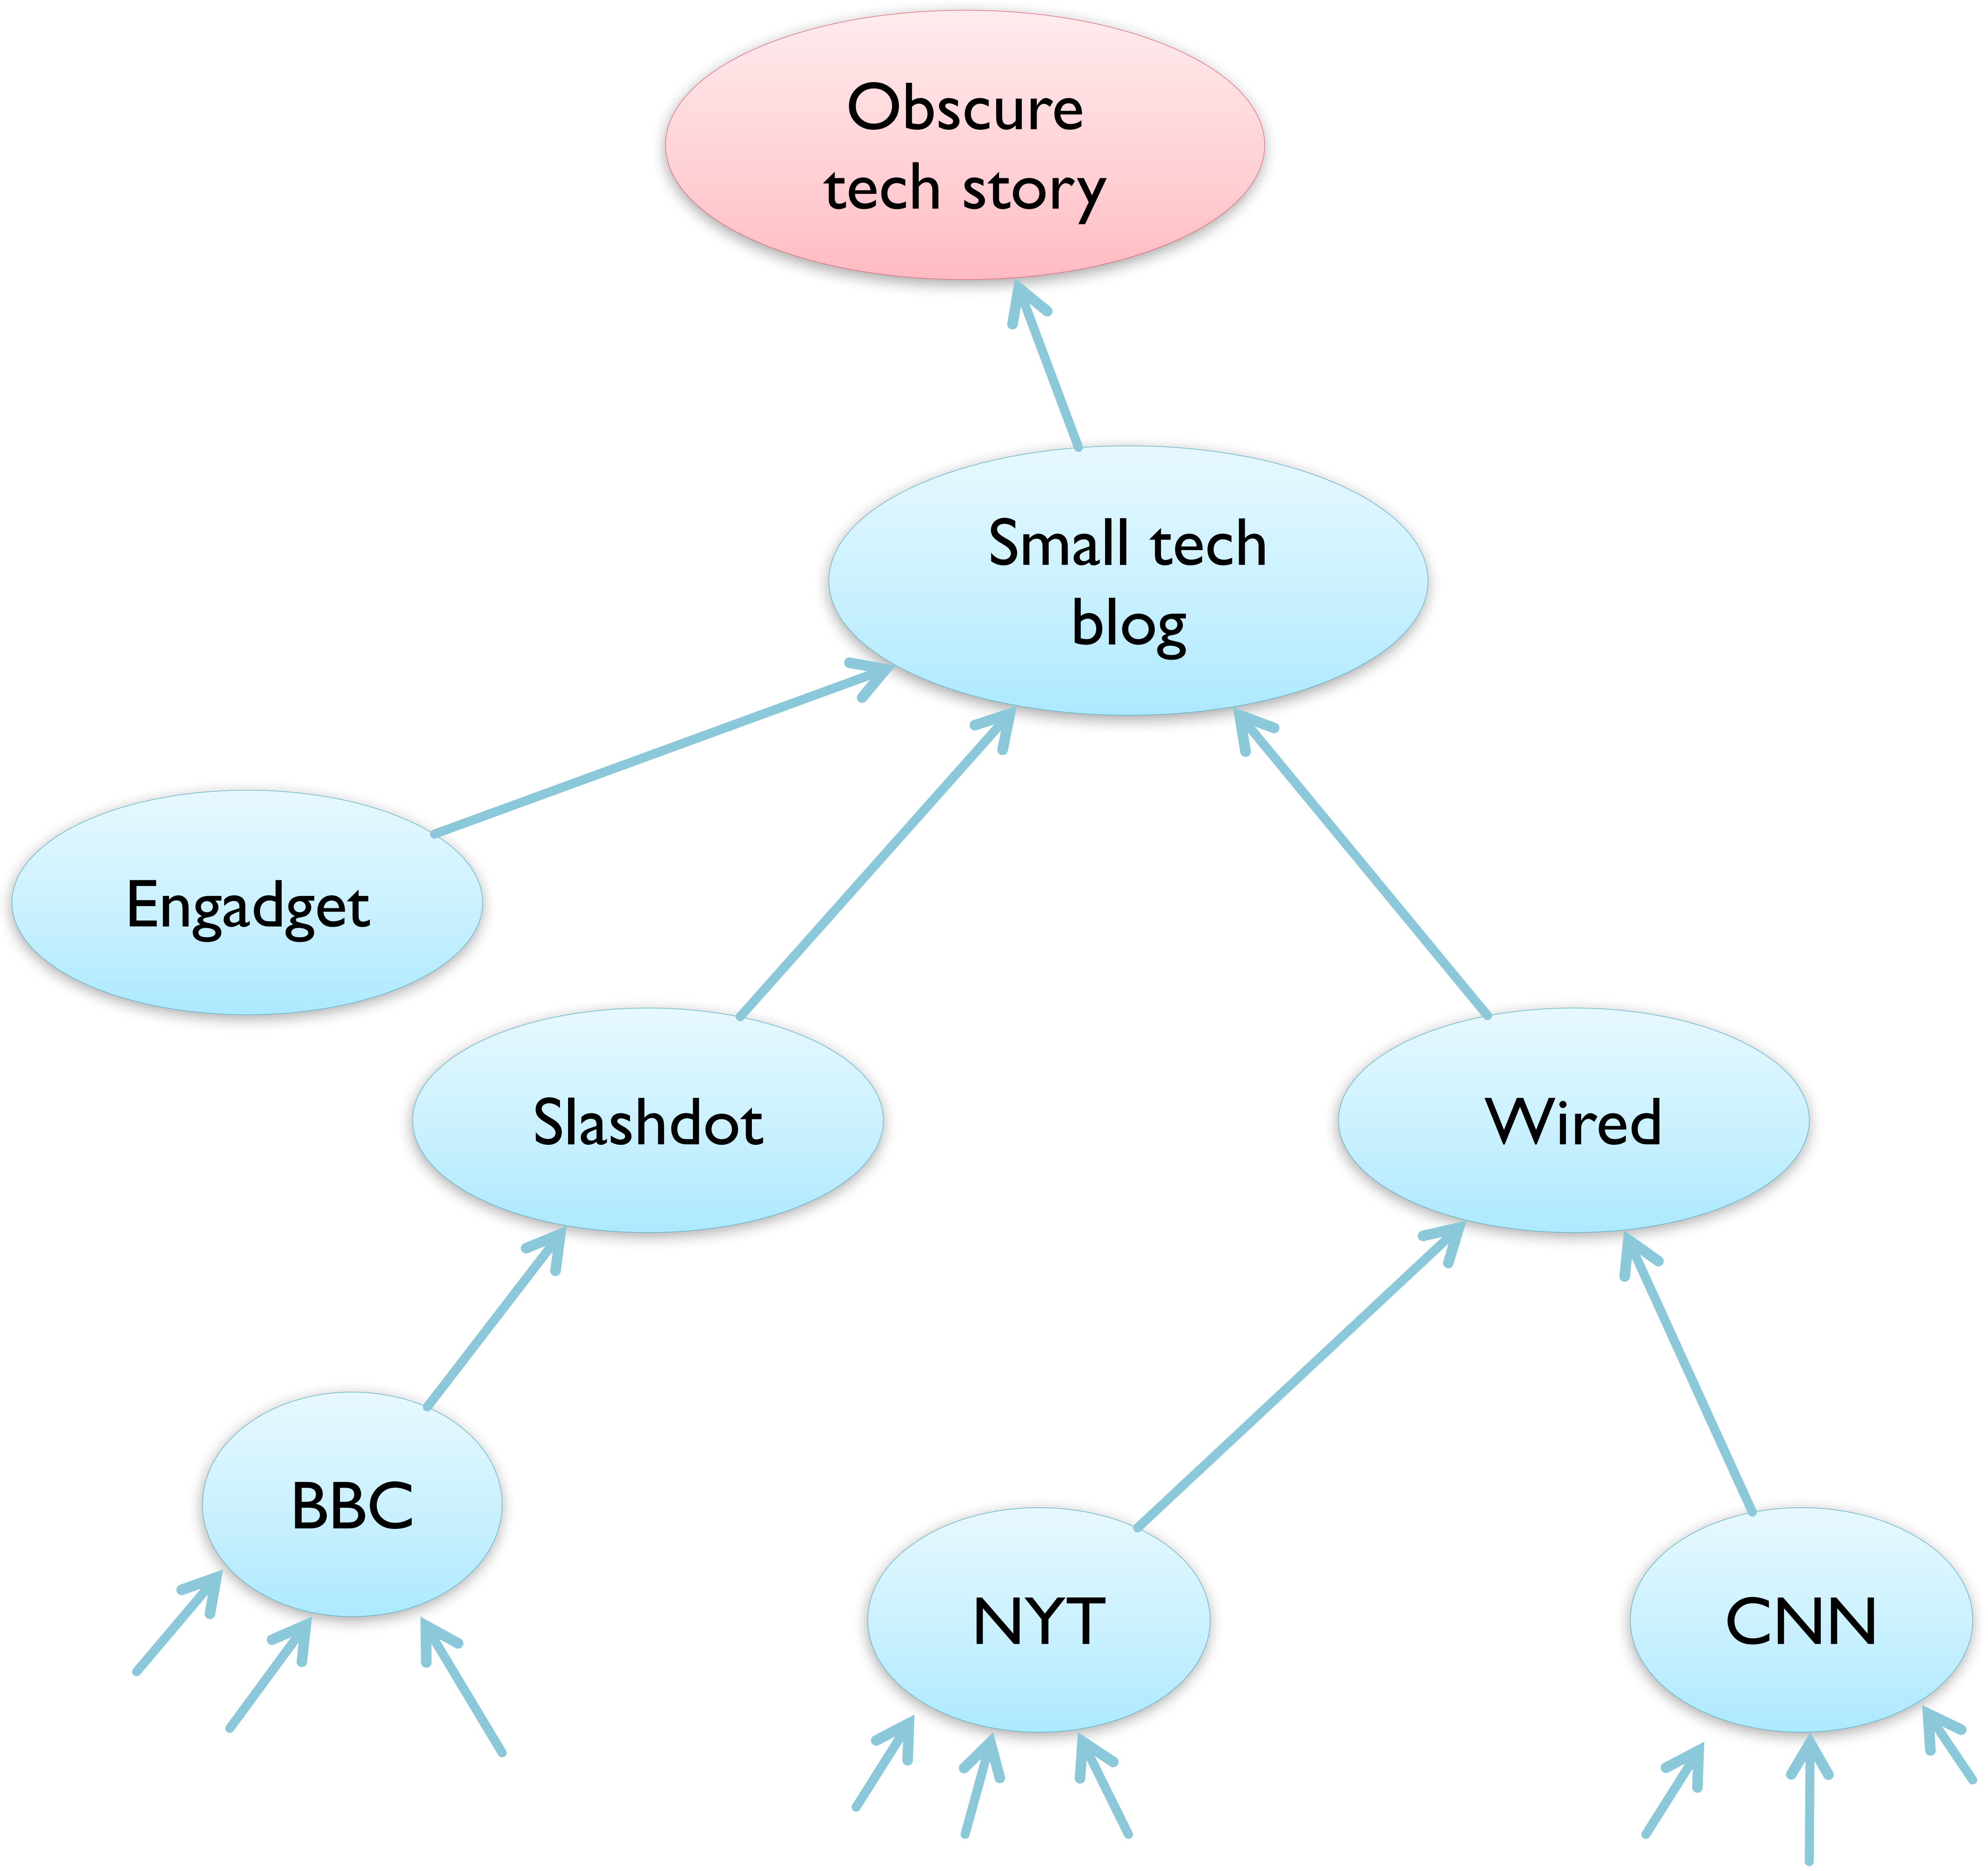

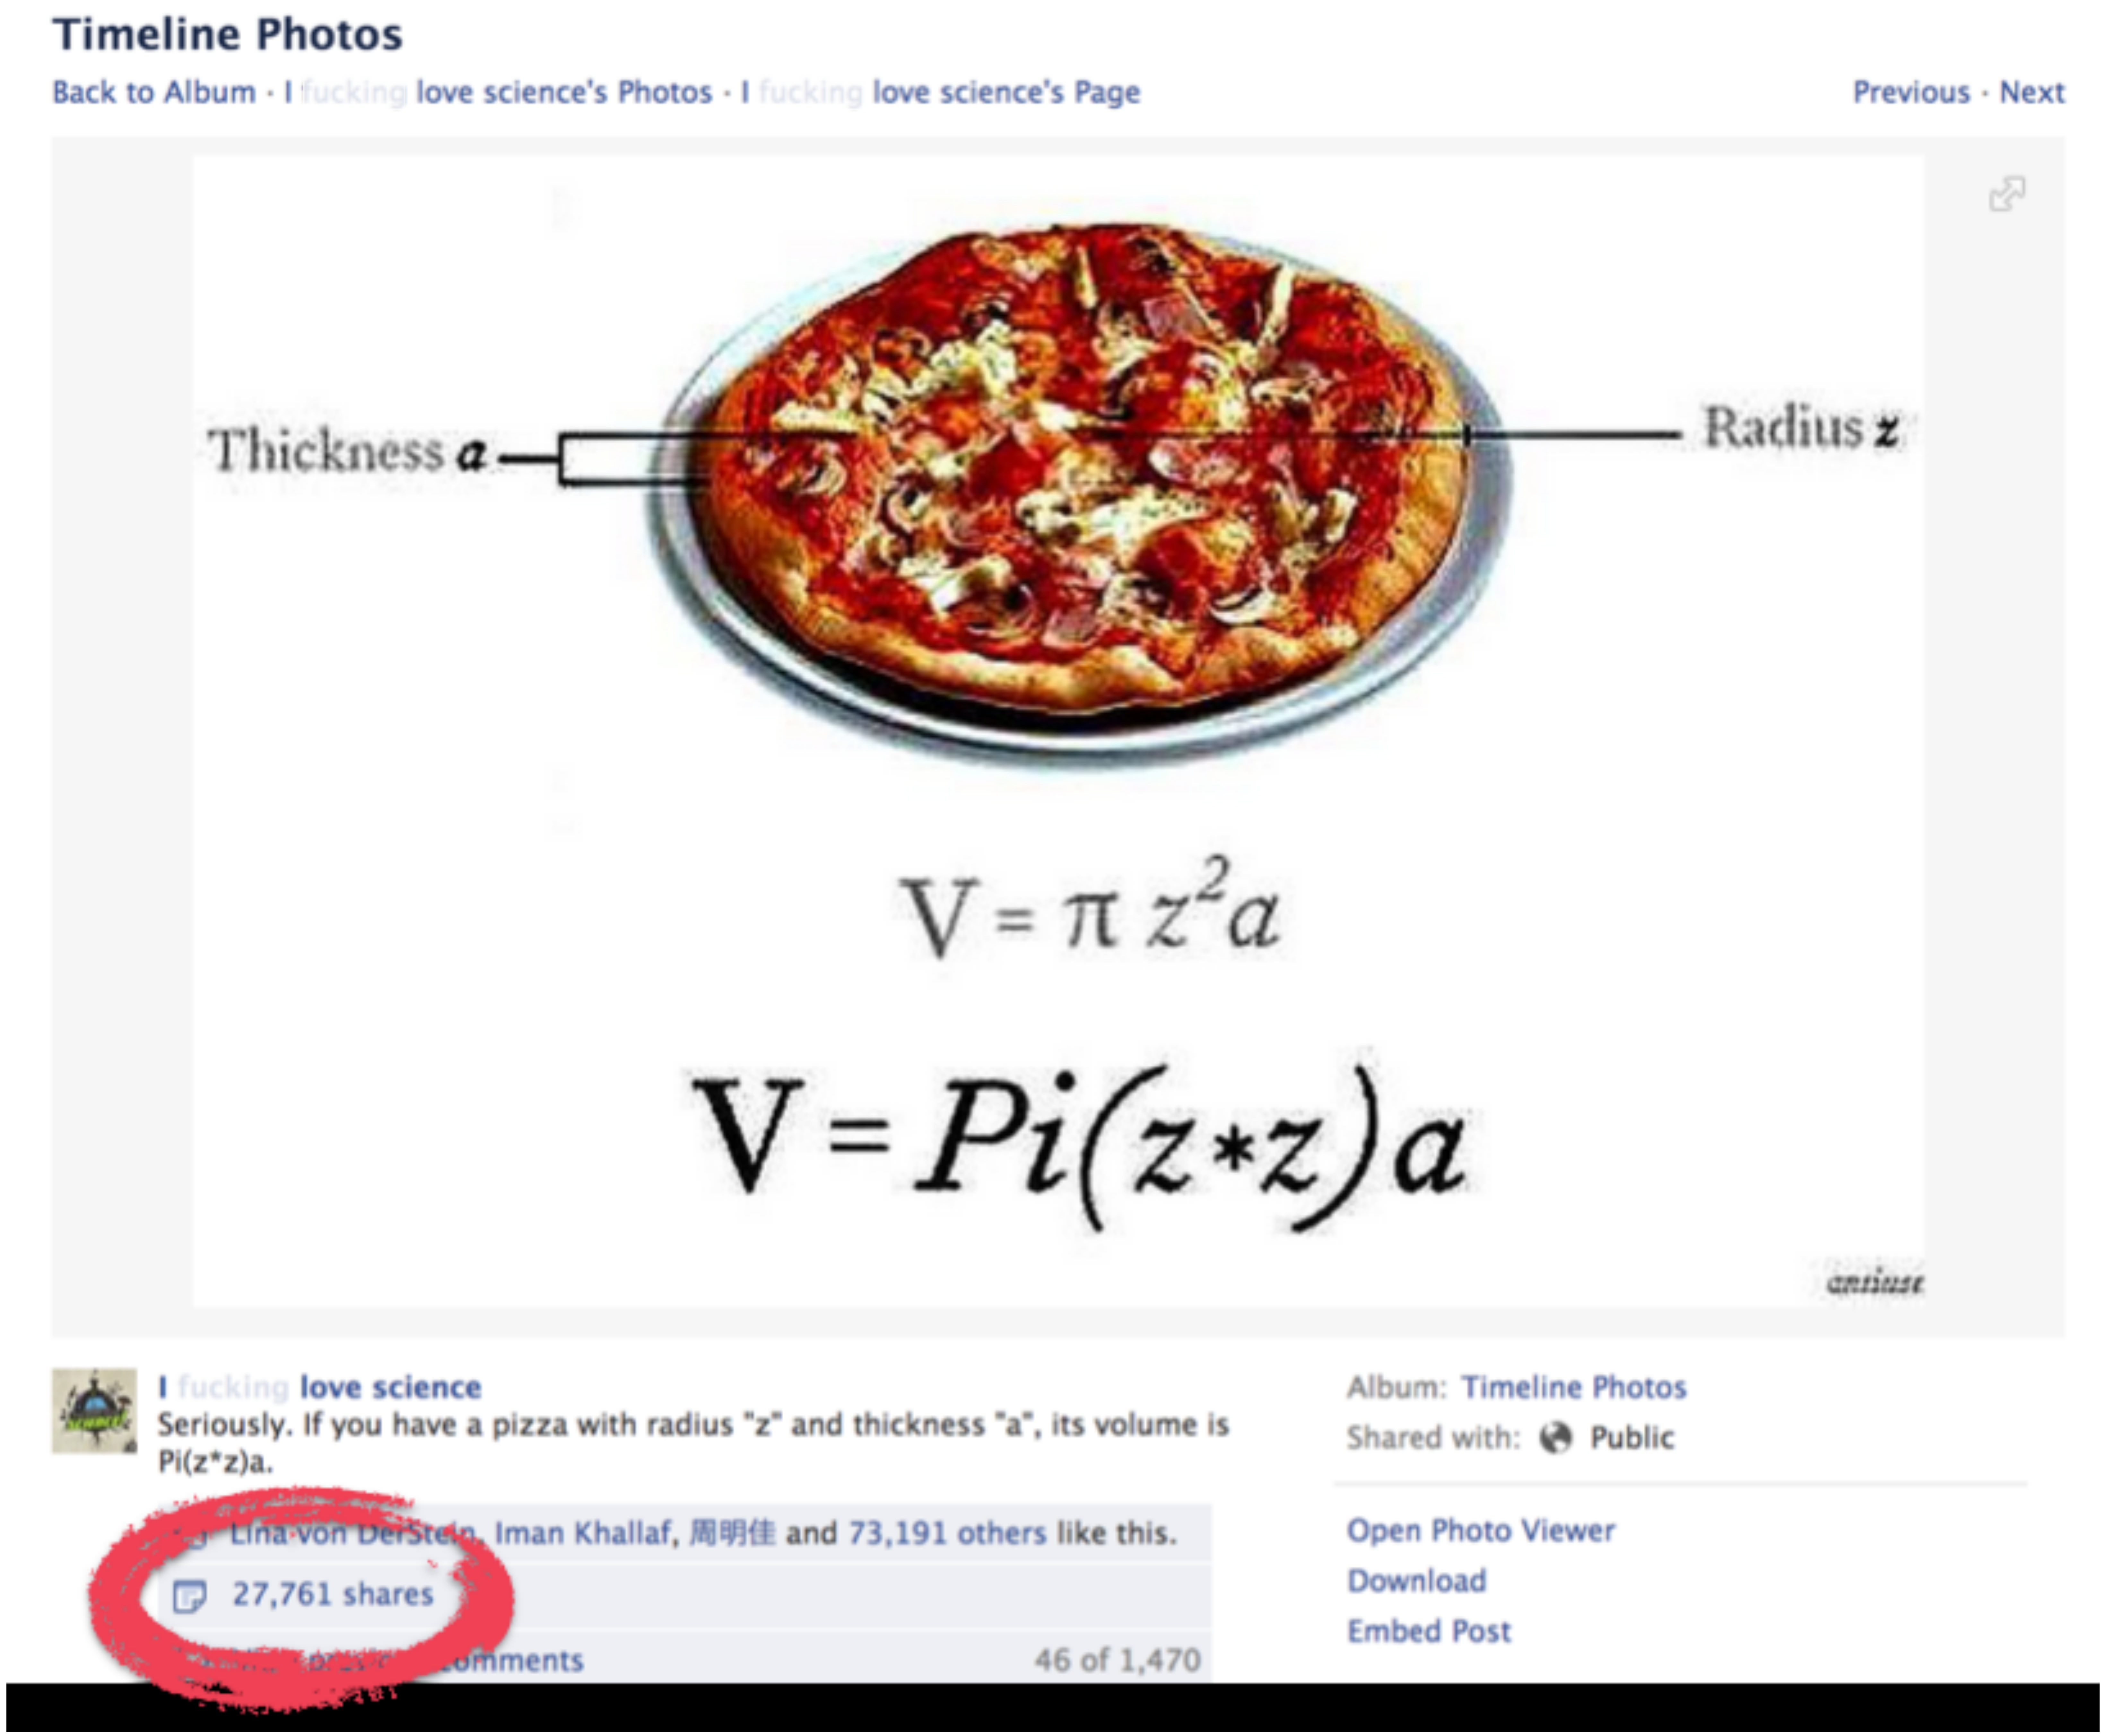
Timeline Photos
Back to Album · I fucking love science's Photos · I fucking love science's Page
Previous · Next
Thickness a
Radius z
$V = \pi z^2 a$
$V = Pi(z*z)a$
I fucking love science
Seriously. If you have a pizza with radius "z" and thickness "a", its volume is Pi(z*z)a.
Lina von DerStein, Iman Khallaf, 周明佳 and 73,191 others like this.
27,761 shares
46 of 1,470
Album: Timeline Photos
Shared with: Public
Open Photo Viewer
Download
Embed Post

Simple Herding Model: Lessons

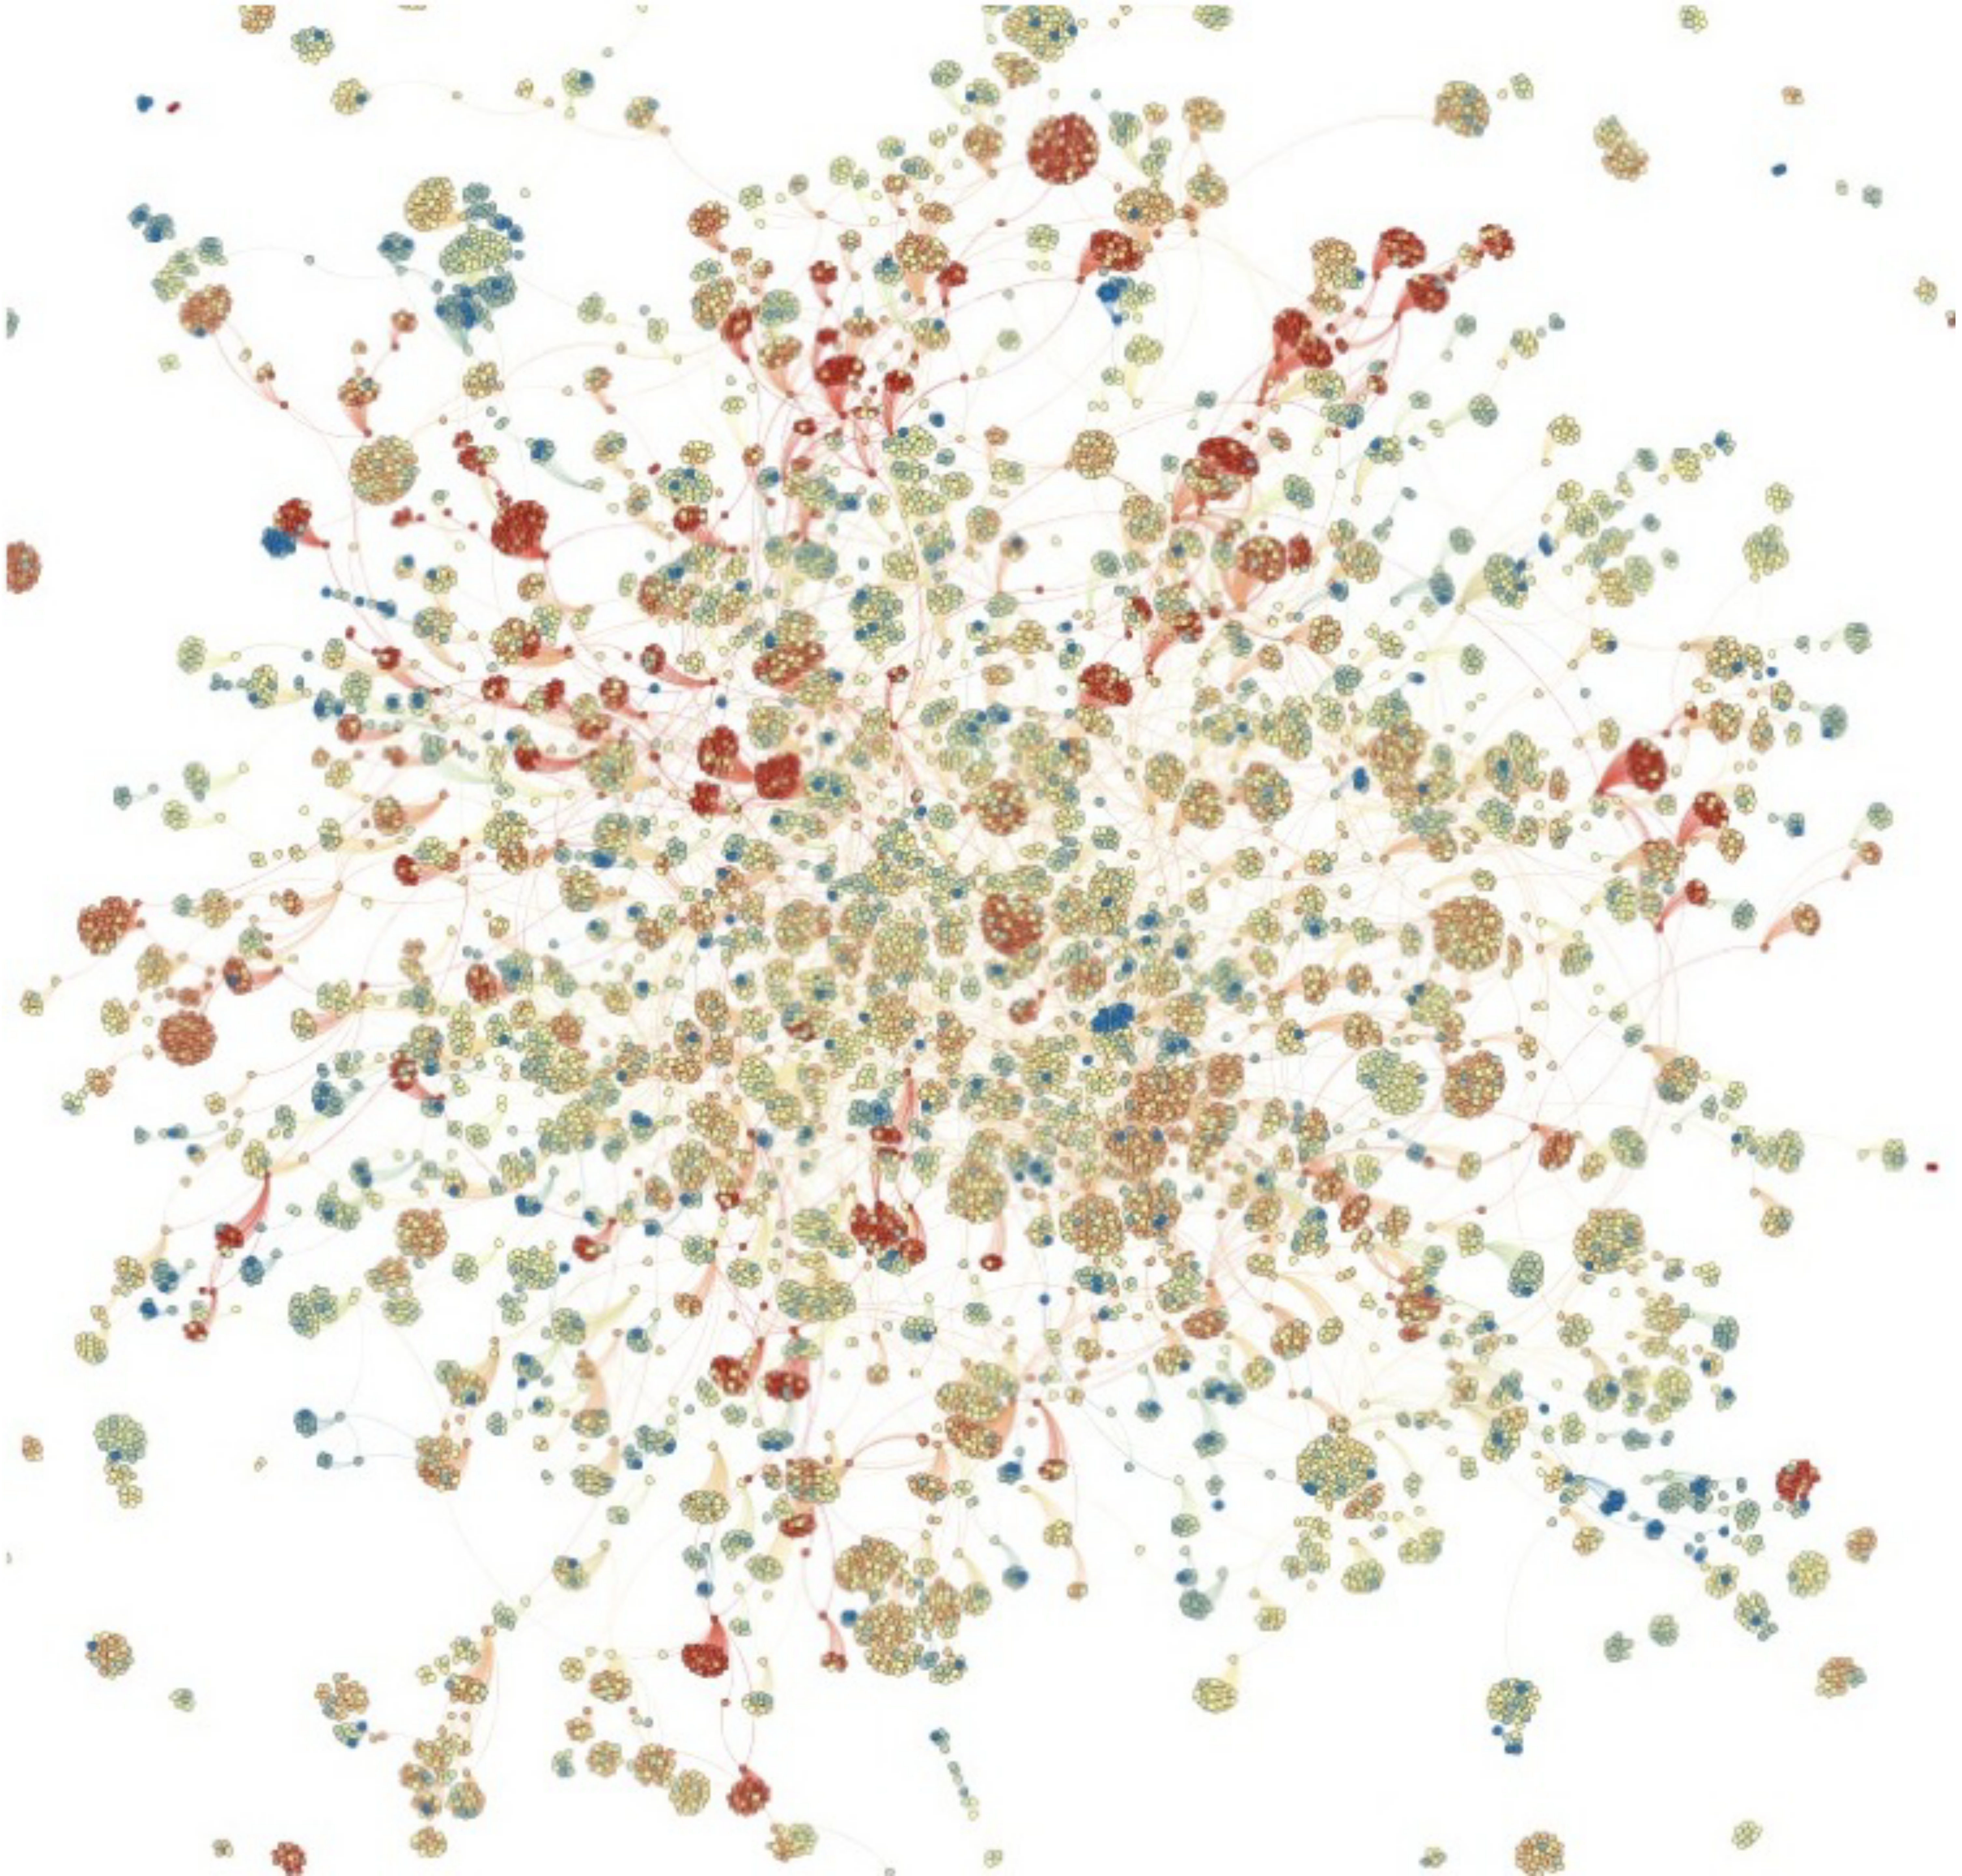

Information-Based Model of Diffusion: Crowd Herding

People influencing each other

Almost infinite number of ways:

- Opinions
- Product purchases
- Political positions
- Technologies used
- etc…

Good reasons for this! Sometimes it's better to follow the crowd than trust your information

A simple example

Going to Yellowknife

Do some research, intend to eat at resto A

But you show up and no one's eating there, instead lots of people are in resto B!

A rational person may reason that those people know something he doesn't, and go with B as well

Sequential decision making

"Information cascade"

Imitation

In this example, people imitate others, but it's not mindless

Kinds of imitation/influence: informational, social pressure to conform, direct benefits

- Sometimes hard to tell apart

Another example: social pressure or informational?

Experiment: bunch of people stand on a street corner and stare up into the sky
What fraction of passersby stop and look up?

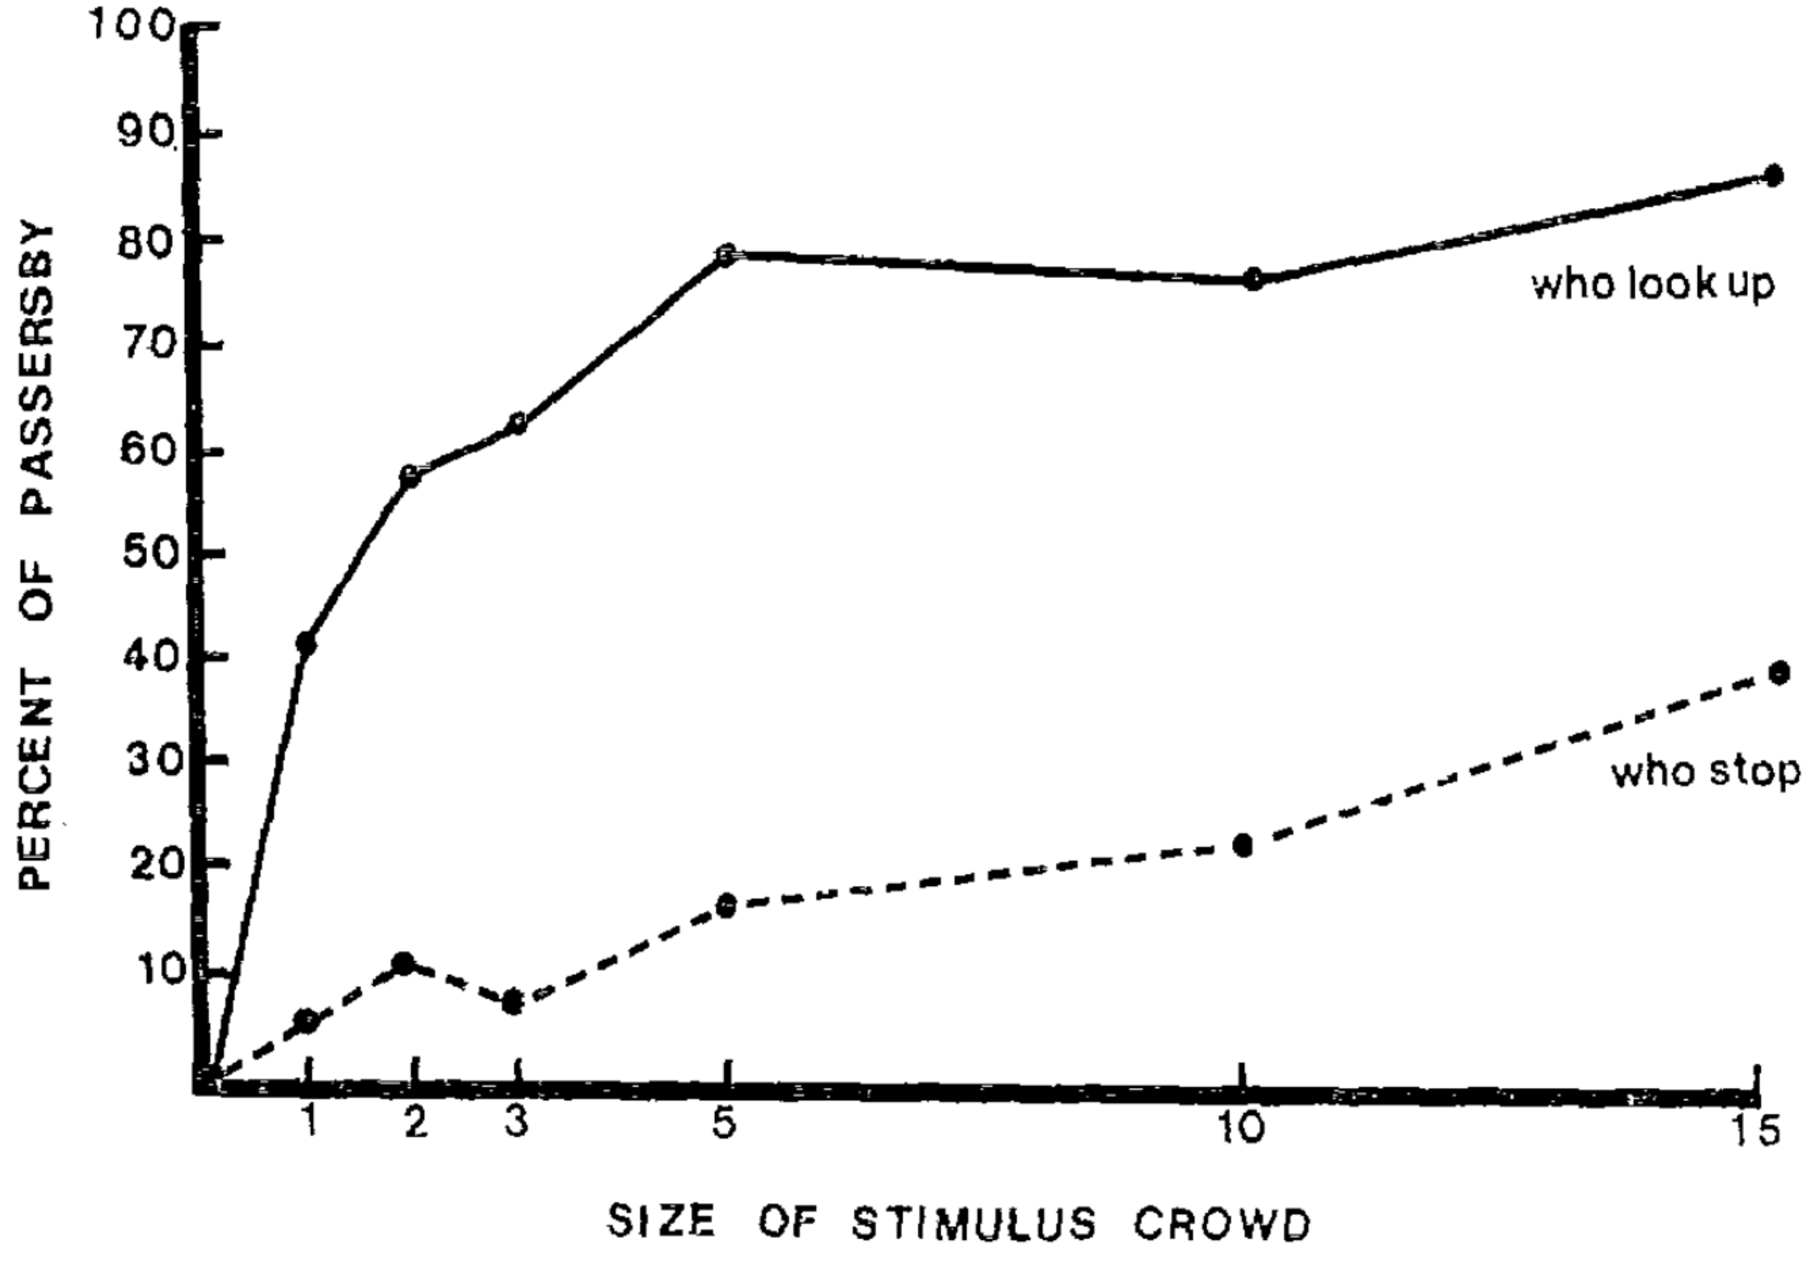

FIG. 1. Mean percentage of passersby who look up and who stop, as a function of the size of the stimulus crowd.

Another example: direct benefits

Joining Facebook
If no one else is on it, useless
But if lots of your friends are on it, helpful
Or fax machines, or WhatsApp, or gaming consoles, etc…


Simple Herding Model

Decision to be made (resto choice, adopt a new technology, support political position, etc)

People decide sequentially, and see all choices of those who acted earlier

Each person has some **private information** that can help guide their decision

People **can't** directly observe what others **know**, but **can** observe what they **do**

Simple Herding Model

Model: n students in a classroom, urn in front

Two urns with marbles:

- "Majority-blue" urn has 2/3 blue, 1/3 red
- "Majority-red" urn has 2/3 red, 1/3 blue
- 50%/50% chance that the urn is majority blue/red

One by one, each student privately gets to look at 1 marble, put it back without showing anyone else, and guess if the urn is Majority-blue or Majority-red

Simple Herding Model

Student 1: Just guess the colour she sees

Student 2:

If same as first person, guess that colour.

But if different from first, then since he knows first guess was what first person saw, then he's indifferent between the two. Guess what he saw

Student 3:

If first 2 are opposite colours, guess what she sees (tiebreaker)

If previous 2 are the same colour (blue) and S3 draws red, then it's like he has drawn three times and gotten two blue, so she should guess majority-blue, **despite her own private information!**

Bayes' Rule

P[A|B] = P[A and B] / P[B]
P[A|B] = P[B|A]P[A] / P[B]
Posterior = Update * Prior

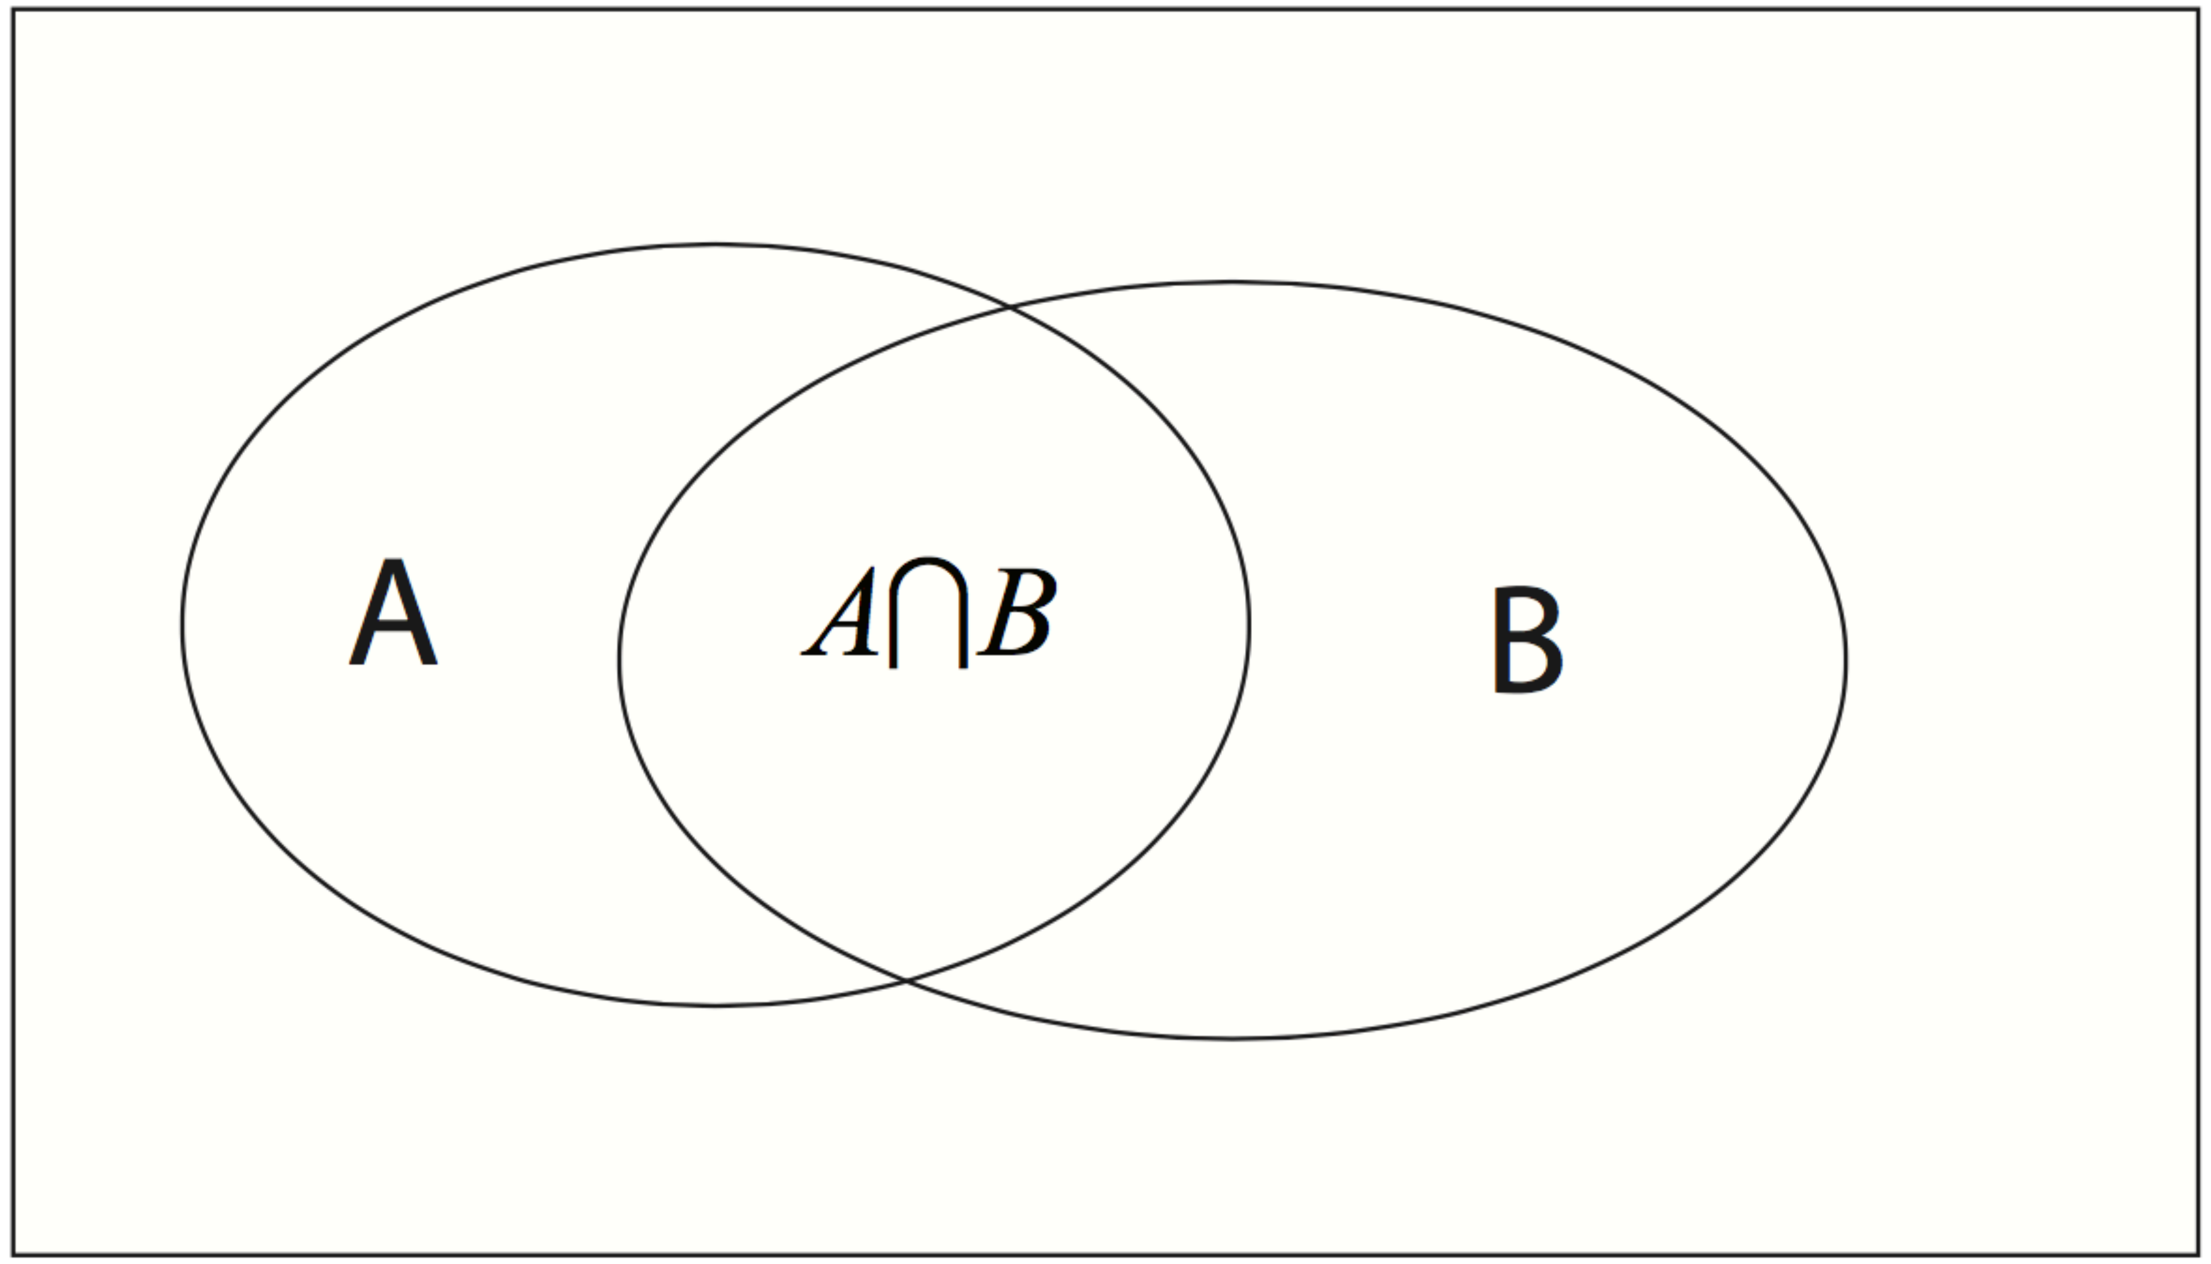

A Student’s Decision

Say you’re one of the students. You go to the urn and pick a marble, say it’s blue.

What should you do?

A Student's Decision

Say you're one of the students. You go to the urn and pick a marble, say it's blue.

What should you do?

Don't just naively guess blue… you've heard a lot of information too! (what if everyone else said red?)

Guess blue if given you what you know **AND** the information you have from others leads you to believe the urn is majority-blue

Simple Herding Model

Student guesses blue if P[majority-blue | what she has seen/heard] > 1/2, red otherwise

Prior: P[majority-blue] = P[majority-red] = 1/2

And because of the marbles in the urns:

P[blue | majority-blue] = P[red | majority-red] = 2/3

Student 1: say she picks blue marble

P[maj-blue | blue] = P[maj-blue]P[blue | maj-blue] / P[blue]

P[blue] = P[blue | maj-blue]P[maj-blue] + P[blue | maj-red]P[maj-red]

= (2/3)(1/2) + (1/3)(1/2) = 1/2

So P[maj-blue | blue] = (1/3)/(1/2) = 2/3

Simple Herding Model

Student 2 same as Student 1 (it's rational to guess what you see), so consider Student 3
Student 3 can reason that first two guesses are what the students actually saw (rationality)
Say she sees different from first two guesses: blue blue red

P[maj-blue | blue blue red]?

= P[maj-blue]P[blue blue red | maj-blue] / P[BBR]
= P[BBR | maj-blue] = (2/3)(2/3)(1/3) = 4/27

P[BBR] = P[BBR|maj-blue]P[maj-blue] + P[BBR|maj-red]P[maj-red]
= (2/3)(2/3)(1/3)(1/2) + (1/3)(1/3)(2/3)(1/2) = 1/9

Plug it all in: 2/3
Student 3 ignores what she sees and goes with what she heard before => **information cascade**

Same for all subsequent students!

Simple Herding Model: Lessons

Cascades can be wrong
Cascades can be based on very little information
Cascades are fragile

Be careful in drawing conclusions from the behaviour of a crowd: we just saw that the crowd can be wrong even if every individual is perfectly rational and takes the same action!

Simple Herding Model: Lessons

erictucker @erictucker · Nov 9

Anti-Trump protestors in Austin today are not as organic as they seem. Here are the busses they came in. #fakeprotests #trump2016 #austin

16K 14K

The Spread of Information

Friends tell their friends stuff

- Rumours/secrets
- Useful information (not homework answers though)
- Beliefs, hopes, desires, fears, …

Social media built to support this:

- Blogs (personal/professional)
- Social networks (Facebook)
- Microblogging (Twitter)

What is the structure of how information spreads?

What does “go viral” mean?

People say stuff goes viral

Person-to-person transmission

Deep branching structures

Hypothesis: an idea, story, joke, etc. spreads like a **virus**, “infecting” minds like viruses infect the body

This implies a certain kind of structure!

What does “go viral” mean?

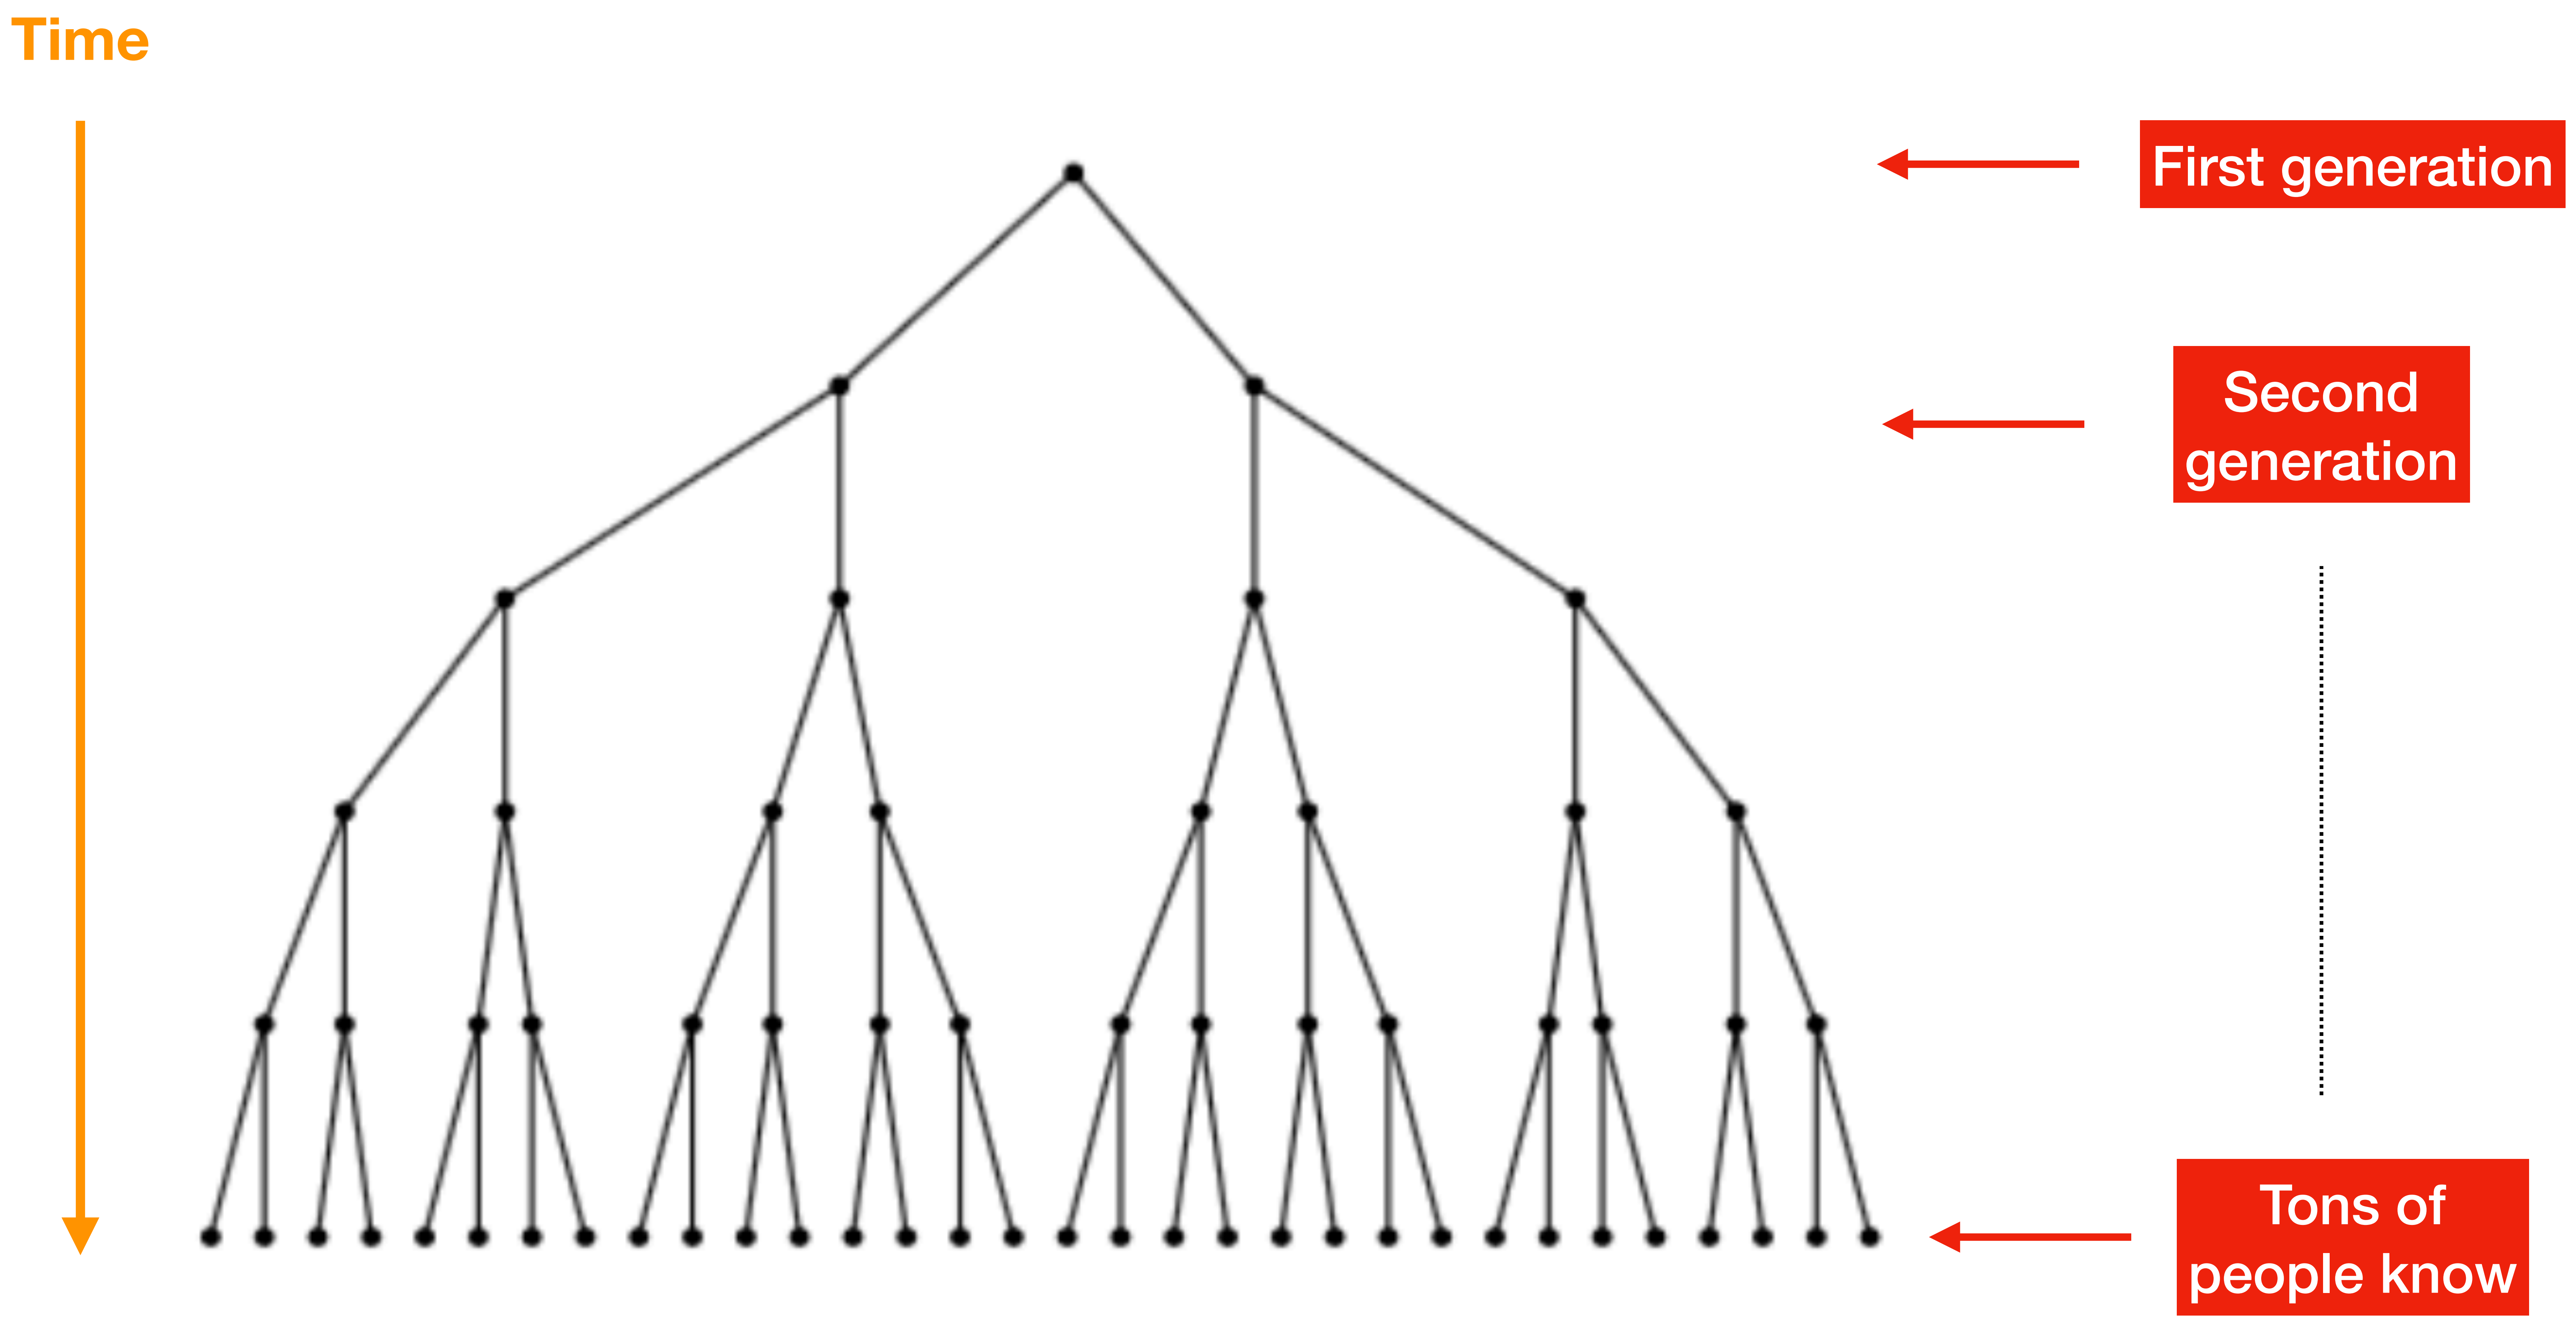

But another way

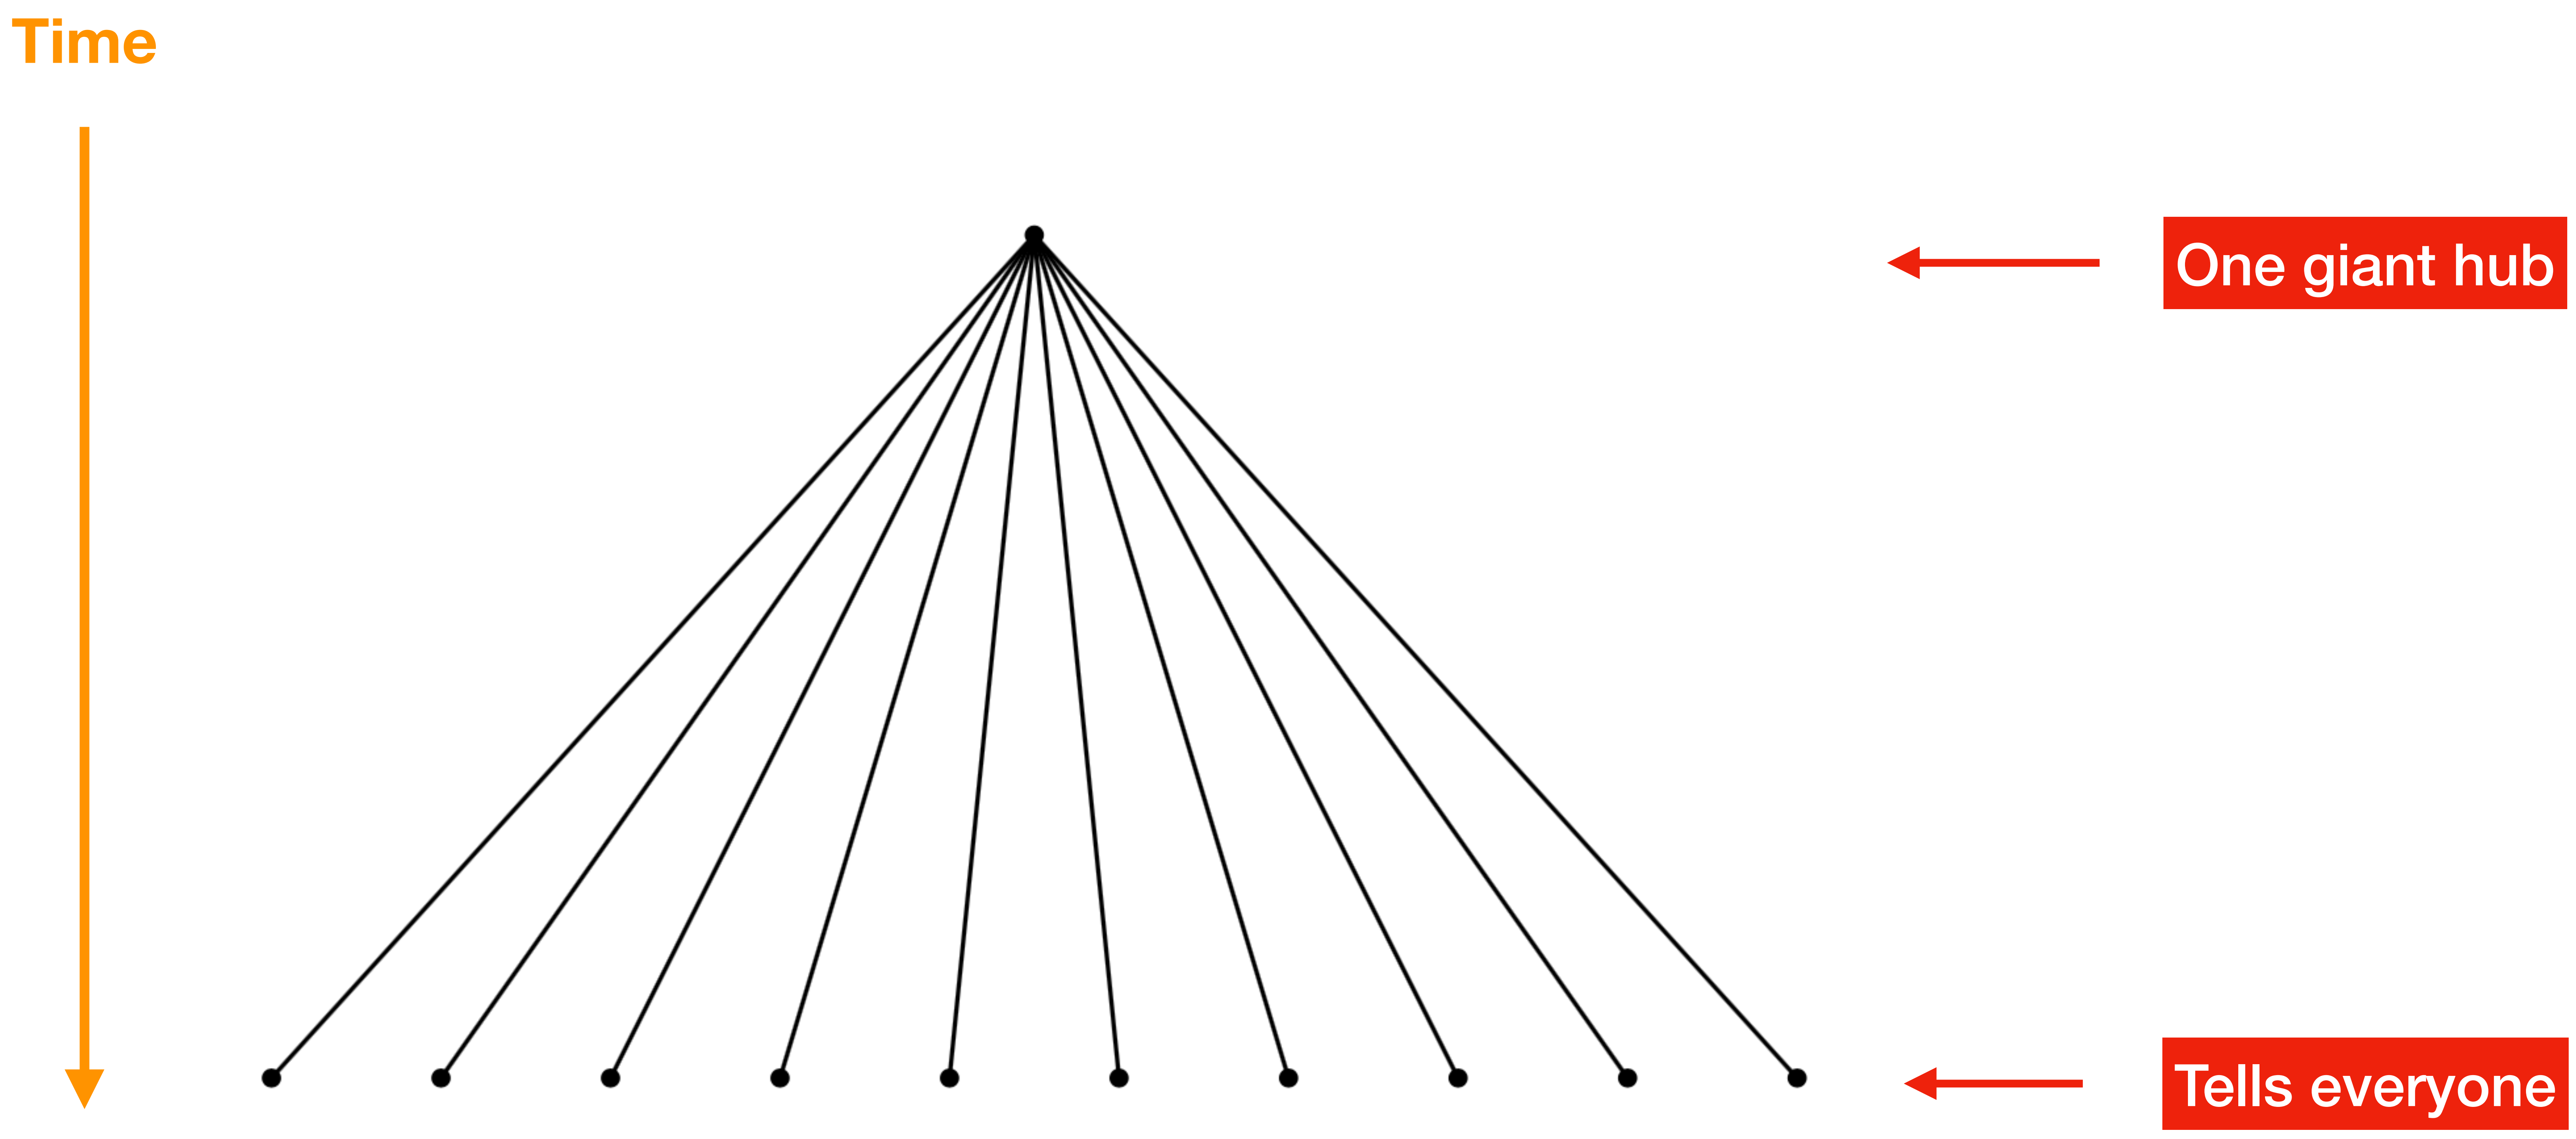

Which is it?

Big media (CNN, BBC, NYT, Fox)
Celebrities (Biebs, Taylor Swift)

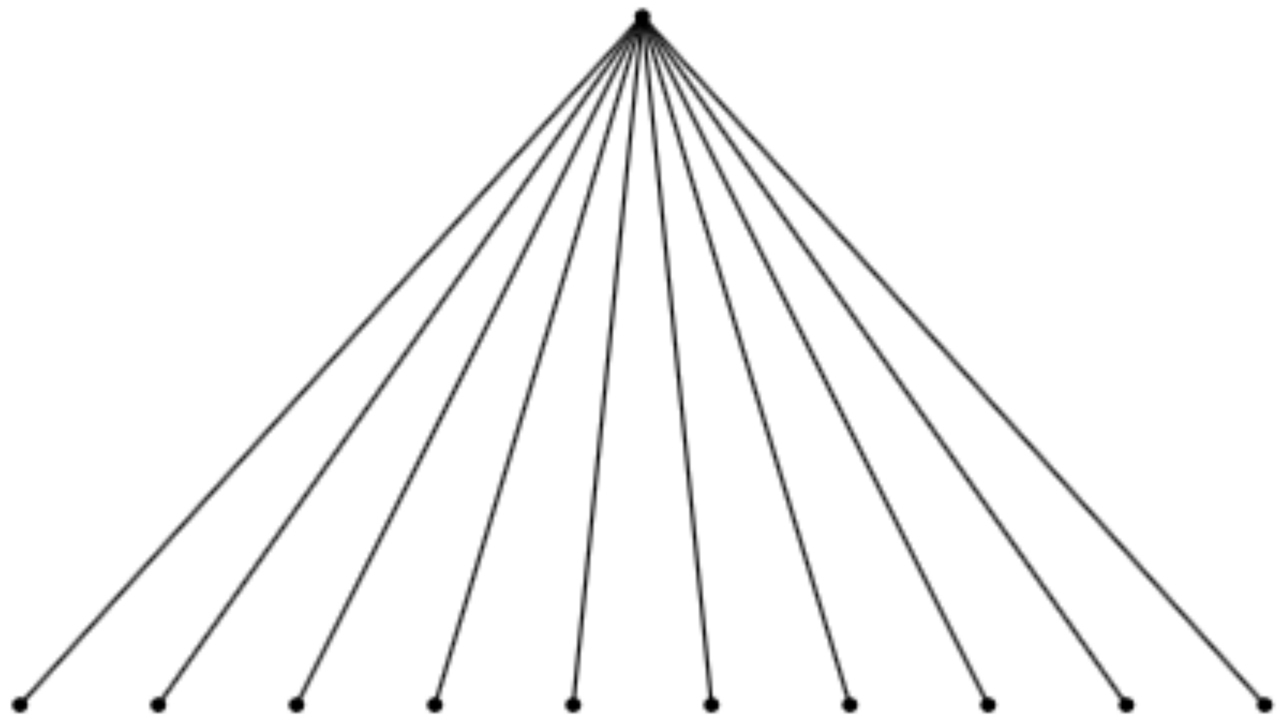

or

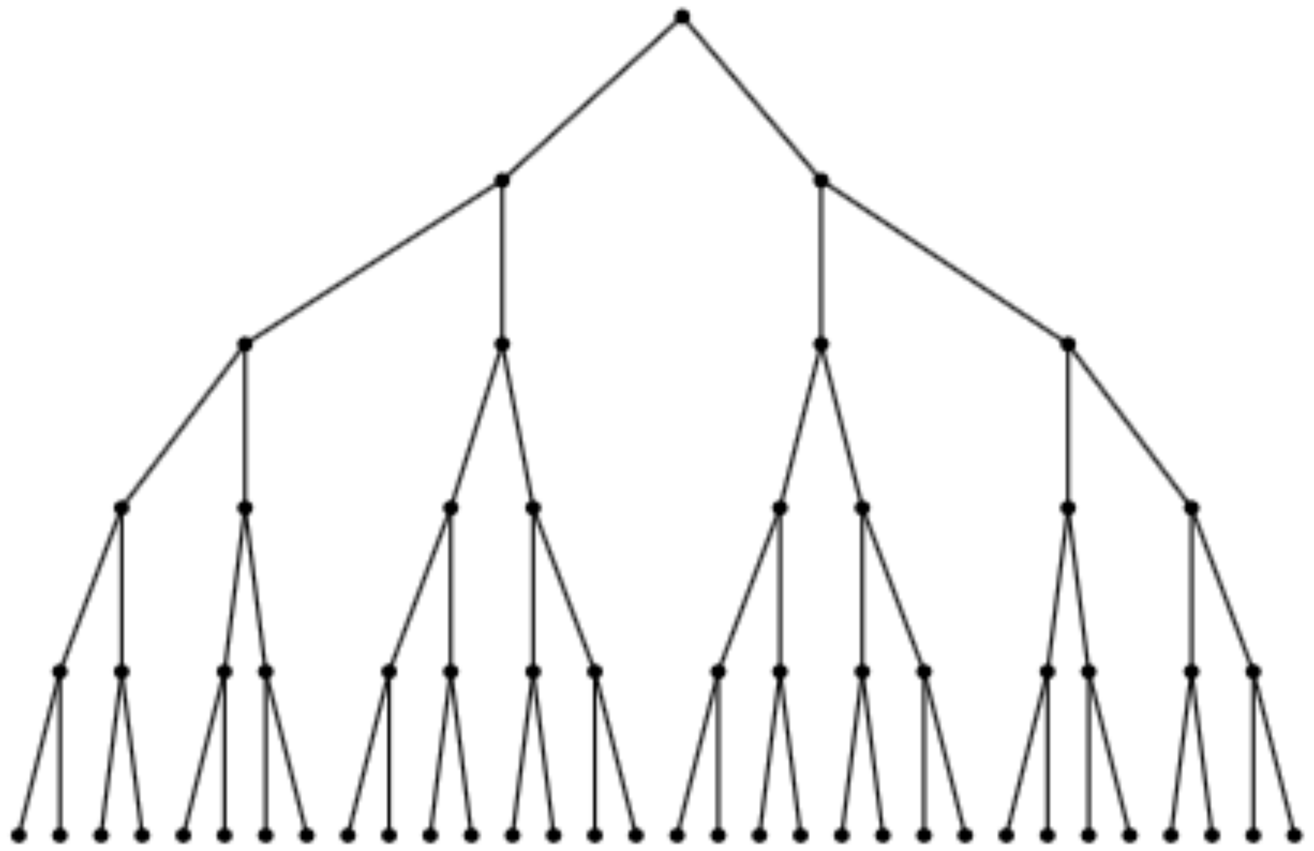

"Broadcast"

"Viral"

- **Organically spreading content**
- **Chain letters**

How to study information spread?

Hard to track "information" spreading from one mind to another

Online proxy: people sharing URLs

Twitter: person A tweets a URL, then a friend B tweets it (or directly retweets)
We say the URL passed from A to B

How to study information spread?

Connect these sharing edges into **trees**

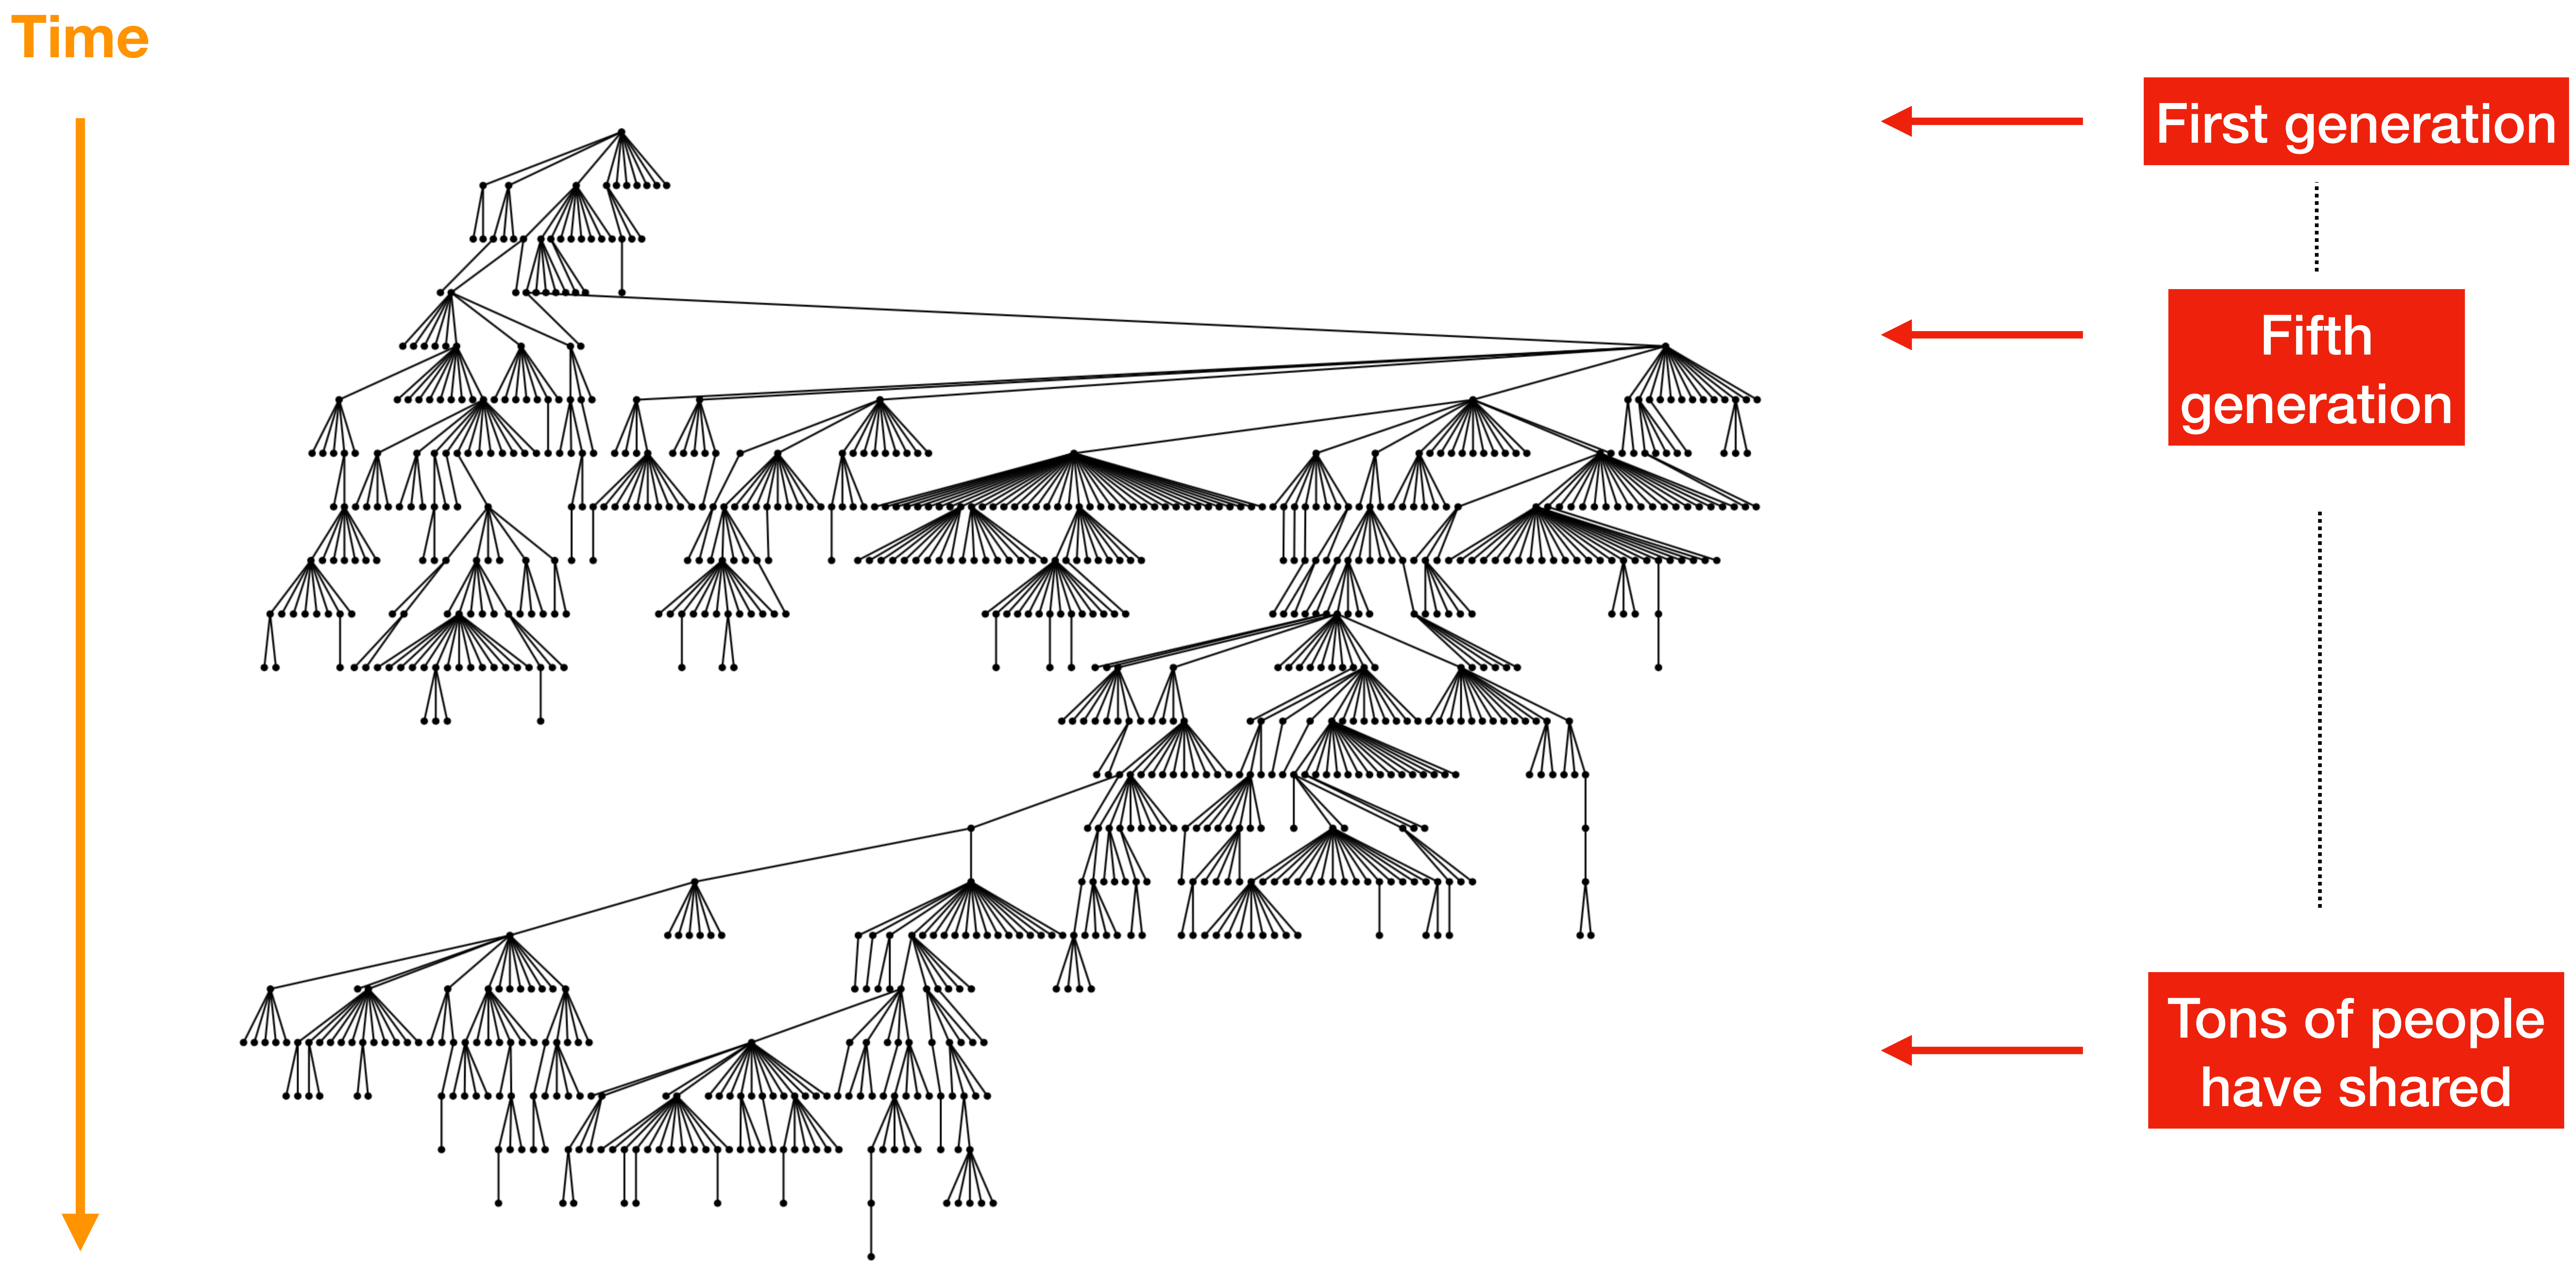

How to measure virality?

How **structurally viral** is a particular cascade?

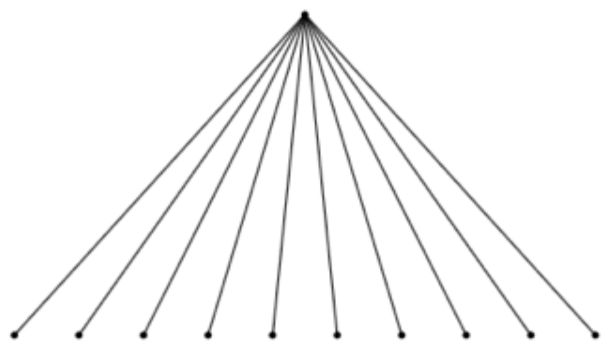

Not viral

?

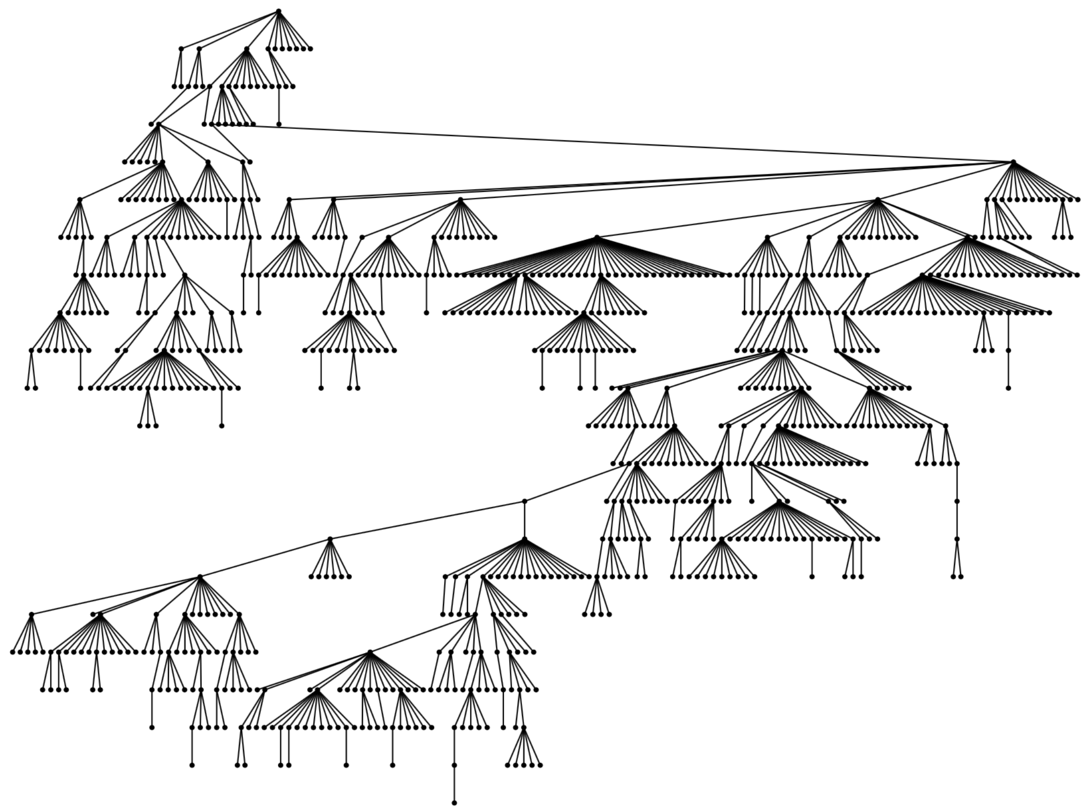

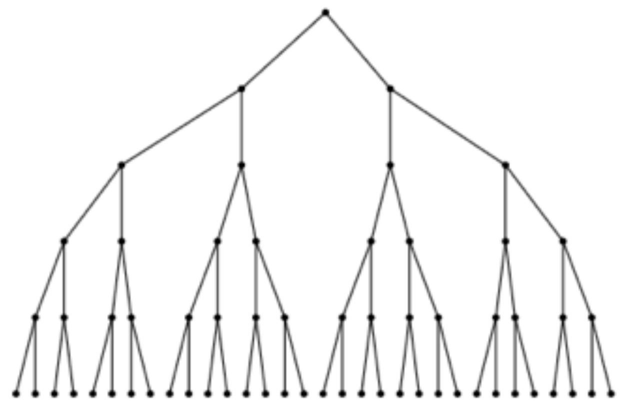

Super viral

How to measure virality?

One idea: **depth of the cascade**

But this is **sensitive to a single long chain**

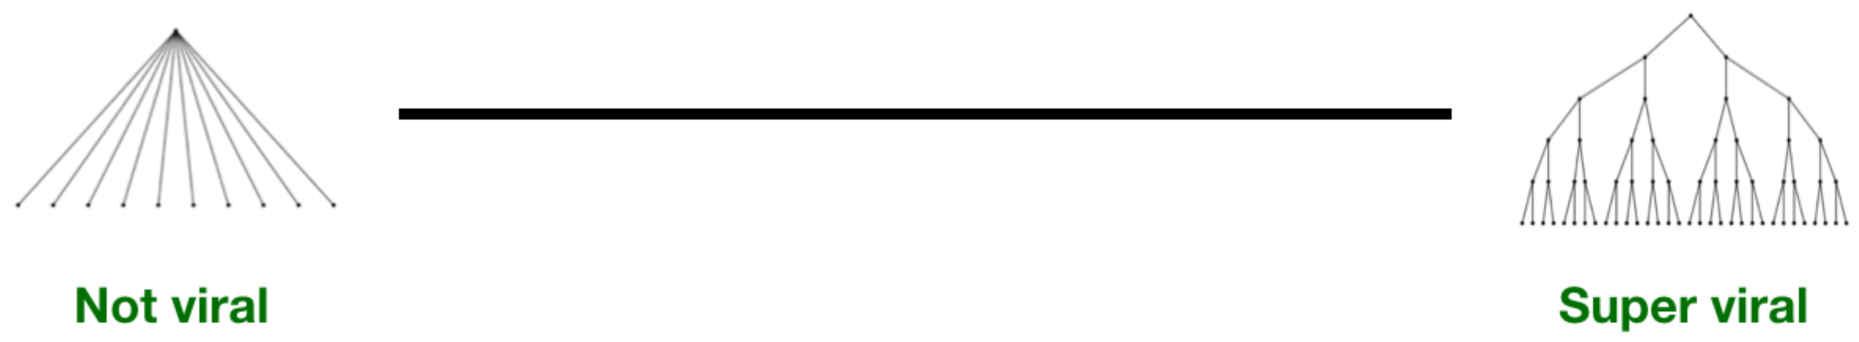

How to measure virality?

Another idea: **average depth of the cascade**

But even this **sometimes fails**: long chain then a big broadcast

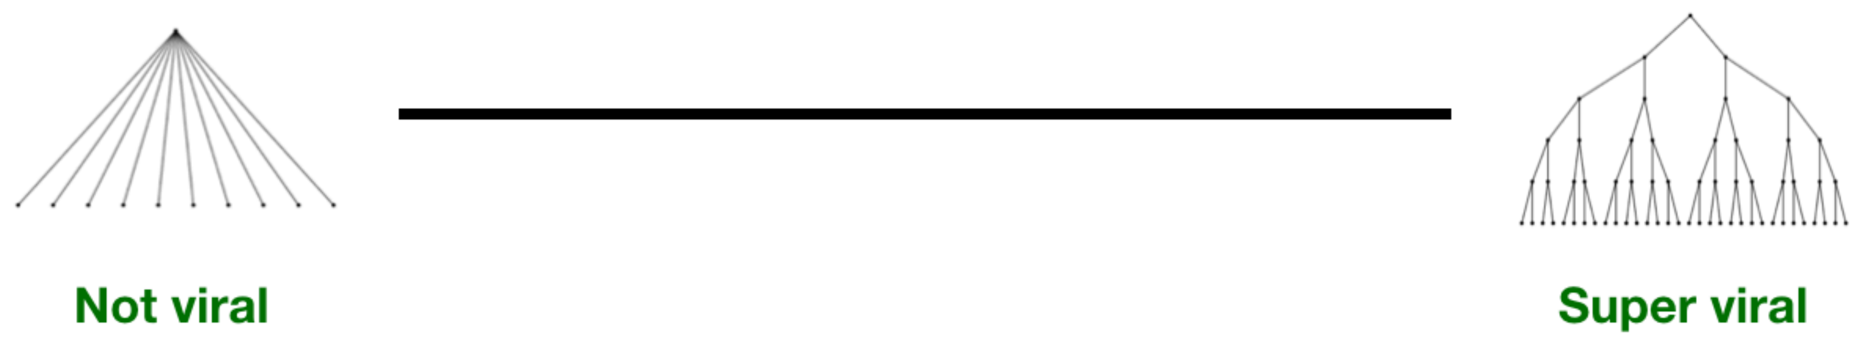

How to measure virality?

Solution: **average path length between nodes**

$$\nu(T) = \frac{1}{n(n-1)} \sum_{i=1}^{n} \sum_{j=1}^{n} d_{ij}$$

Simple average!

Originally studied in mathematical chemistry [Wiener 1947] => “Wiener index”

Not viral **Super viral**

Measure virality in data!

Now we have a way to construct information cascades on Twitter

And for each cascade we can compute a number that determines how "structurally viral" it is

So **how often does stuff go viral?**

Not viral

$$\nu(T) = \frac{1}{n(n-1)} \sum_{i=1}^{n} \sum_{j=1}^{n} d_{ij}$$

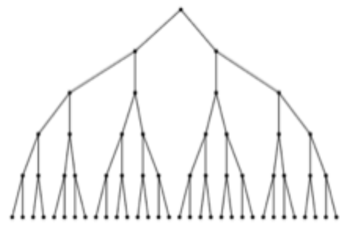

Super viral

Measure virality in data!

Looked at an **entire year of Twitter data**

622 million unique URLs, 1.2 billion “adoptions” (tweets) of these URLs

Every URL is associated with a forest of trees

Not viral

$$\nu(T) = \frac{1}{n(n-1)}\sum_{i=1}^{n}\sum_{j=1}^{n} d_{ij}$$

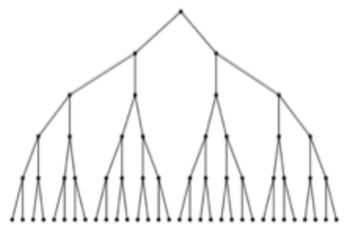

Super viral

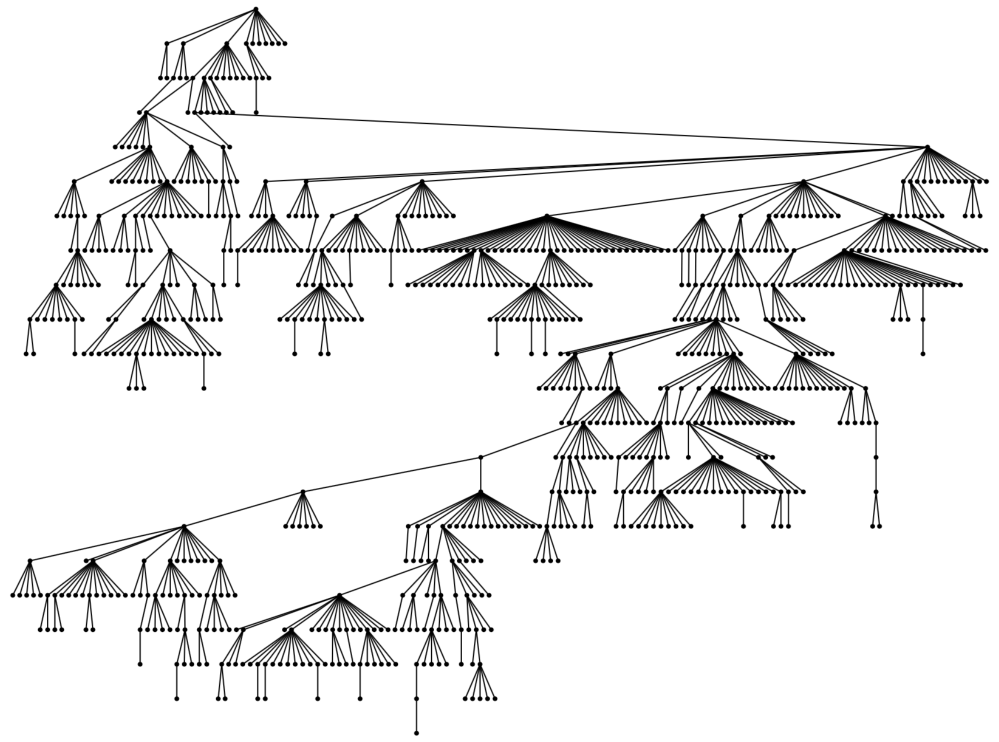

Measure virality in data!

First conclusion: most stuff goes nowhere

Average cascade size: 1.3

Not very interesting cascades: focus on trees of size at least 100 (empirically 1/4000)

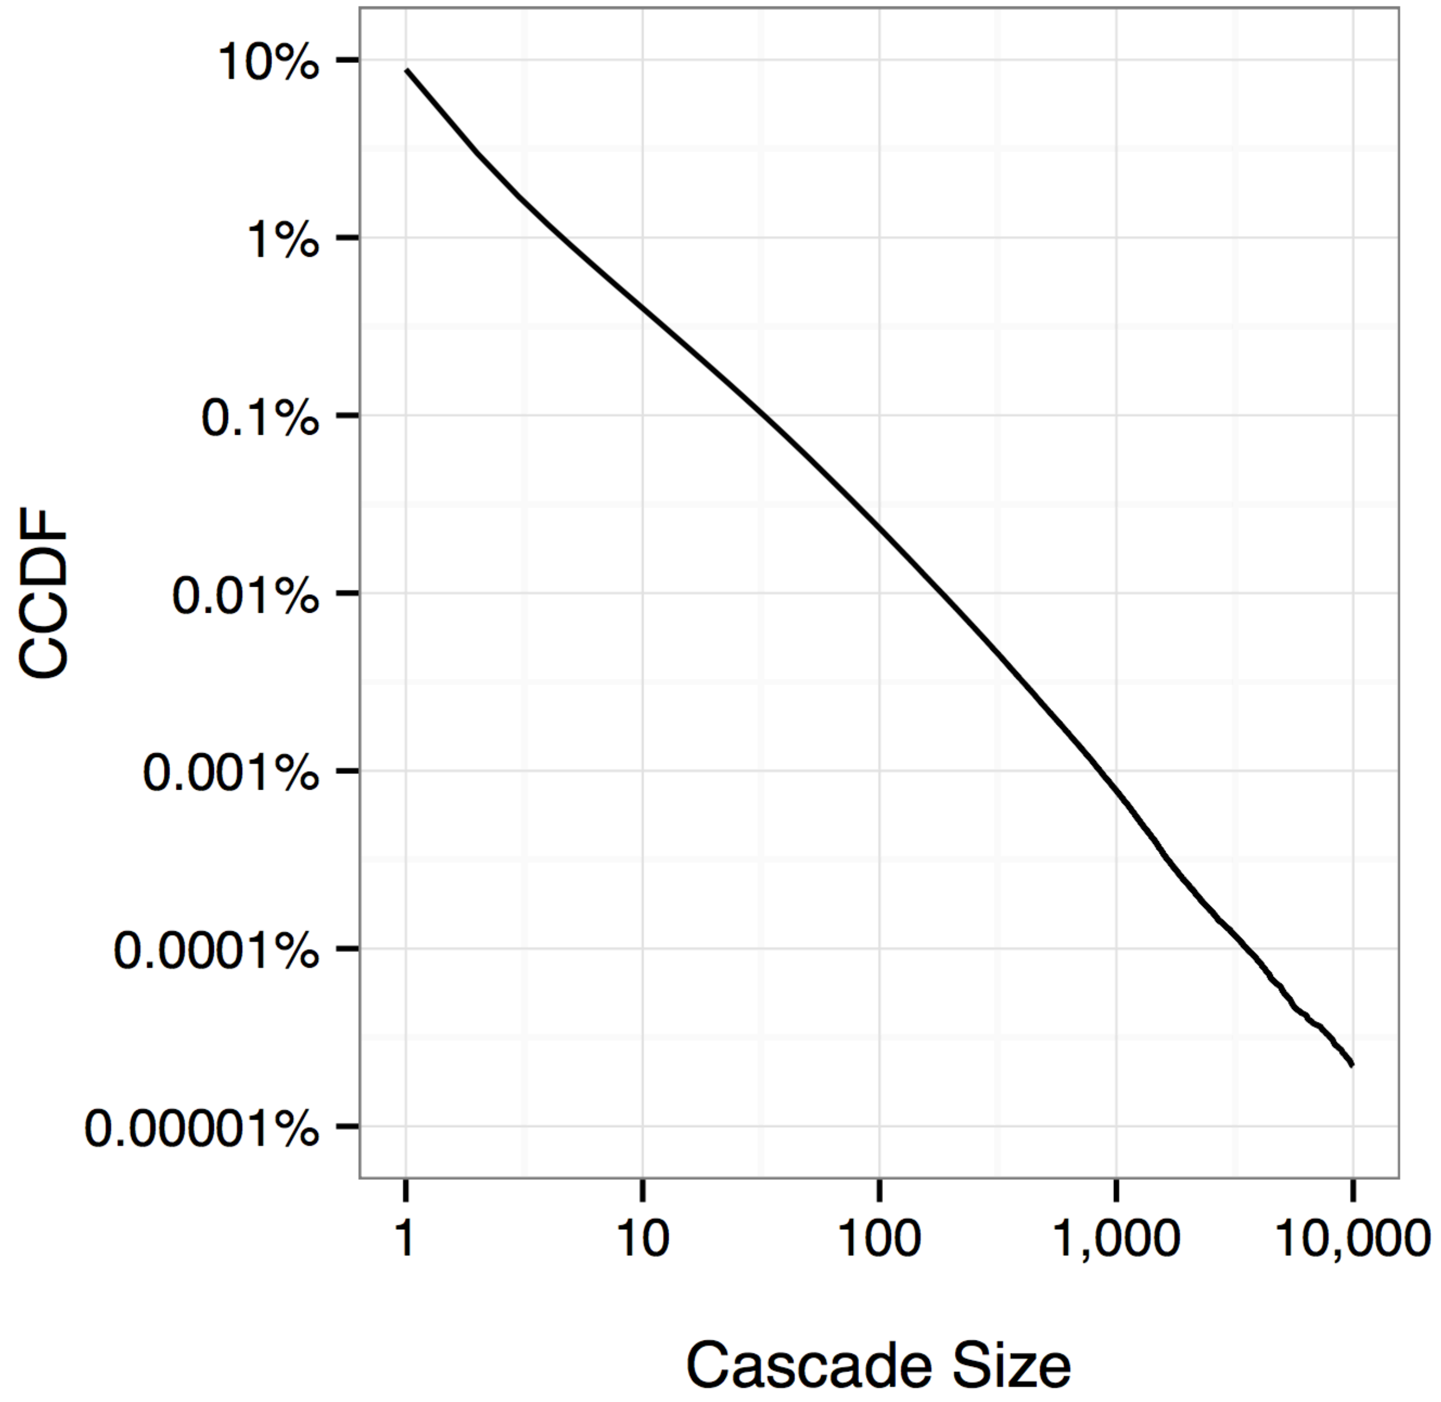

Power law!

Surprising diversity

Broadcast

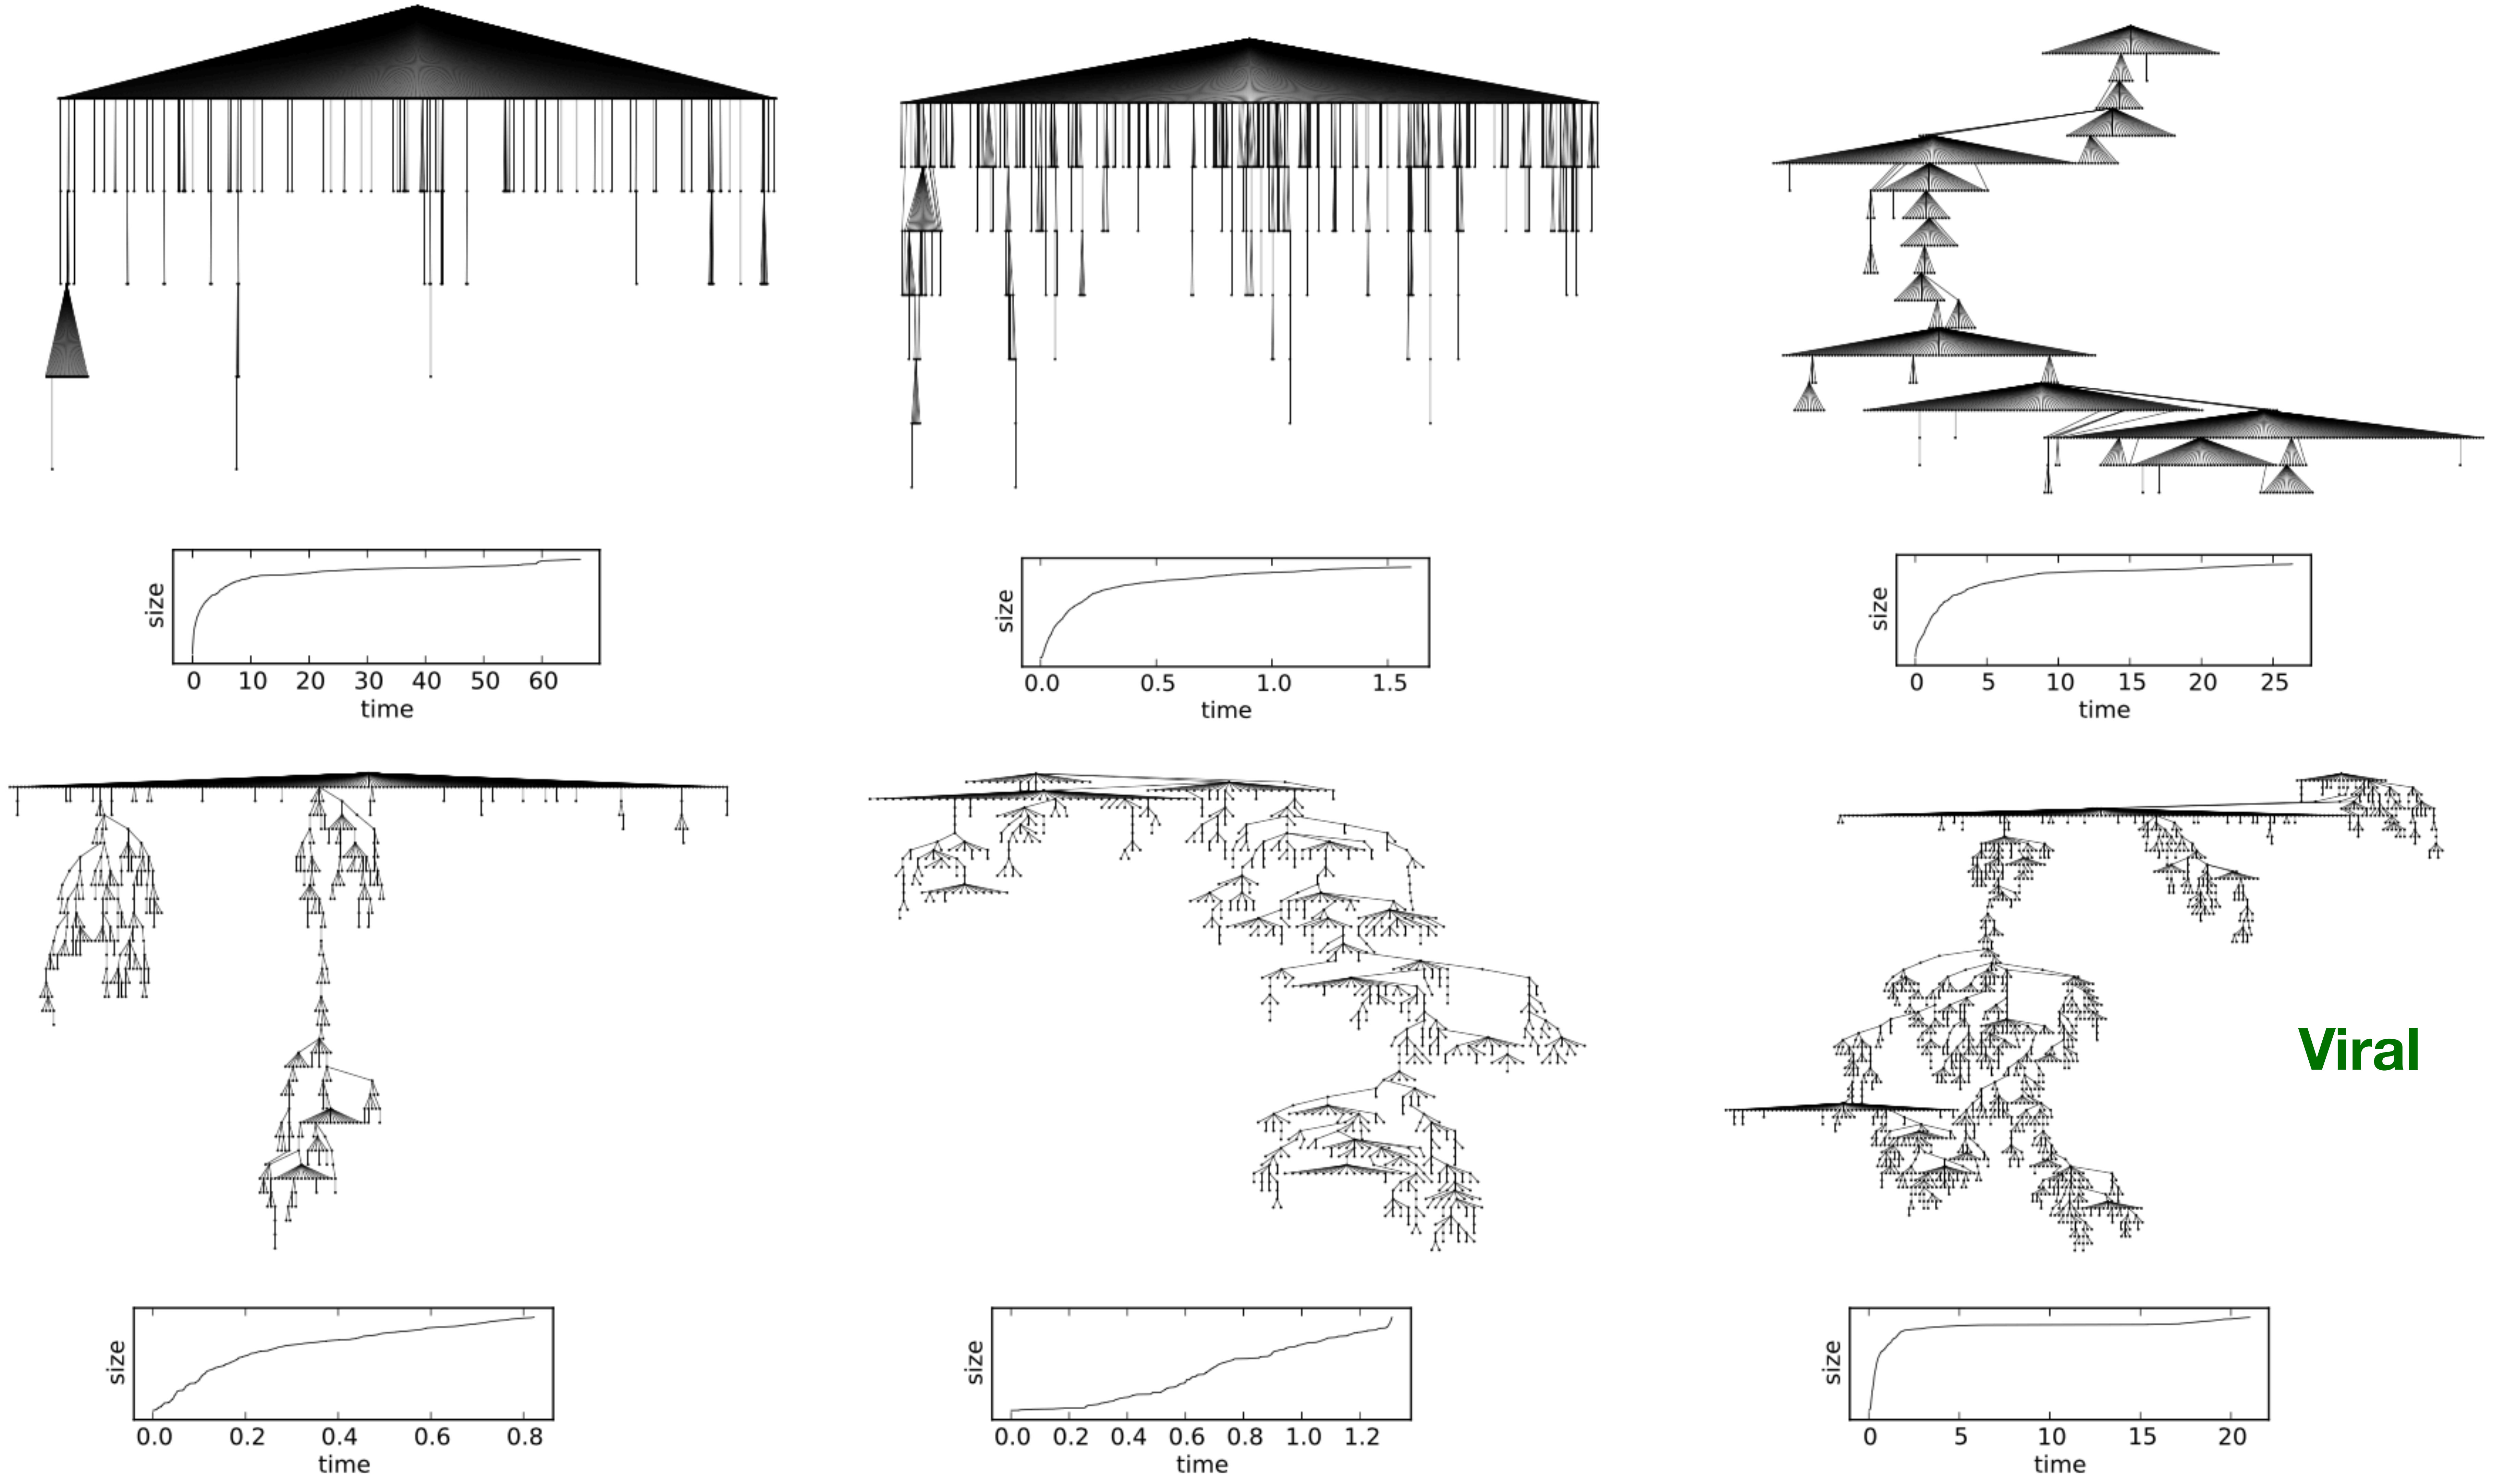

Viral

Surprising diversity

Even big trees **aren't structurally viral** either

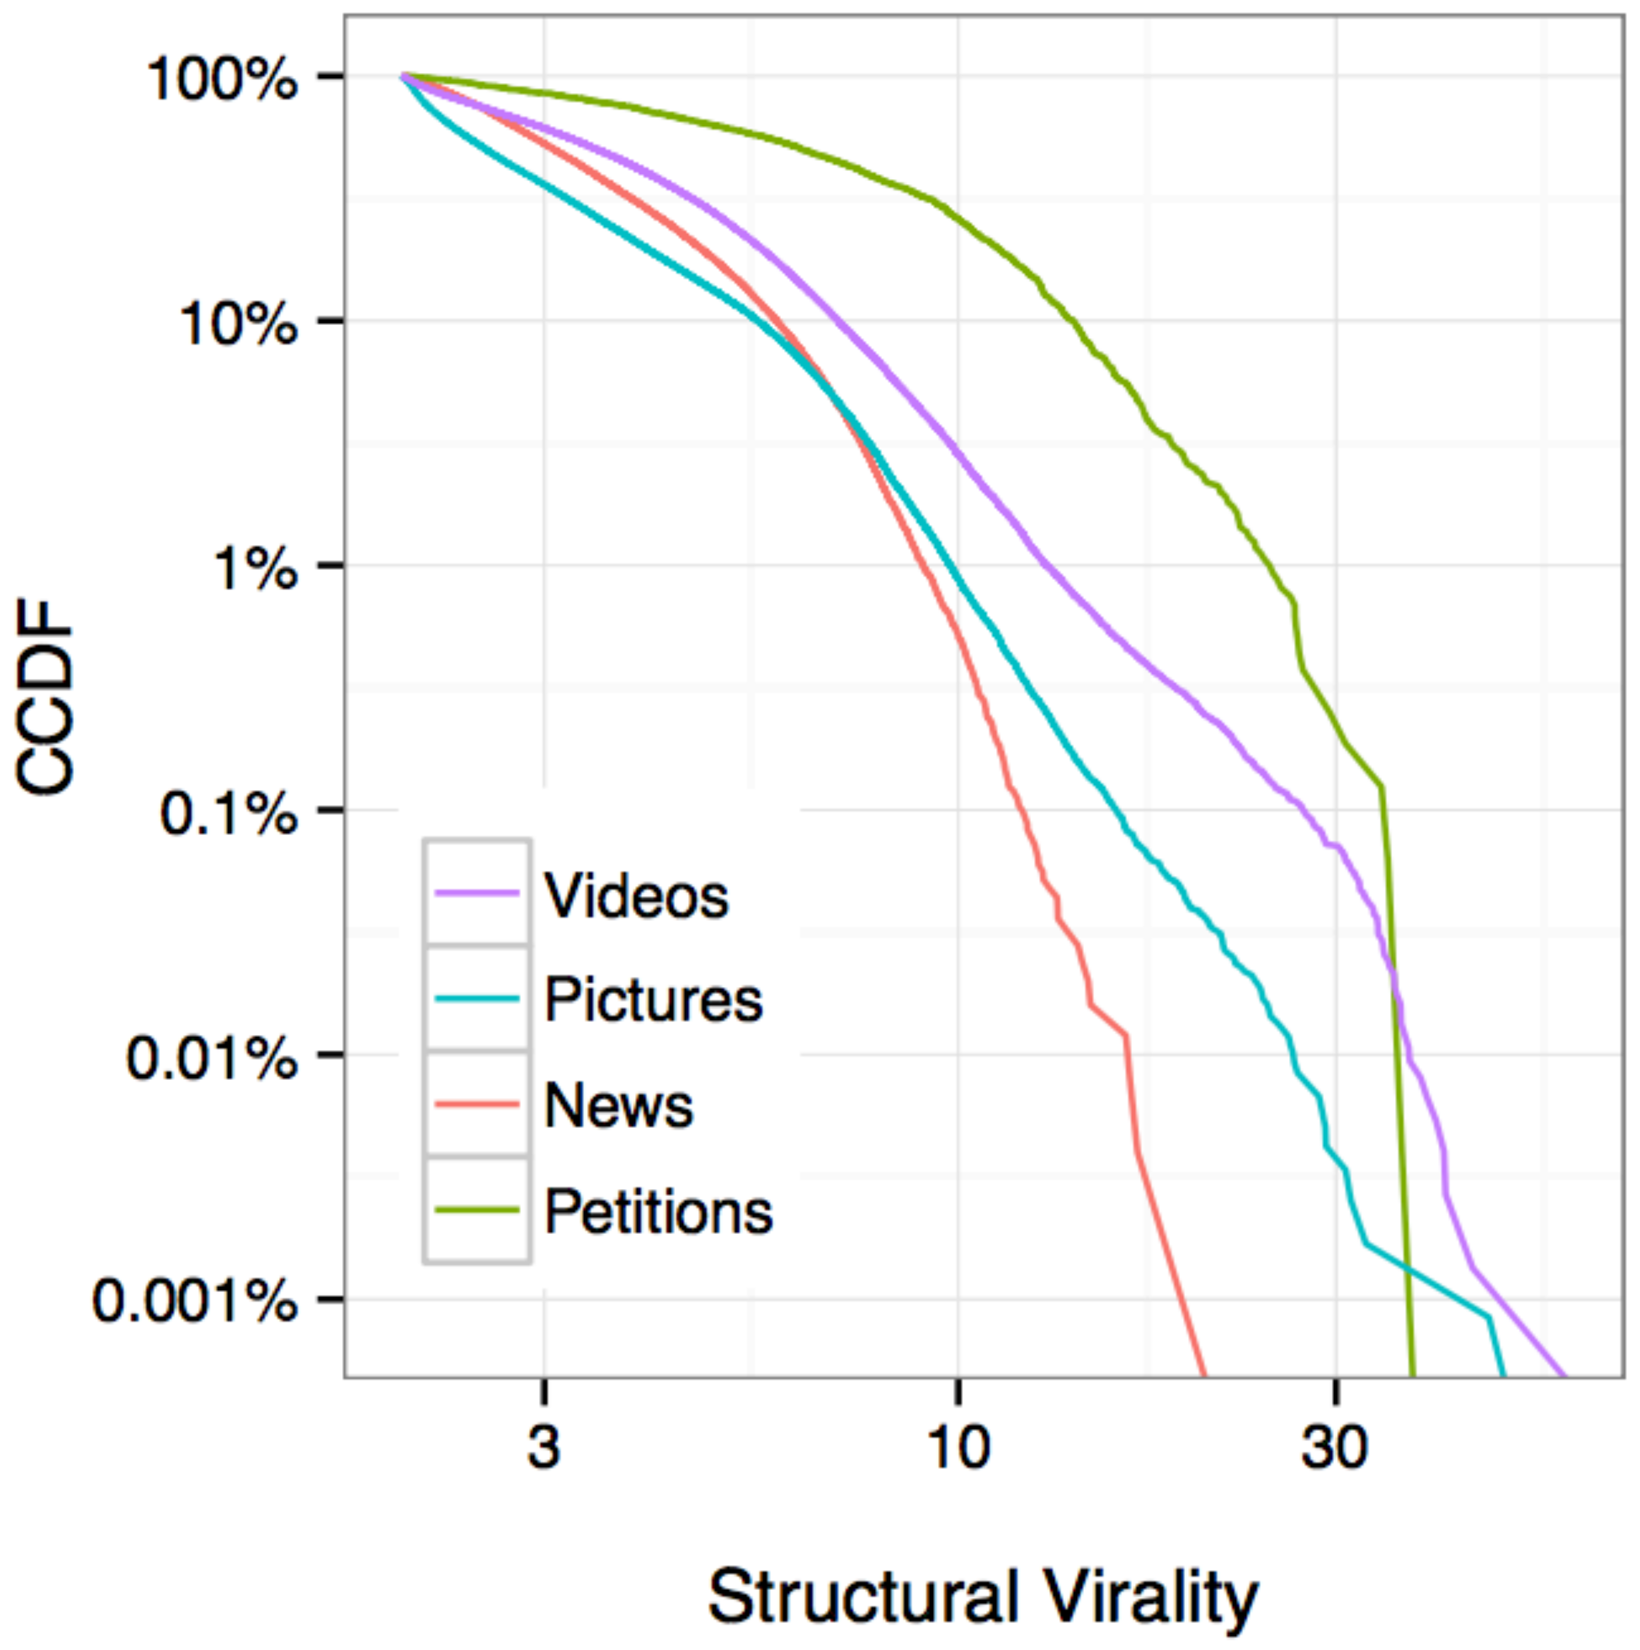

Surprising diversity at every scale

Across domains and across sizes, we see **lots of different types of structures** from broadcast to viral

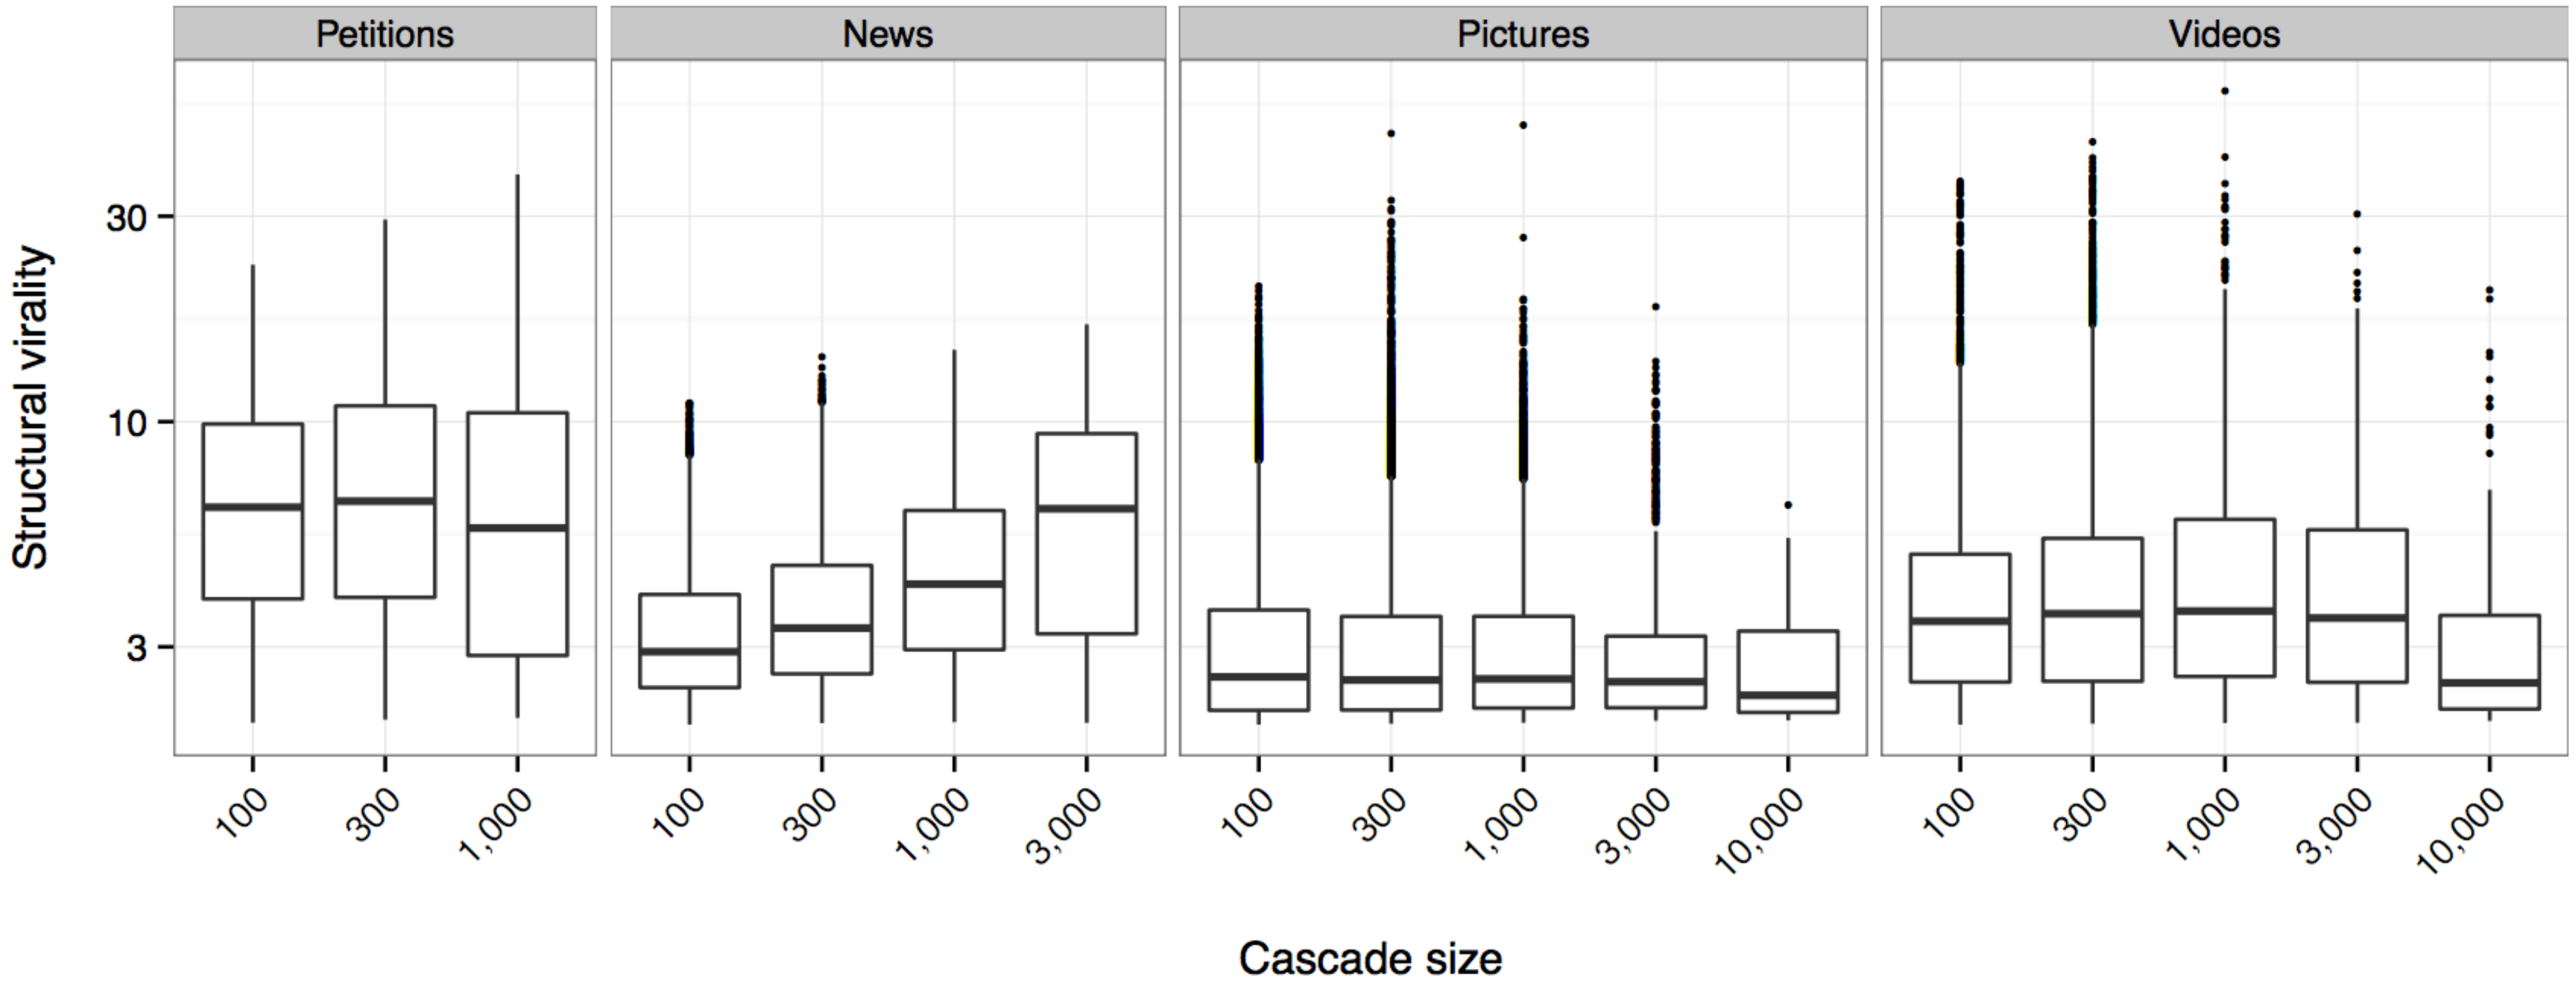

Surprising diversity at every scale

Across domains and across sizes, we see **lots of different types of structures** from broadcast to viral

Very low correlation between size and virality!

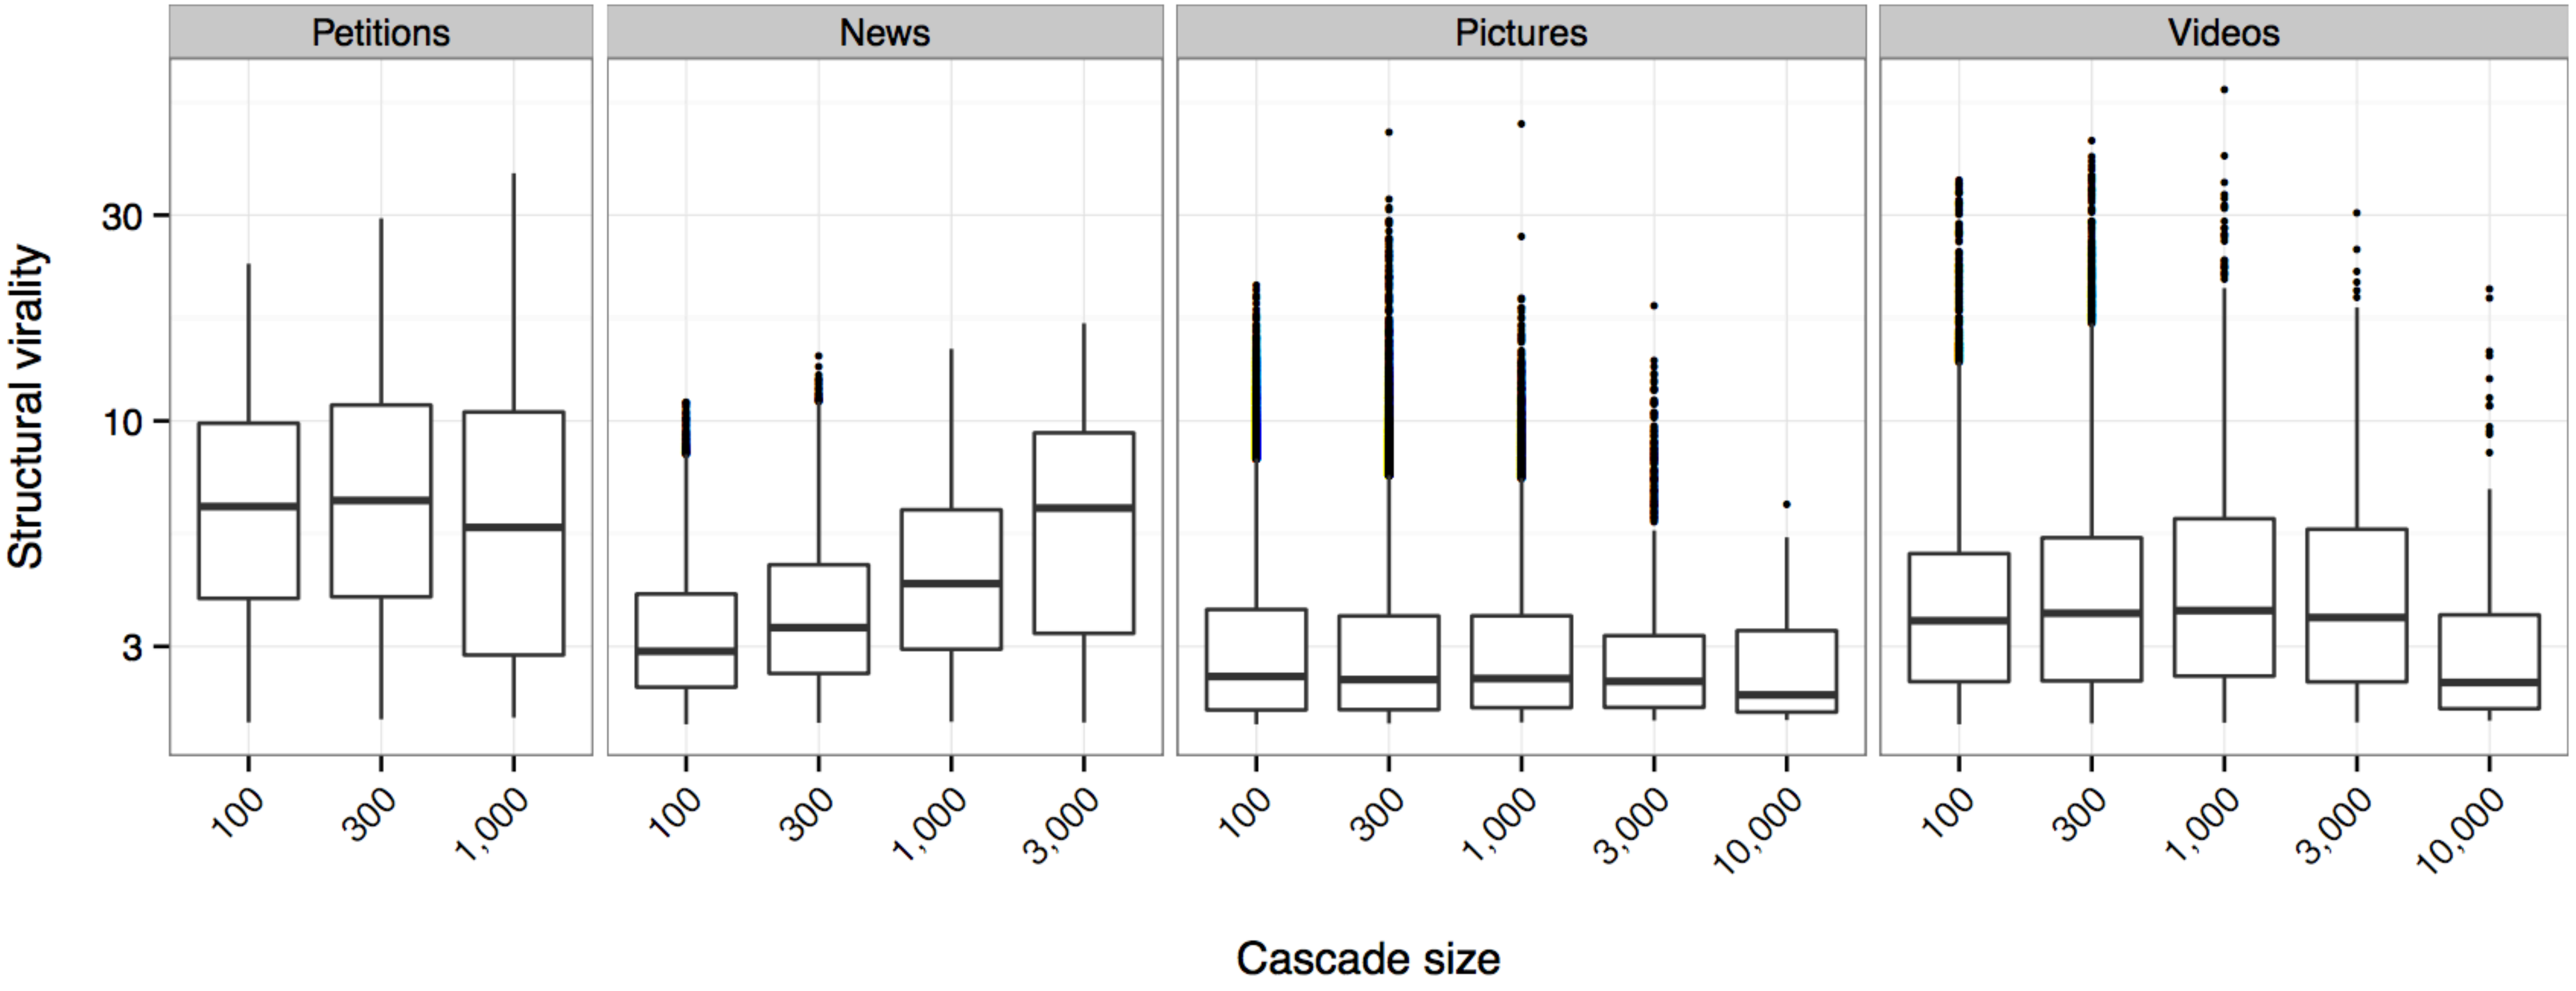

Today: Game Theory in the Wild and Influence Through Networks

- If people are connected through a network, it's possible for them to influence each other's behaviour and actions
- Today: why?
  - Informational
  - Direct benefit
  - Social conformity