Network Effects and Cascading Behavior

CS224W: Machine Learning with Graphs

Jure Leskovec, Stanford University

http://cs224w.stanford.edu

Spreading Through Networks

- **Spreading through networks:**
  - Cascading behavior
  - Diffusion of innovations
  - Network effects
  - Epidemics
- **Behaviors that cascade from node to node like an epidemic**
- **Examples:**
  - **Biological:**
    - Diseases via contagion
  - **Technological:**
    - Cascading failures
    - Spread of information
  - **Social:**
    - Rumors, news, new technology
    - Viral marketing

Information Diffusion: Media

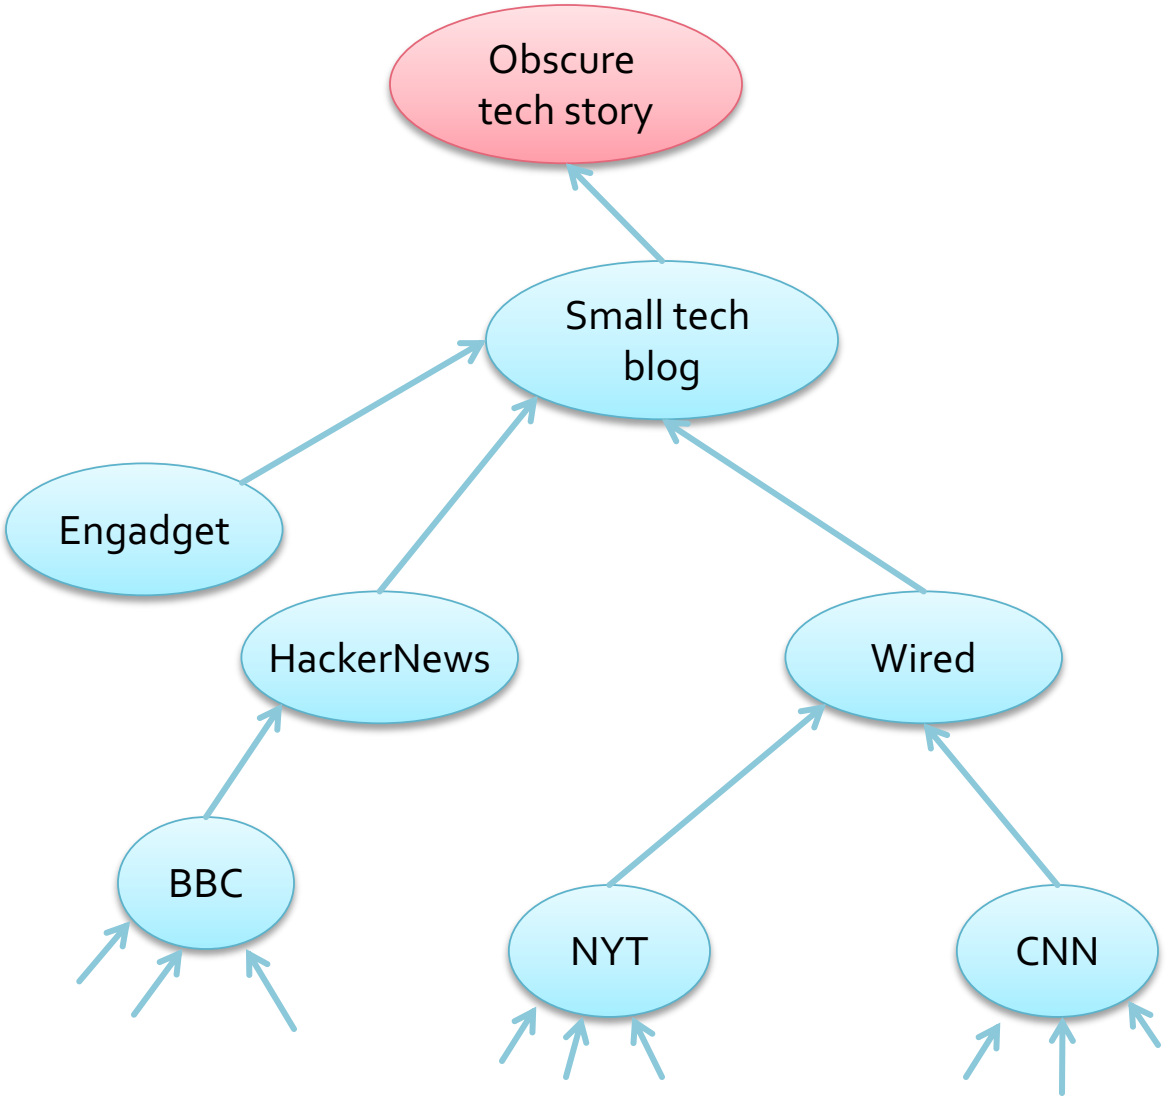

Twitter & Facebook post sharing

Timeline Photos

Back to Album · I fucking love science's Photos · I fucking love science's Page

Previous · Next

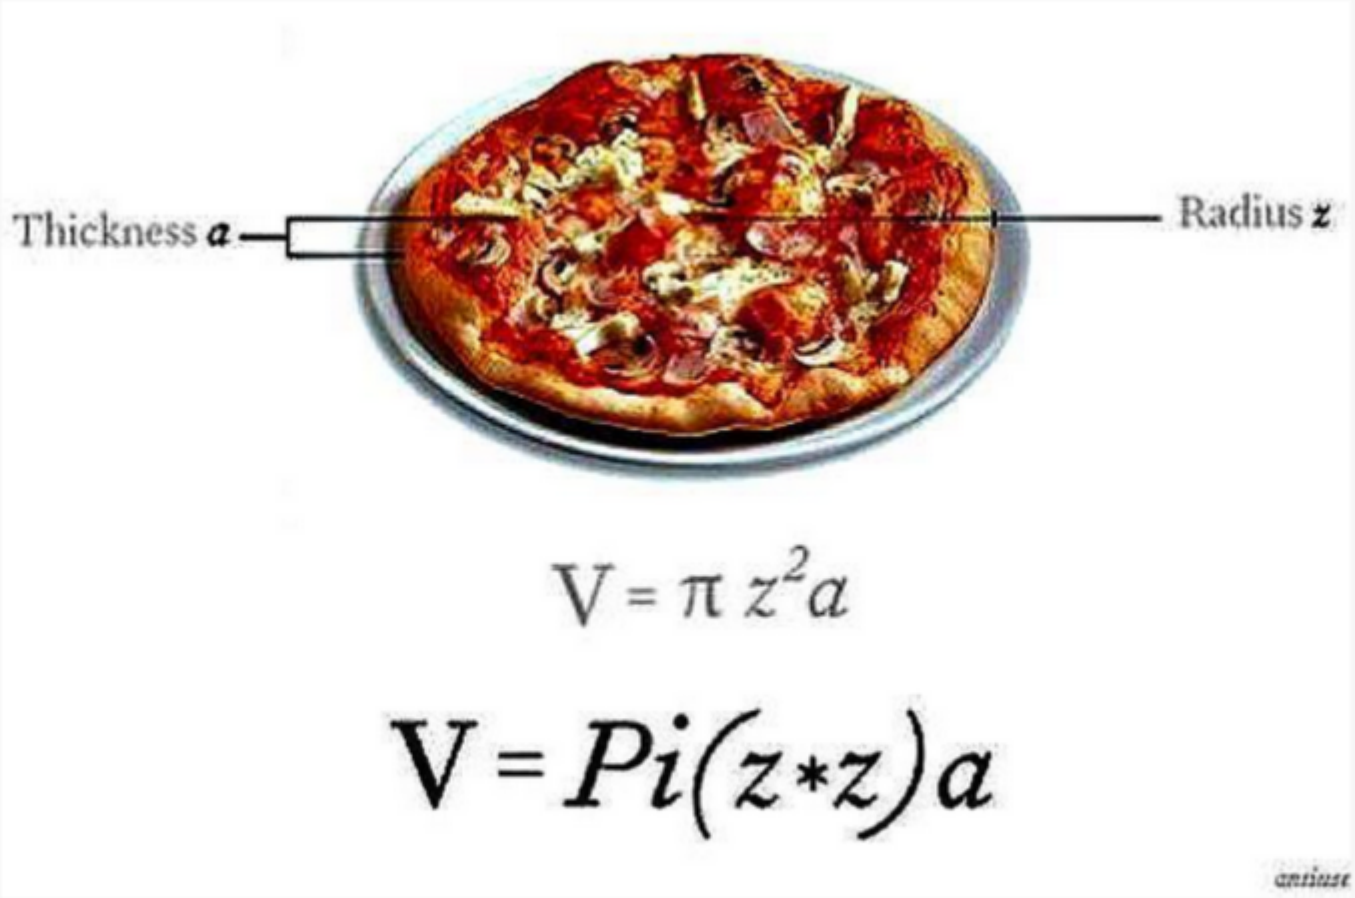

I fucking love science

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

Diffusion in Viral Marketing

- **Product adoption:**
  - **Senders and followers of recommendations**

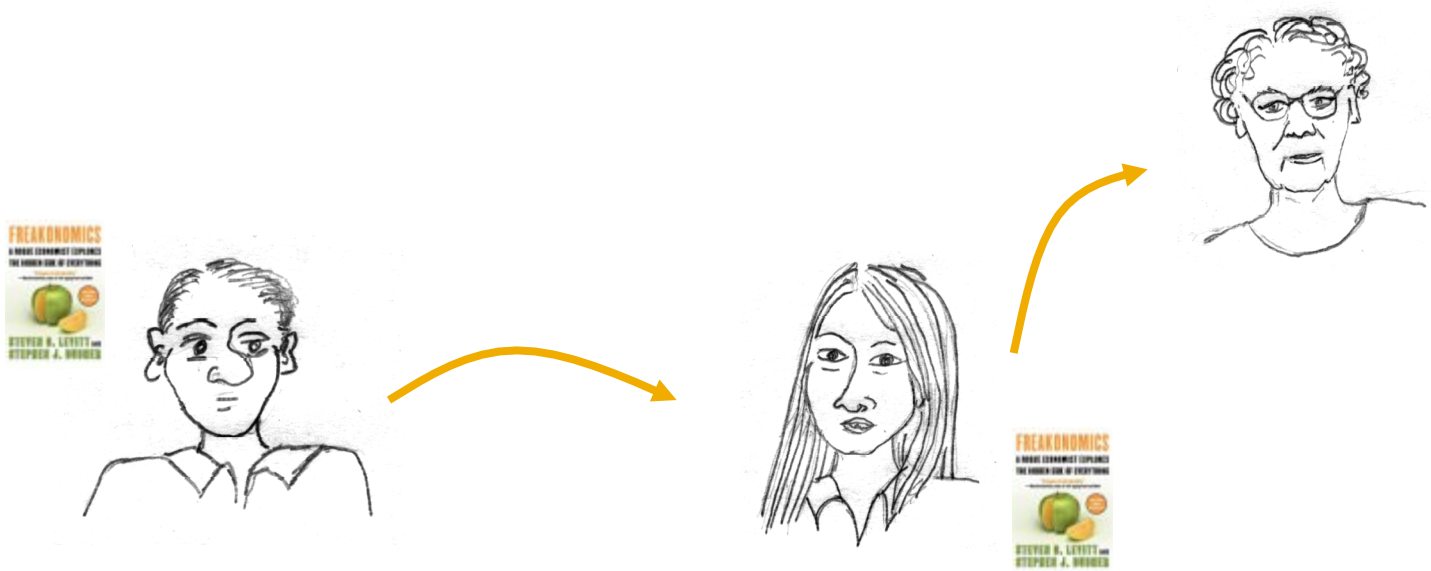

Spread of Diseases (e.g., Ebola)

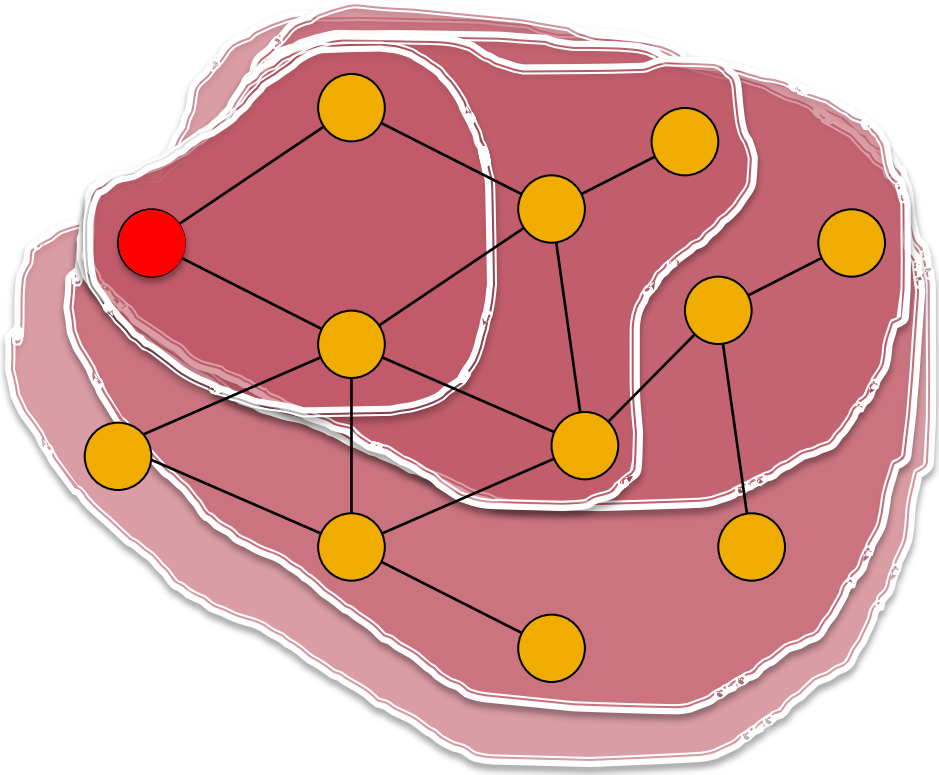

Network Cascades

- **Contagion that spreads over the edges of the network**
- It creates a propagation tree, i.e., **cascade**

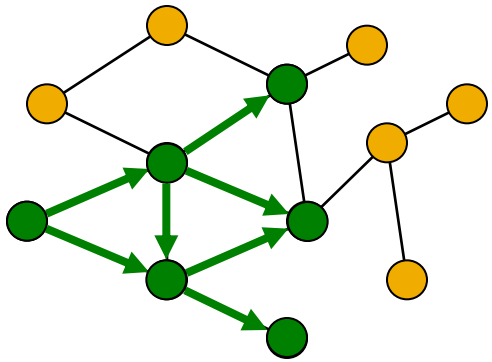

Network

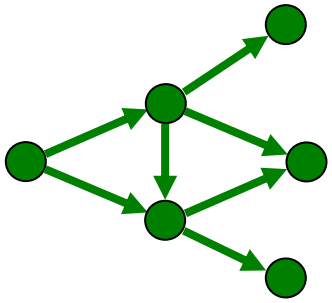

Cascade
(propagation tree)

Terminology:

- What spreads: Contagion
- "Infection" event: Adoption, infection, activation
- Main players: Infected/active nodes, adopters

How Do We Model Diffusion?

- **Decision based models (today!):**
  - Models of product adoption, decision making
    - A node observes decisions of its neighbors and makes its own decision
  - **Example:**
    - You join demonstrations if *k* of your friends do so too
- **Probabilistic models (on Tuesday):**
  - **Models of influence or disease spreading**
    - An infected node tries to "push" the contagion to an uninfected node
  - **Example:**
    - You "catch" a disease with some prob. from each active neighbor in the network

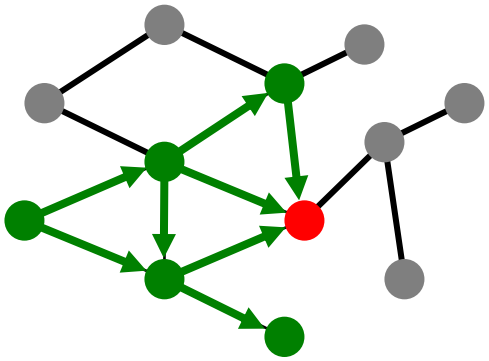

Decision Based Model of Diffusion

[Morris 2000]

Game Theoretic Model of Cascades

- **Based on 2 player coordination game**
  - 2 players – each chooses technology A or B
  - Each player can only adopt **one** "behavior", **A** or **B**
  - **Intuition**: you (node v) gain more payoff if your friends have adopted the **same** behavior as you

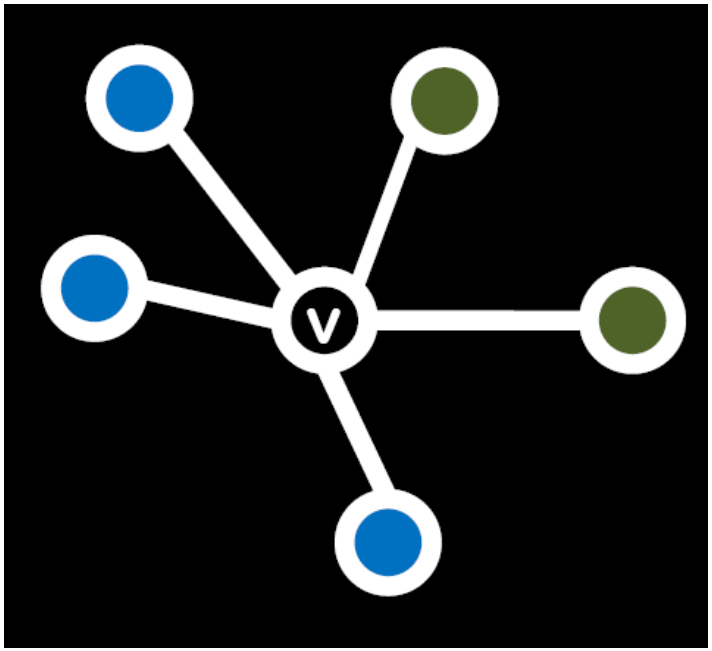

Local view of the network of node v

Example: VHS vs. BetaMax

Example: BlueRay vs. HD DVD

Blu-ray Disc

HD DVD

The Model for Two Nodes

- ***Payoff matrix:***
  - If both **v** and **w** adopt behavior **A**, they each get payoff ***a > 0***
  - If ***v*** and ***w*** adopt behavior ***B***, they reach get payoff ***b > 0***
  - If ***v*** and ***w*** adopt the opposite behaviors, they each get ***0***

- **In some large network:**
  - Each node ***v*** is playing a copy of the game with each of its neighbors
  - **Payoff**: sum of node payoffs over all games

Calculation of Node *v*

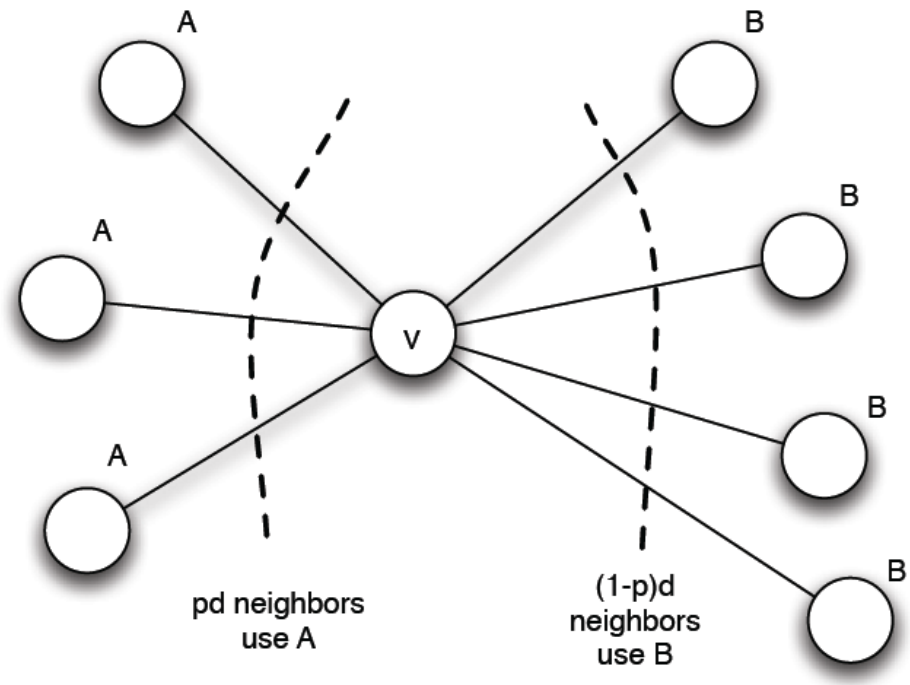

Threshold:

v chooses ***A*** if

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

p… frac. v's nbrs. with A

***q*… payoff threshold**

- Let **v** have ***d*** neighbors
- Assume fraction ***p*** of ***v***'s neighbors adopt ***A***
  - $Payoff_v = a \cdot p \cdot d$ if *v* chooses A
    $= b \cdot (1-p) \cdot d$ if *v* chooses B
- **Thus: *v* chooses A if: *p* > *q***

Example Scenario

Scenario:

- Graph where everyone starts with all ***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 q=50% of my friends take A I'll also take A.

  This means: **a = b-ε** (ε>0, small positive constant) and then **q=1/2**

Example Scenario

$S = \{u, v\}$

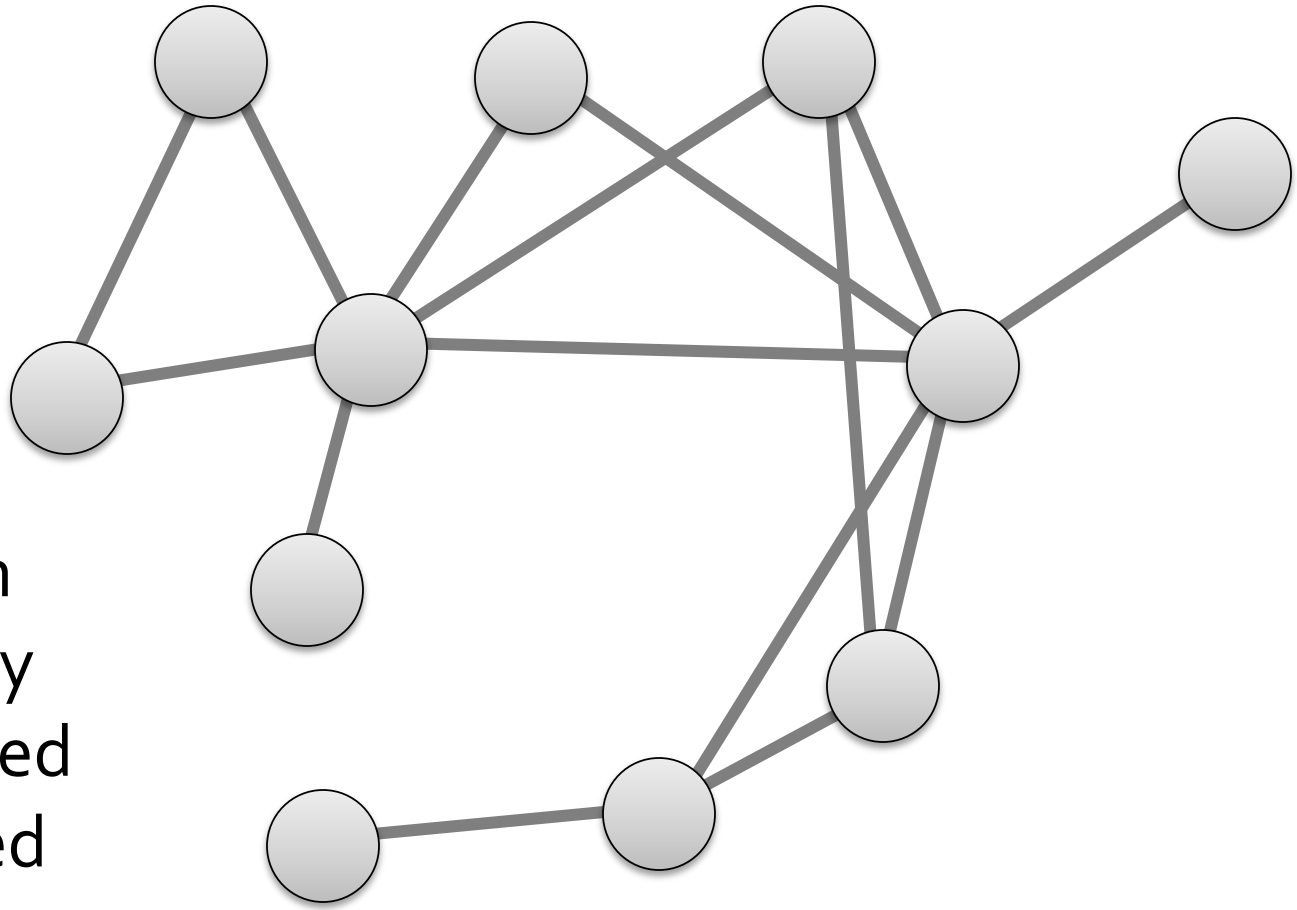

If **more** than **q=50%** of my friends are red I'll also be red

Example Scenario

$S = \{u, v\}$

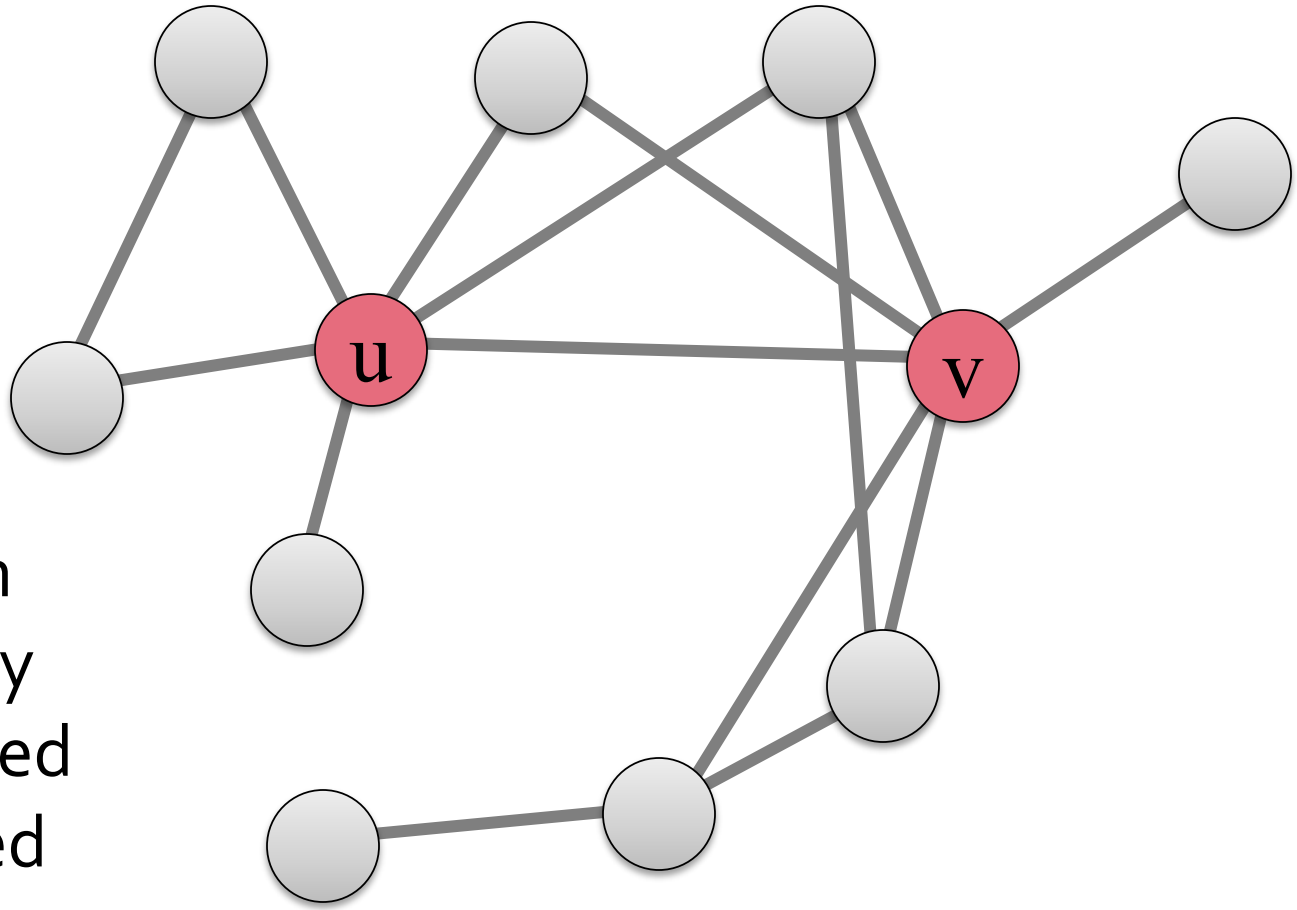

If **more** than **q=50%** of my friends are red I'll also be red

Example Scenario

$S = \{u, v\}$

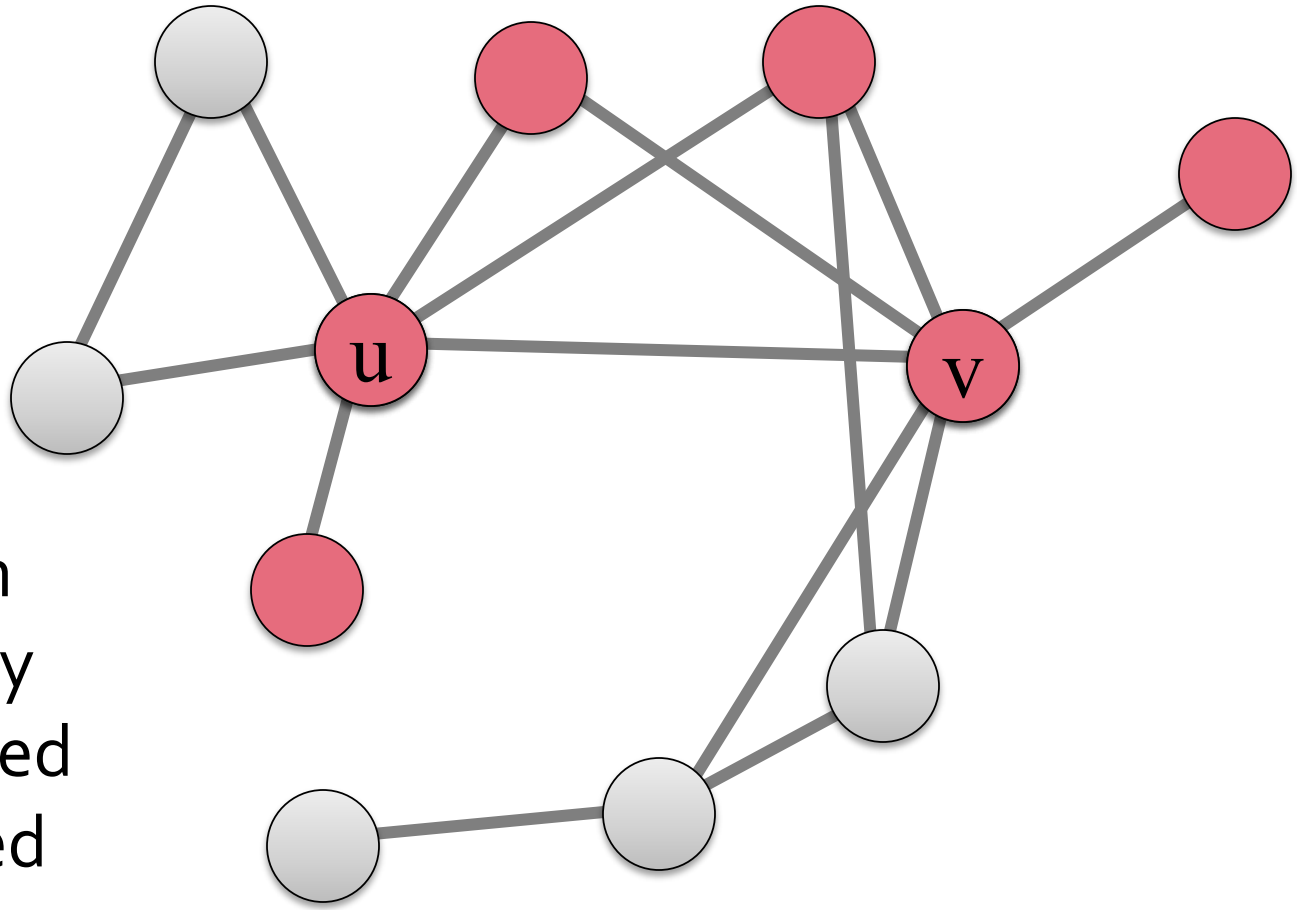

If **more** than **q=50%** of my friends are red I'll also be red

Example Scenario

$S = \{u, v\}$

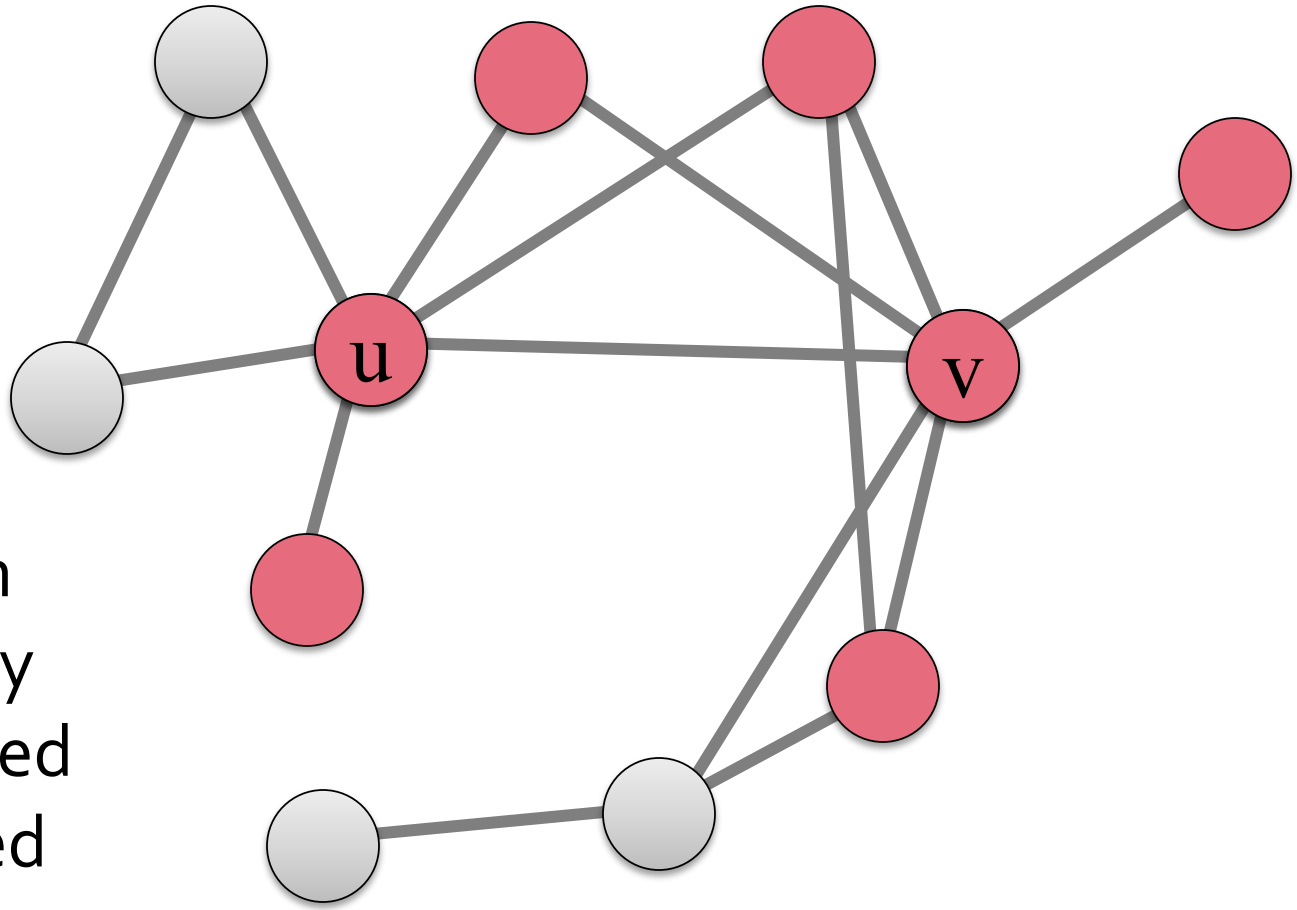

If **more** than **q=50%** of my friends are red I'll also be red

Example Scenario

$S = \{u, v\}$

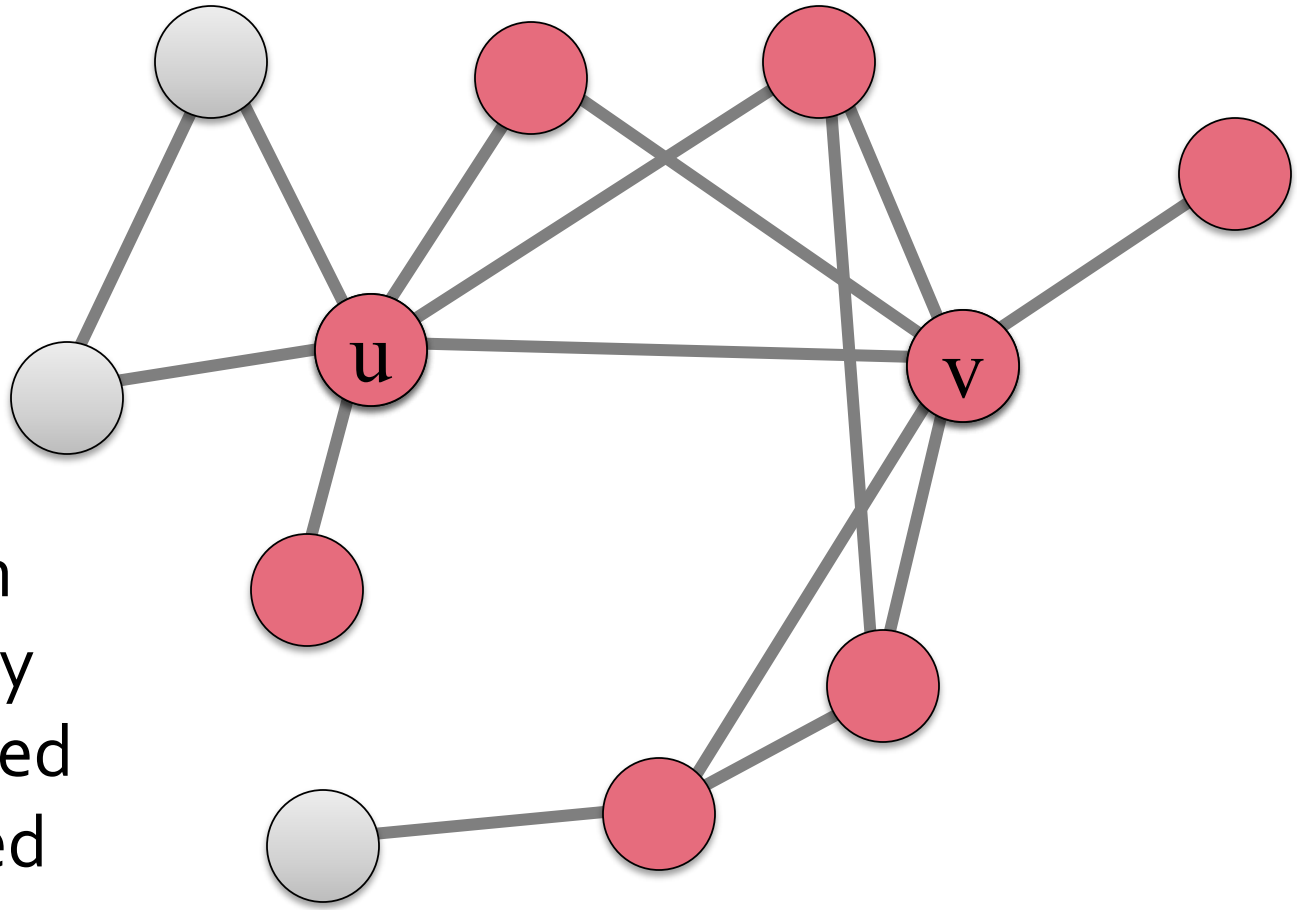

If **more** than **q=50%** of my friends are red I'll also be red

Example Scenario

$S = \{u, v\}$

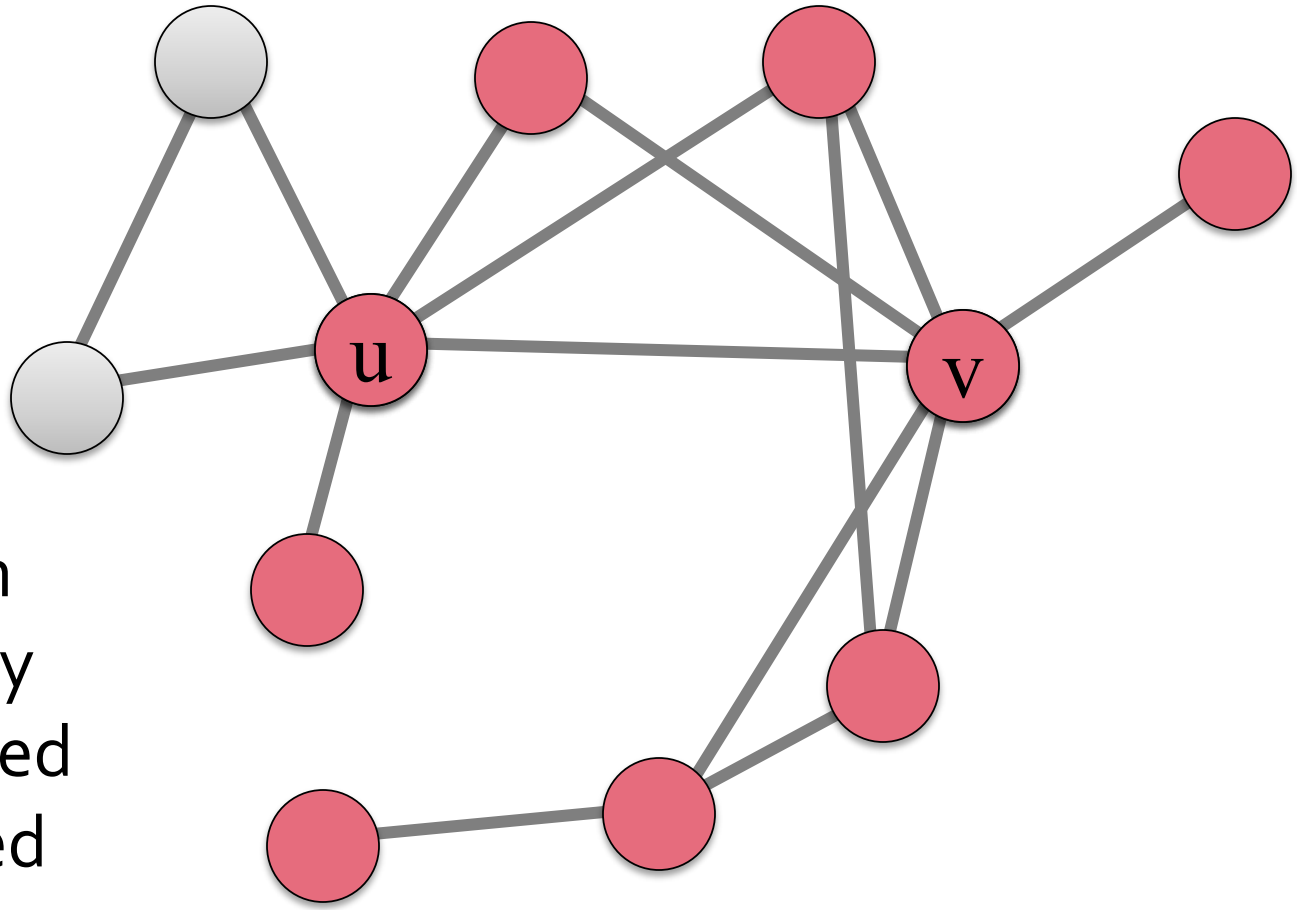

If **more** than **q=50%** of my friends are red I'll also be red

Application: Modeling protest recruitment on social networks

The Dynamics of Protest Recruitment through an Online Network

Bailon et al. Nature Scientific Reports, 2011

The Spanish 'Indignados' Movement

- **Anti-austerity protests in Spain May 15-22, 2011 as a response to the financial crisis**
- Twitter was used to organize and mobilize users to participate in the protest

Data collected using hashtags

- **Researchers identified 70 hashtags that were used by the protesters**

#hashtags

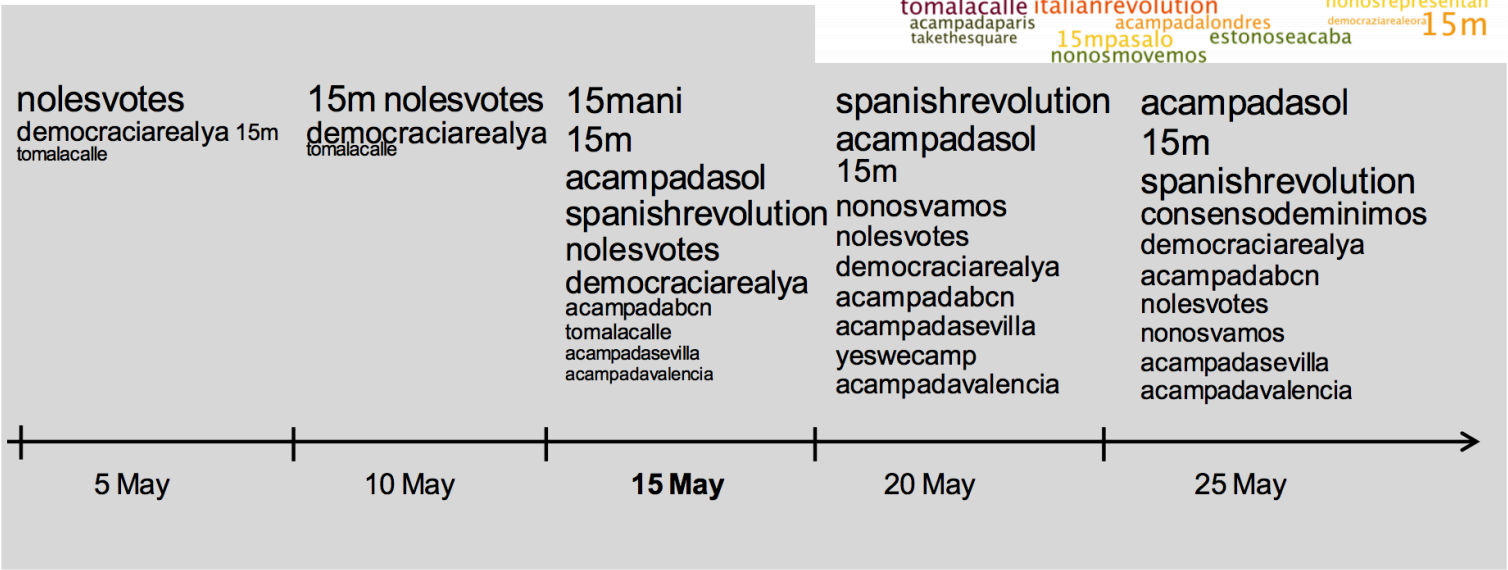

Dataset

- **70 hashtags were crawled for 1 month period**
  - Number of tweets: 581,750
- **Relevant users:** Any user who tweeted any relevant hashtag and their followers + followees
  - Number of users: 87,569
- **Created two undirected follower networks:**
  1. **Full network:** with all Twitter follow links
  2. **Symmetric network** with only the reciprocal follow links ($i \rightarrow j$ and $j \rightarrow i$)
     - This network represents "strong" connections only.

Definitions

- **User activation time:** Moment when user starts tweeting protest messages
- k_{in} = The total number of neighbors when a user became active
- k_a = Number of active neighbors when a user became active
- **Activation threshold = k_a/k_{in}**
  - The fraction of active neighbors at the time when a user becomes active

Recruitment & Activation Threshold

- If $k_a/k_{in} \approx 0$, then the user joins the movement when very few neighbors are active ⇒ no social pressure
- If $k_a/k_{in} \approx 1$, then the user joins the movement after most of its neighbors are active ⇒ high social pressure

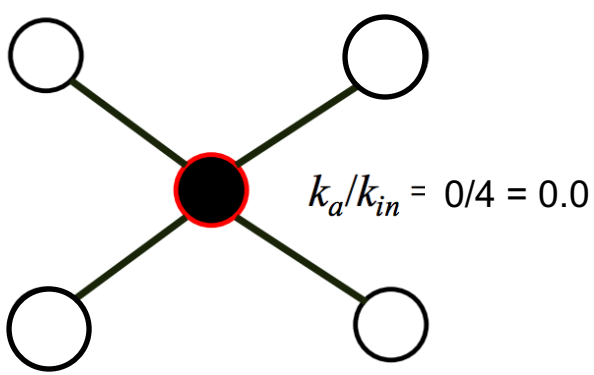

No social pressure for middle node to join

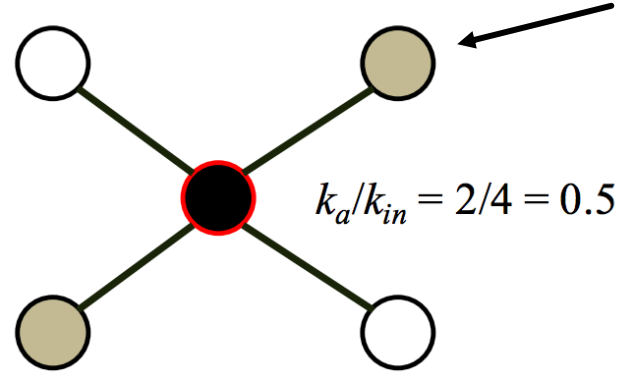

Non-zero social pressure for middle node to join

Distribution of activation thresholds

- Mostly uniform distribution of activation threshold in both networks, except for two local peaks

0 activation threshold users: Many self-active users.

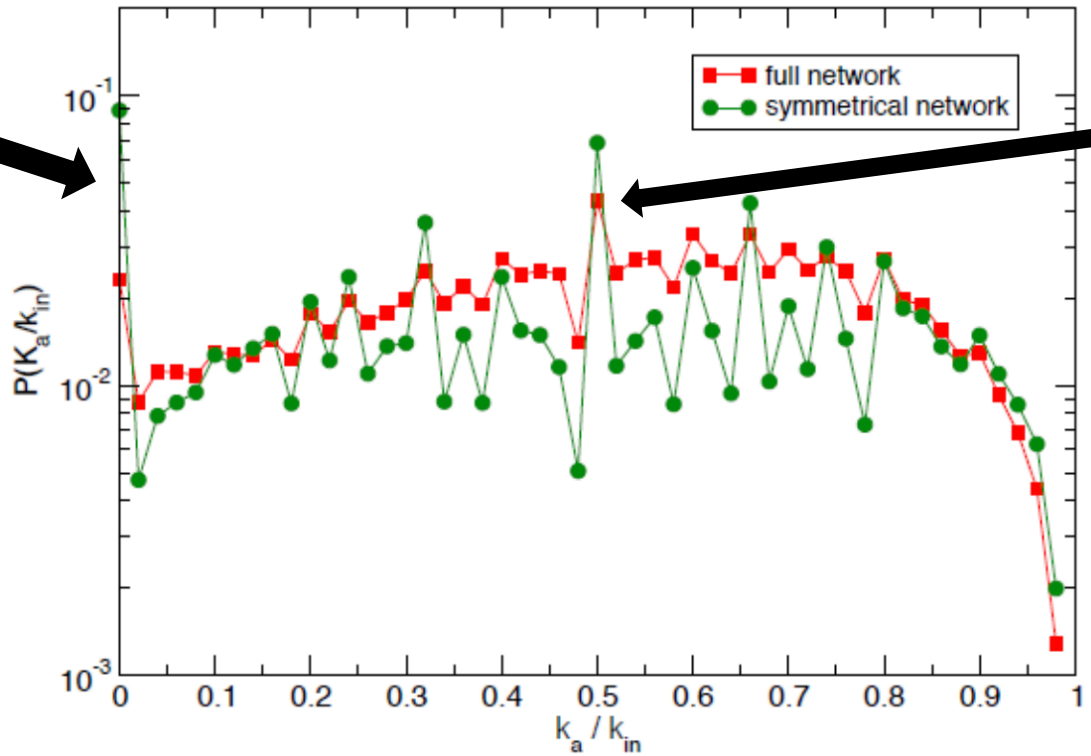

0.5 activation threshold users: Many users who join after half their neighbors do.

Effect of neighbor activation time

- **Hypothesis:** If several neighbors become active in a short time period, then a user is more likely to become active
- **Method:** Calculate the burstiness of active neighbors as

$$\Delta k_a/k_a = (k_a^{t+1} - k_a^{t})/k_a^{t+1}$$

Low threshold users are insensitive to recruitment bursts.

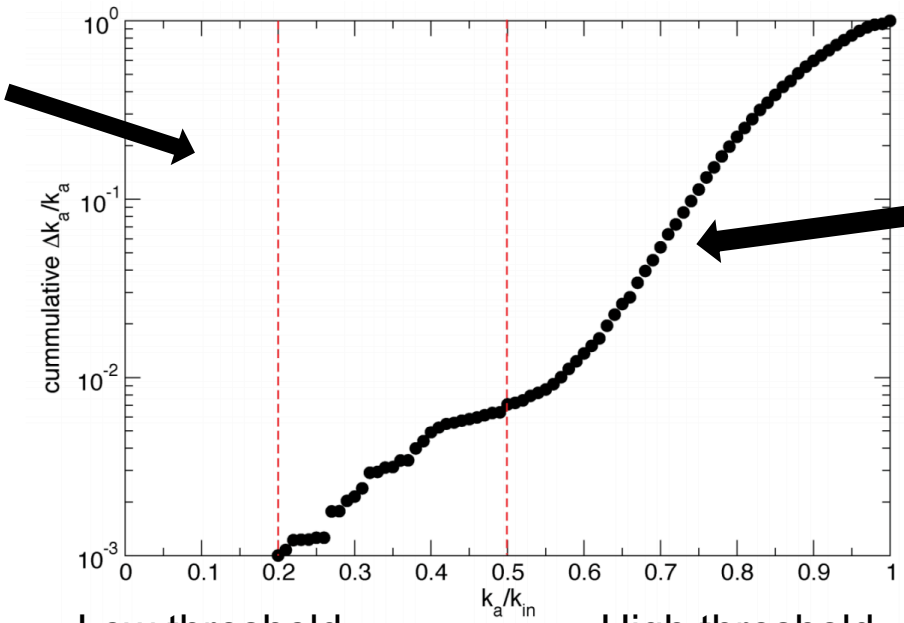

High threshold users join after sudden bursts in neighborhood activation

Low threshold users

High threshold users

Information cascades

- **No cascades are given in the data**
- **So cascades were identified as follows:**
  - If a user tweets a message at time *t* and one of its followers tweets a message in (t, t+**Δ**t), then they form a cascade.
  - **E.g., 1 ➞ 2 ➞ 3 below form a cascade:**

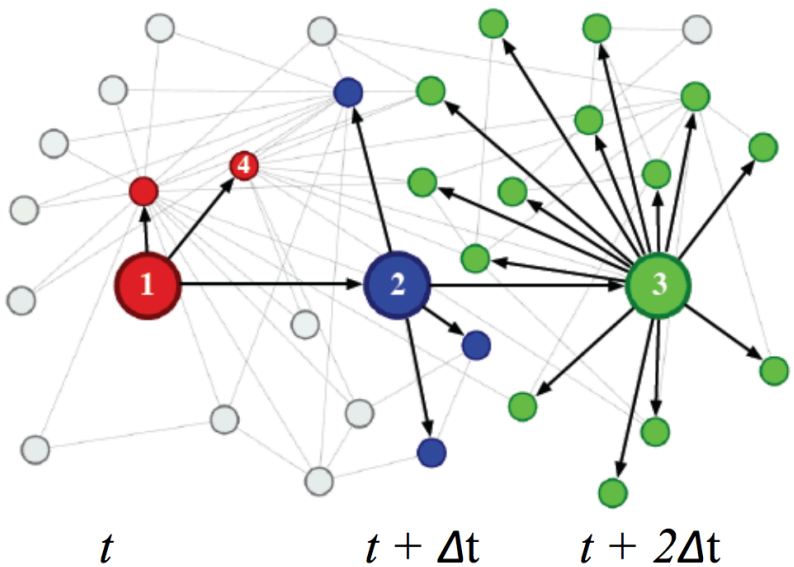

Size of information cascades

- **Size** = number of nodes in the cascade
- **Most cascades are small:**

Fraction of cascades with size at least S

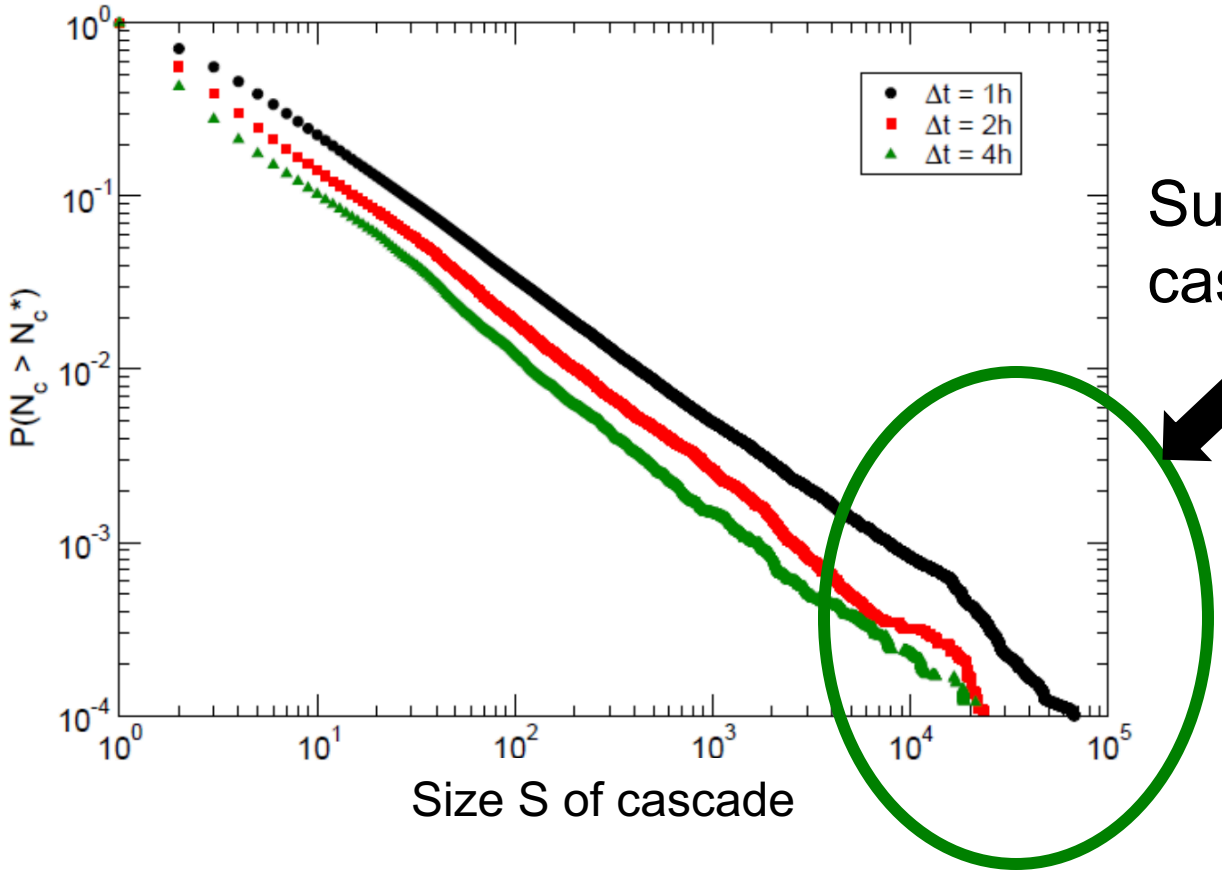

Successful cascades

Who starts successful cascades?

- **Are starters of successful cascades more central in the network?**
- **Method:** *k*-core decomposition
  - *k*-core: biggest connected subgraph where every node has at least degree k
  - Method: repeatedly remove all nodes with degree less than *k*
  - Higher *k*-core number of a node means it is more central

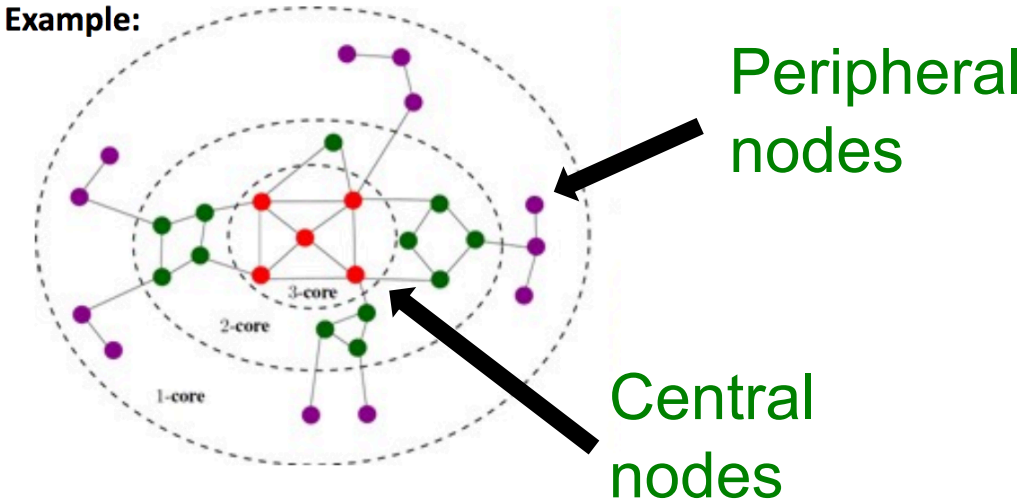

Who starts the successful cascades?

- K-core decomposition of follow network
  - Red nodes start successful cascades
  - Red nodes have higher *k*-core values
    - So, successful cascade starters are central and connected to equally well connected users

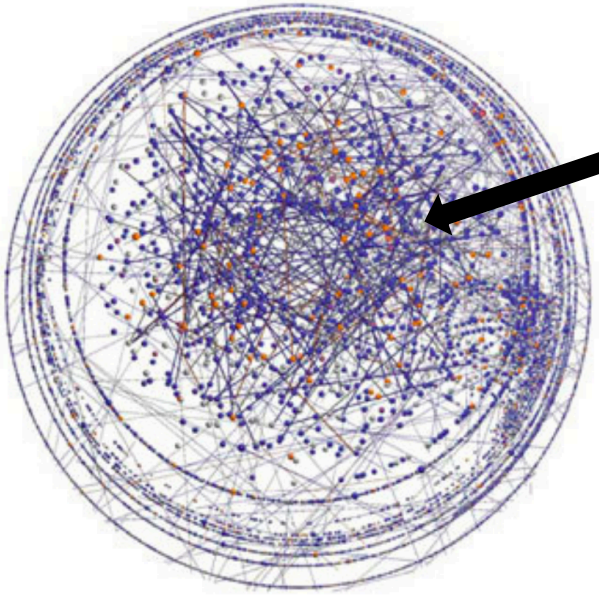

Successful cascade starters are central (higher k-core number)

Summary: Cascades on Twitter

- Uniform activation threshold for users, with two local peaks
- Most cascades are short
- Successful cascades are started by central (more core) users

Models of Cascading Behavior

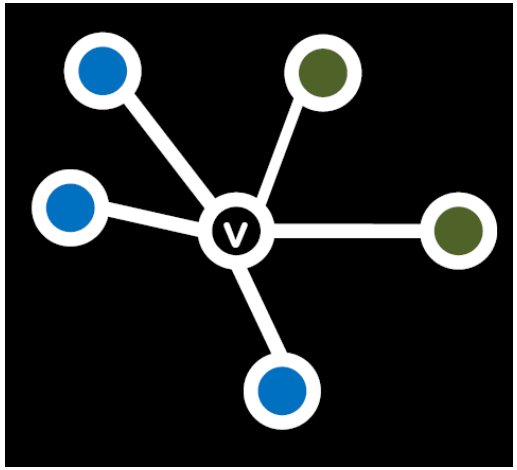

- **So far:**
  Decision Based Models
  - Utility based
  - Deterministic
  - "Node" centric: A node observes decisions of its neighbors and makes its own decision
- **Next: Extending decision based models to multiple contagions**

Extending the Model: Allow People to Adopt A and B

Extending the model

- **So far:**
  - Behaviors ***A*** and ***B*** compete
  - Can only get utility from neighbors of same behavior: ***A*-*A* get a**, ***B*-*B* get *b***, ***A*-*B* get *0***
- **For example:**
  - **Using Skype vs. WhatsApp**
    - Can only talk using the same software
  - **Having a VHS vs. BetaMax player**
    - Can only share tapes with people using the same type of tape

  - **But one can buy 2 players or install 2 programs**

Cascades & Compatibility

- **So far:**
  - Behaviors ***A*** and ***B*** compete
  - Can only get utility from neighbors of same behavior: ***A-A*** **get a**, ***B-B*** **get *b***, ***A-B*** **get *0***
- **Let's add an extra strategy "*AB*"**
  - ***AB-A*** **:** gets ***a***
  - ***AB-B*** **:** gets ***b***
  - ***AB-AB*** **:** gets **max(*a*, *b*)**
  - **Also:** Some **cost *c*** for the effort of maintaining both strategies (summed over all interactions)
    - Note: a given node can receive *a* from one neighbor and *b* from another by playing AB, which is why it could be worth the cost *c*

Cascades & Compatibility: Model

- Every node in an infinite network starts with ***B***
- Then a finite set ***S*** initially adopts ***A***
- Run the model for ***t=1,2,3,…***
  - Each node selects behavior that will optimize payoff (given what its neighbors did in at time ***t-1***)

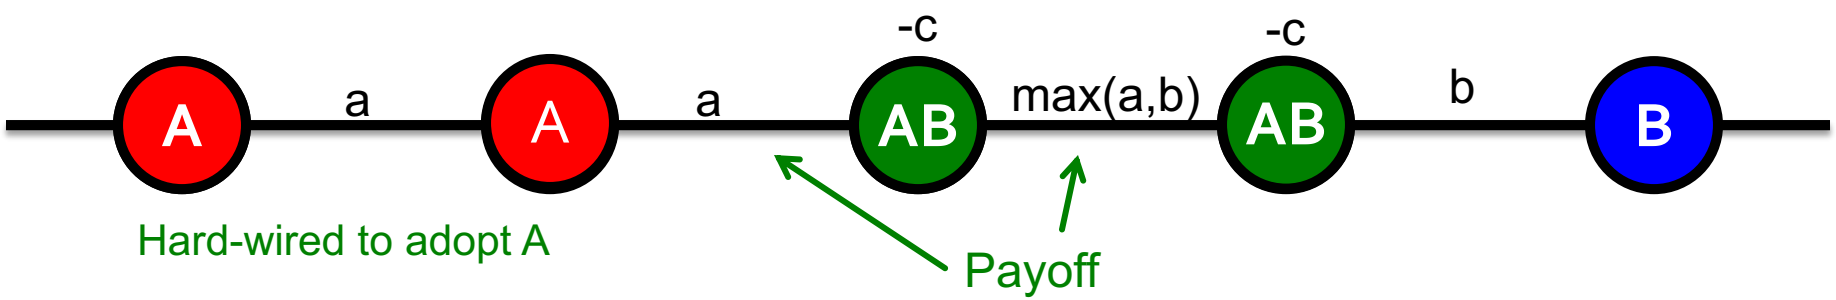

- **How will nodes switch from *B* to *A* or *AB*?**

Example: Path Graph (1)

- **Path graph:** Start with **B**s, **a** > **b** (**A** is better)
- **One node switches to A – what happens?**
  - **With just A, B: A** spreads if ***a*** > ***b***
  - **With A, B, AB: Does A spread?**
- **Example: a=3, b=2, c=1**

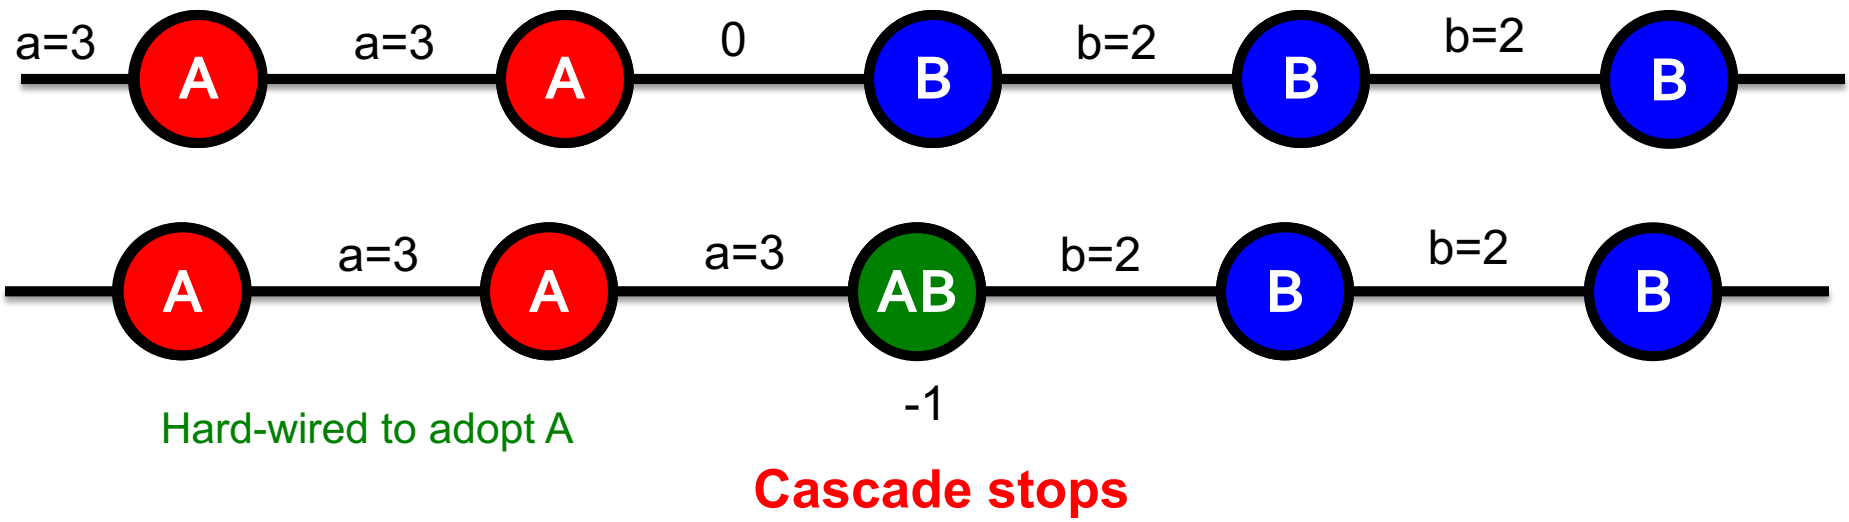

Example: Path Graph (2)

- **Example: a=5, b=3, c=1**

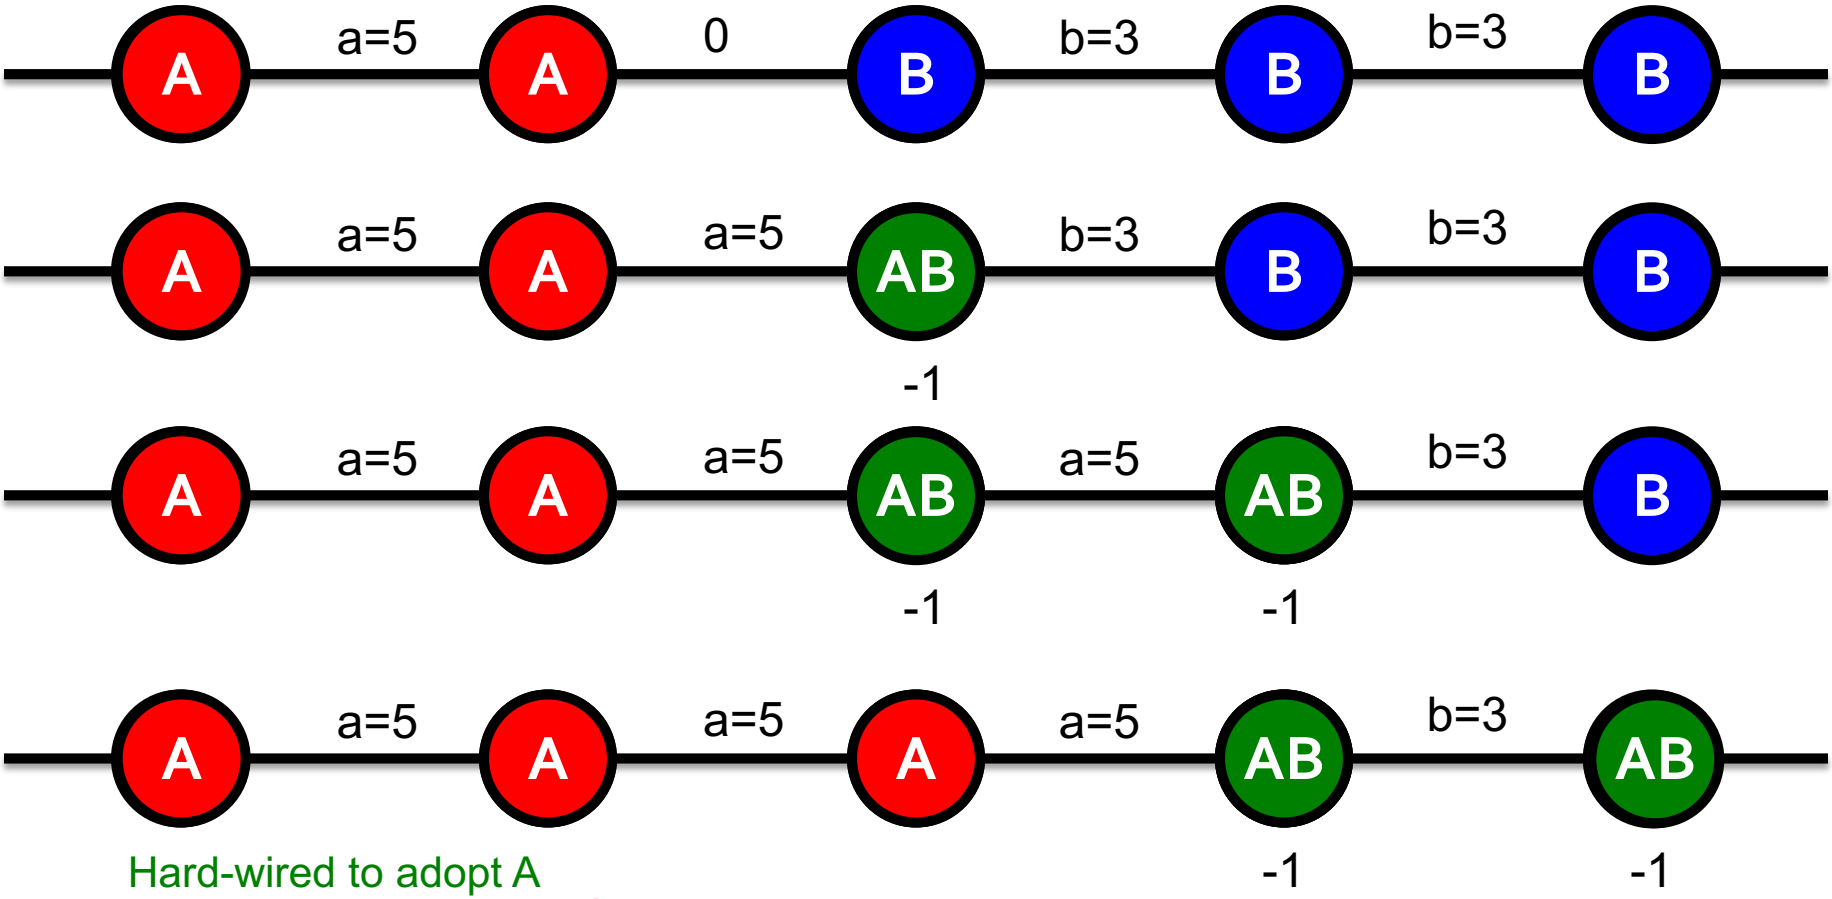

Cascade never stops!

What about in a general case?

- Let's solve the model in a general case:
  - **Infinite path, start with all Bs**
  - **Payoffs for *w*:** **A**:a, **B**:1, **AB**:a+1-c
- **For what pairs (c,a) does A spread?**
  - **We need to analyze two cases for node *w*:** Based on the values of a and c, what would w do?

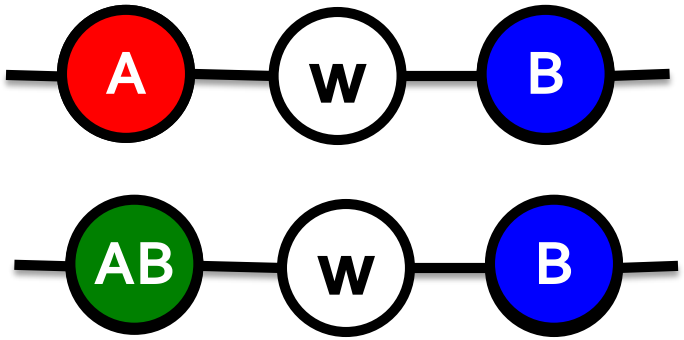

For what pairs (c,a) does A spread?

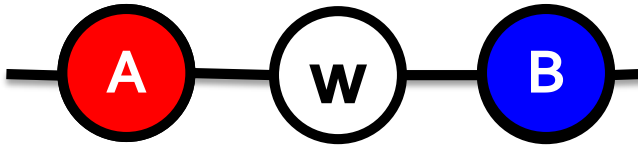

- **Infinite path, start with Bs**
- **Payoffs for *w*: A**:a, **B**:1, **AB**:a+1-c
- What does node ***w*** adopt?

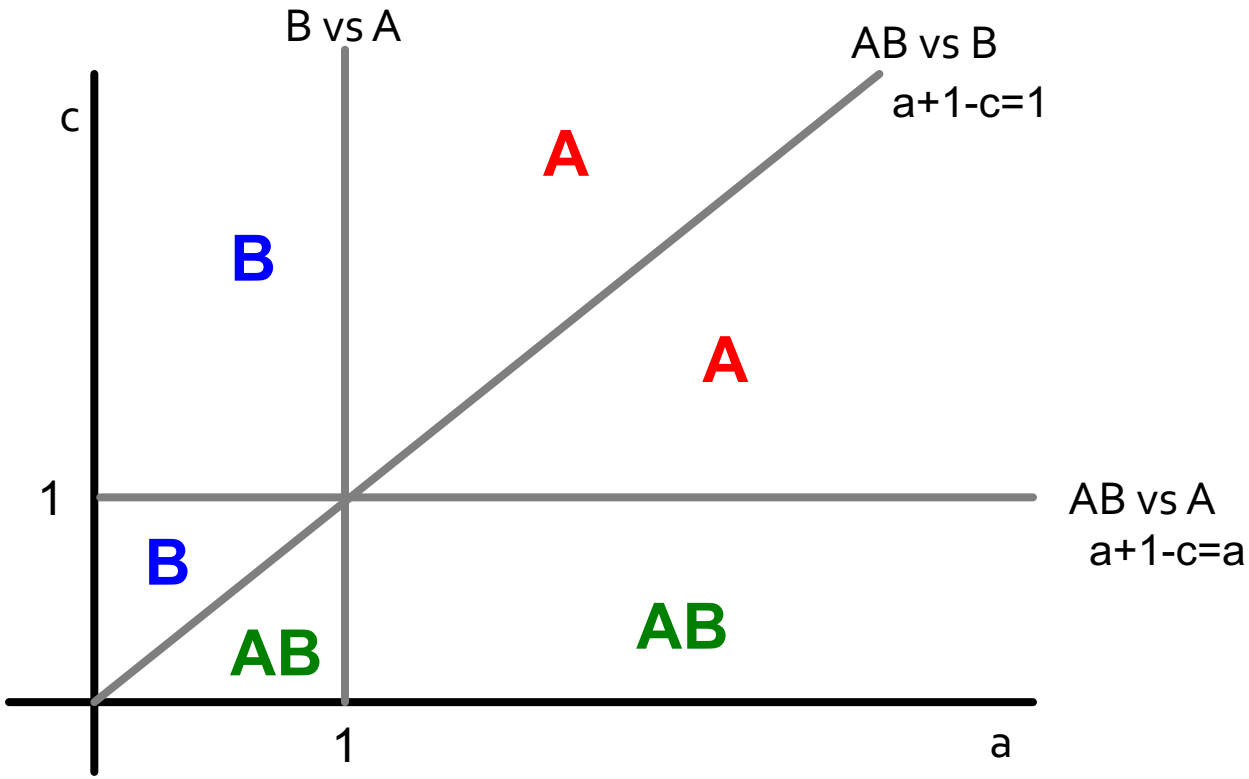

For what pairs (c,a) does A spread?

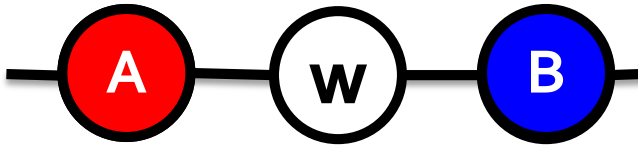

- **Infinite path, start with Bs**
- **Payoffs for *w*:** **A**:a, **B**:1, **AB**:a+1-c
- What does node ***w*** in **A-*w*-B** do?

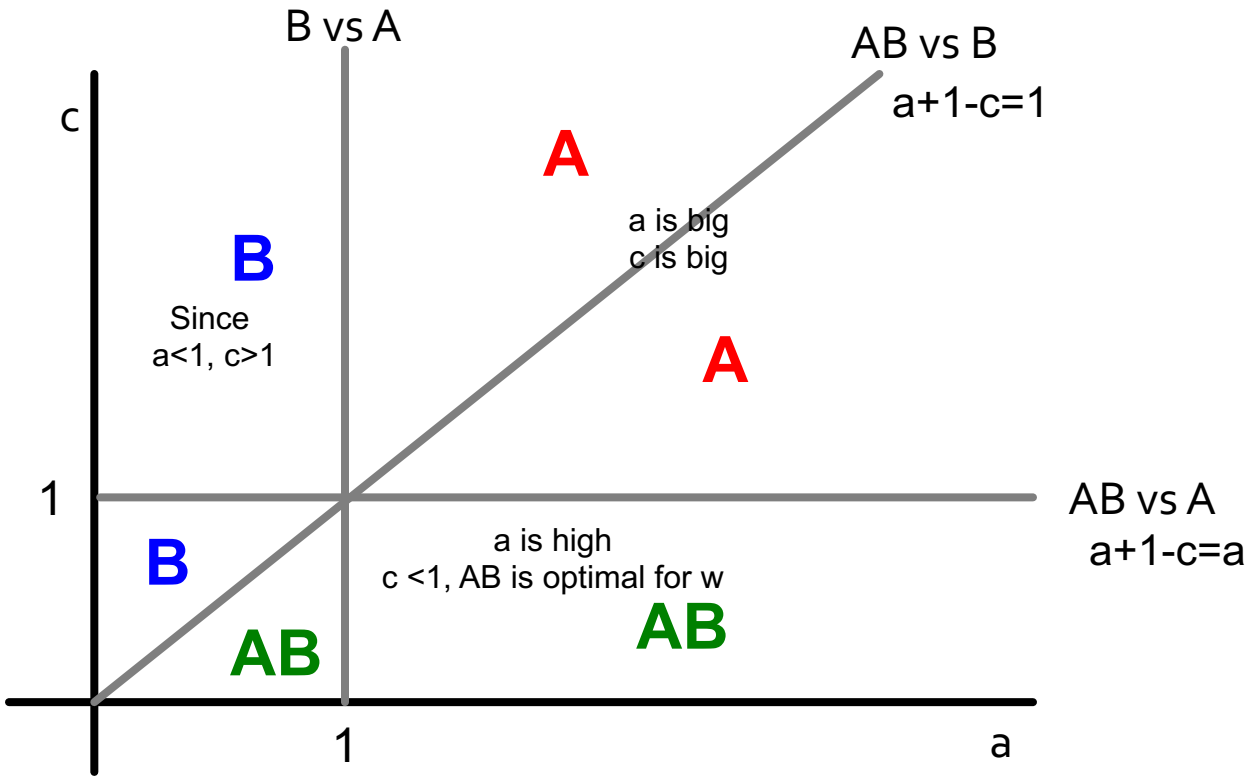

For what pairs (c,a) does A spread?

- **Same reward structure as before but now payoffs for *w* change:** **A**:a, **B**:1+1, **AB**:a+1-c
- **Notice: Now also AB spreads**
- What does node ***w*** in **AB-*w*-B** do?

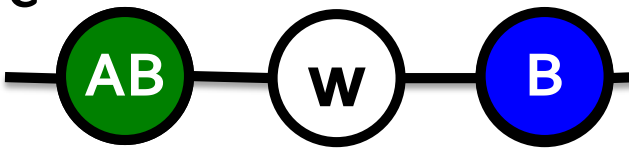

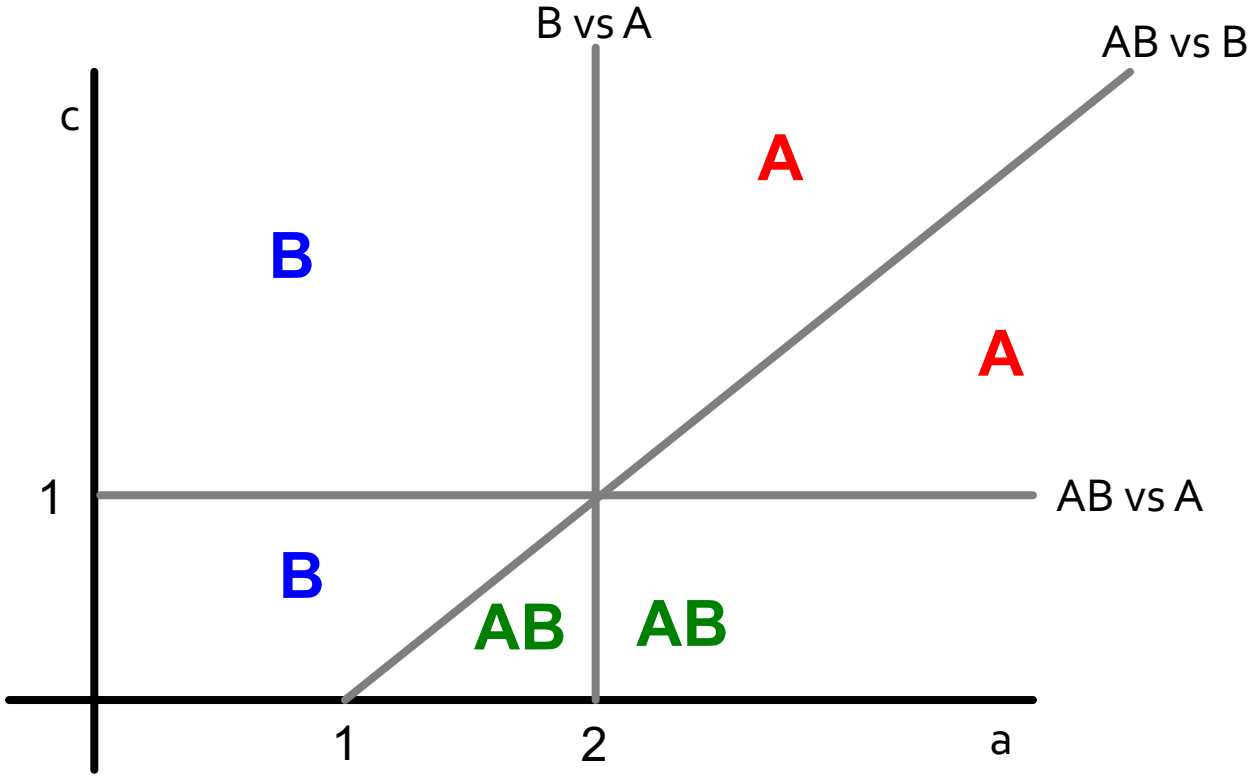

For what pairs (c,a) does A spread?

- **Same reward structure as before but now payoffs for *w* change:** **A**:a, **B**:1+1, **AB**:a+1-c
- **Notice: Now also AB spreads**
- What does node ***w*** in **AB-*w*-B** do?

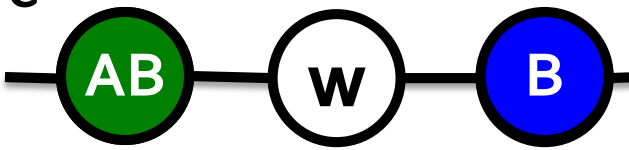

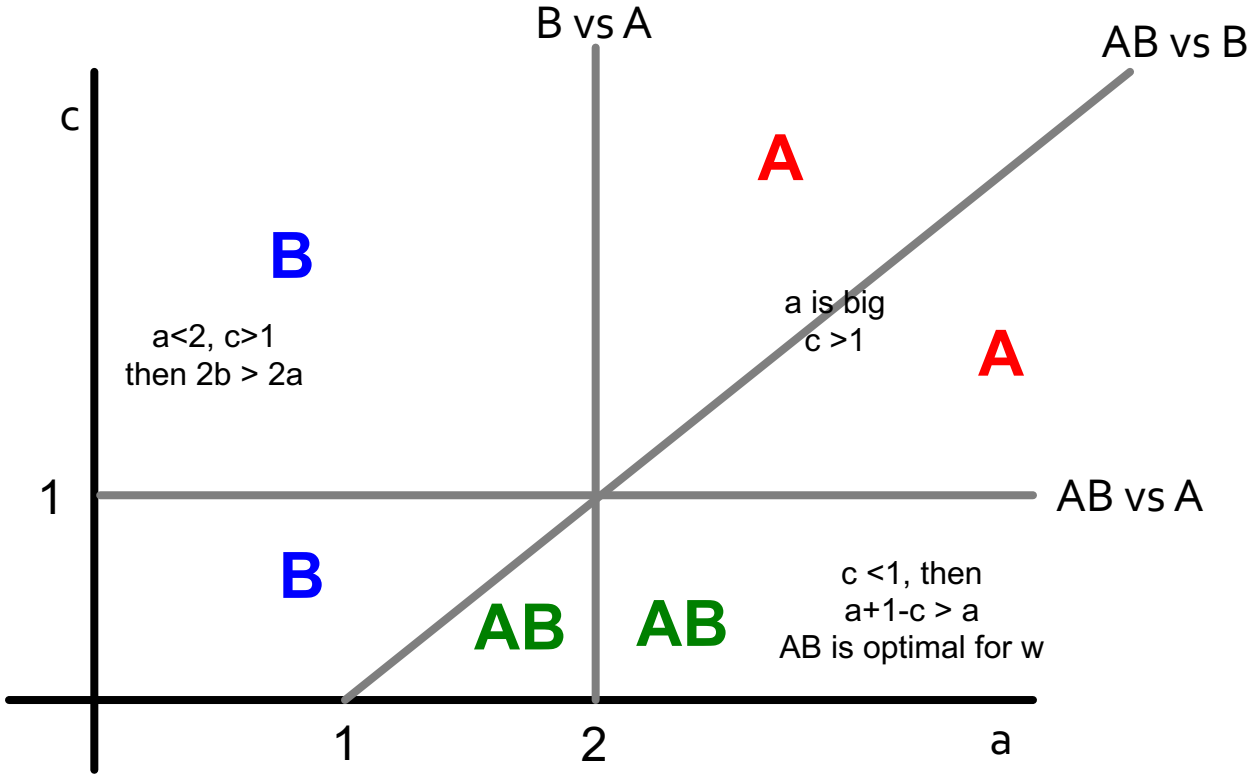

For what pairs (c,a) does A spread?

- **Joining the two pictures:**

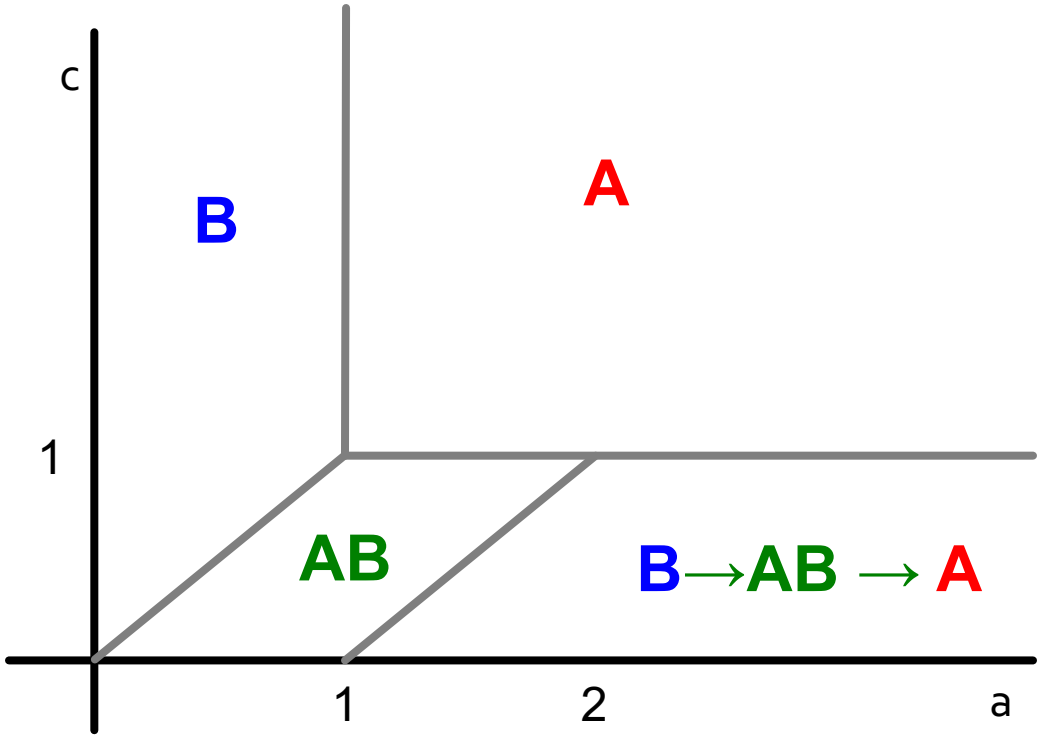

Lesson

- ***B* is the default throughout the network until new/better *A* comes along. What happens?**
  - **Infiltration:** If **B** is too compatible then people will take on both and then drop the worse one **(B)**
  - **Direct conquest:** If **A** makes itself not compatible – people on the border must choose. They pick the better one (**A**)
  - **Buffer zone:** If you choose an optimal level then you keep a static "buffer" between **A** and **B**

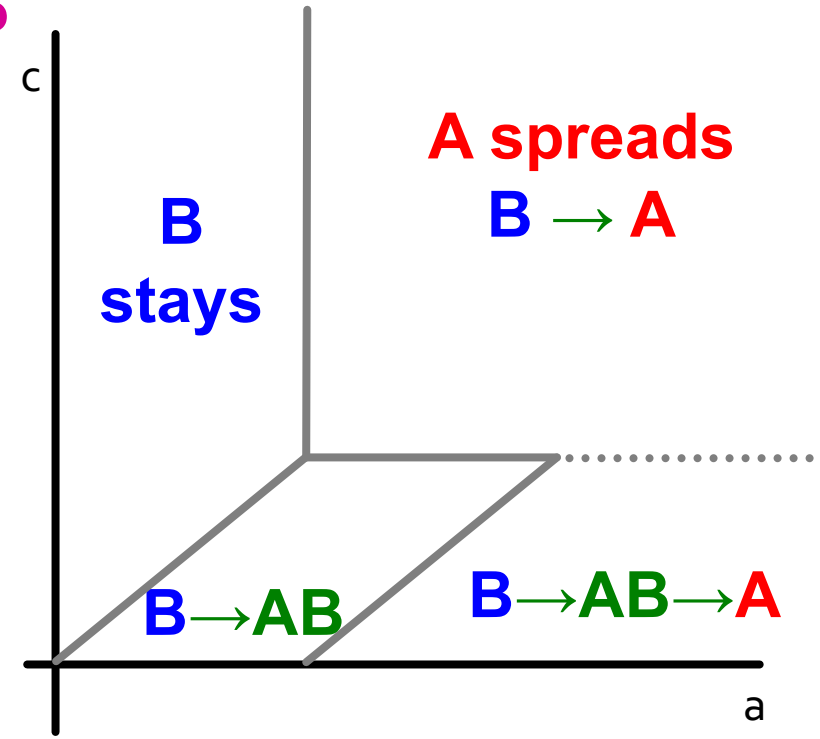

Models of Cascading Behavior

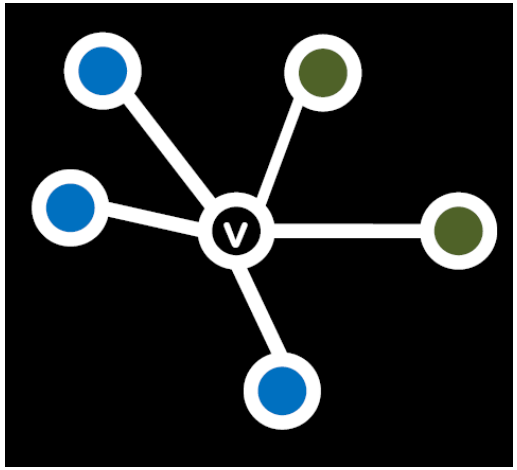

- **So far:**
  Decision Based Models
  - Utility based
  - Deterministic
  - "Node" centric: A node observes decisions of its neighbors and makes its own decision
  - Require us to know too much about the data
- **Next: Probabilistic Models**
  - Lets you do things by observing data
  - **Limitation**: we can't model causality

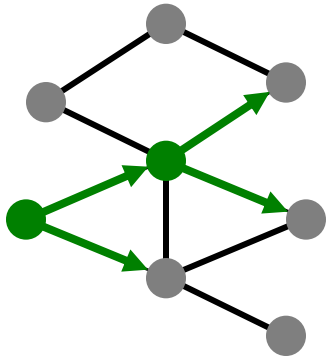

TRAILER:

Probabilistic Contagion and Models of Influence

Epidemics vs Cascade Spreading

- In decision-based models nodes make decisions based on pay-off benefits of adopting one strategy or the other.
- **In epidemic spreading:**
  - Lack of decision making
  - Process of contagion is complex and unobservable
    - In some cases it involves (or can be modeled as) randomness

Simple model: Branching Process

- **First wave**: A person carrying a disease enters the population and transmits to all she meets with probability q. She meets d people, a portion of which will be infected.
- **Second wave**: Each of the d people goes and meets d different people. So we have a second wave of $d * d = d^2$ people, a portion of which will be infected.
- **Subsequent waves**: same process

Example with k=3

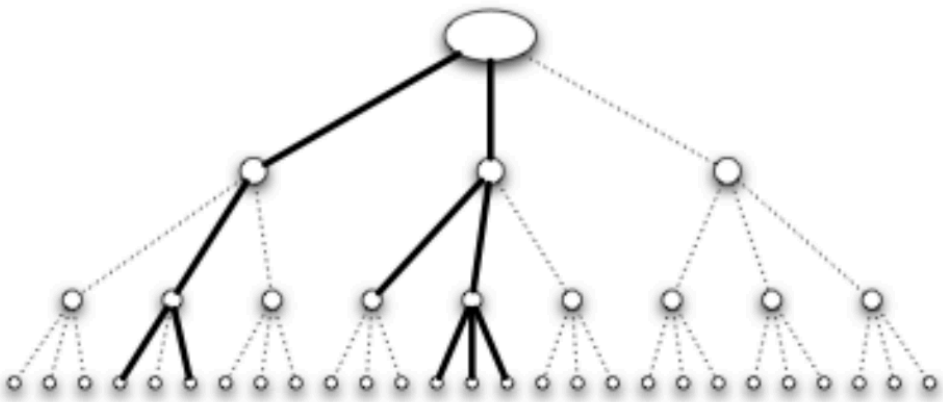

High contagion probability:
The disease spreads

Low contagion probability:
The disease dies out

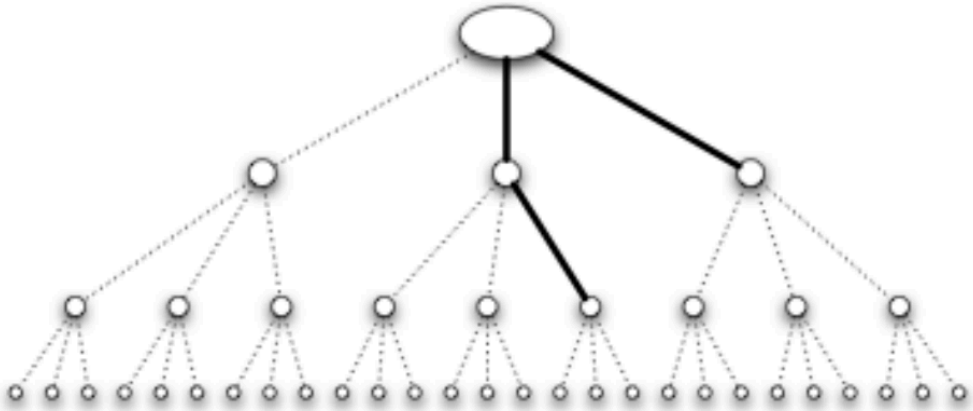

Probabilistic Spreading Models

- **Epidemic Model based on Random Trees**
  - **(a variant of a branching processes)**
  - A patient meets ***d*** other people
  - With probability ***q > 0*** she infects each of them
- **Q: For which values of *d* and *q* does the epidemic run forever?**
  - **Run forever:** $\lim_{h\to\infty} P\begin{bmatrix} \boldsymbol{a\ node\ is\ infected} \\ \boldsymbol{at\ depth\ h} \end{bmatrix} > \mathbf{0}$
  - **Die out:** -- || -- $= 0$