Decision Based Models of Cascades

CS224W: Social and Information Network Analysis
Jure Leskovec, Stanford University
http://cs224w.stanford.edu

RECAP: Game Theoretic Model of Cascades

[Morris 2000]

Game Theoretic Model of Cascades

- **Based on 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

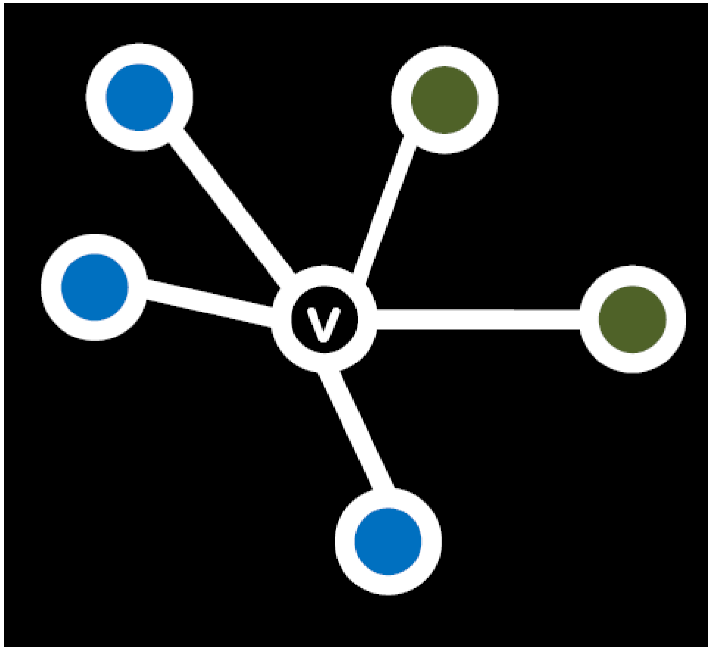

Local view of the network of node v

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 per game

Calculation of Node v

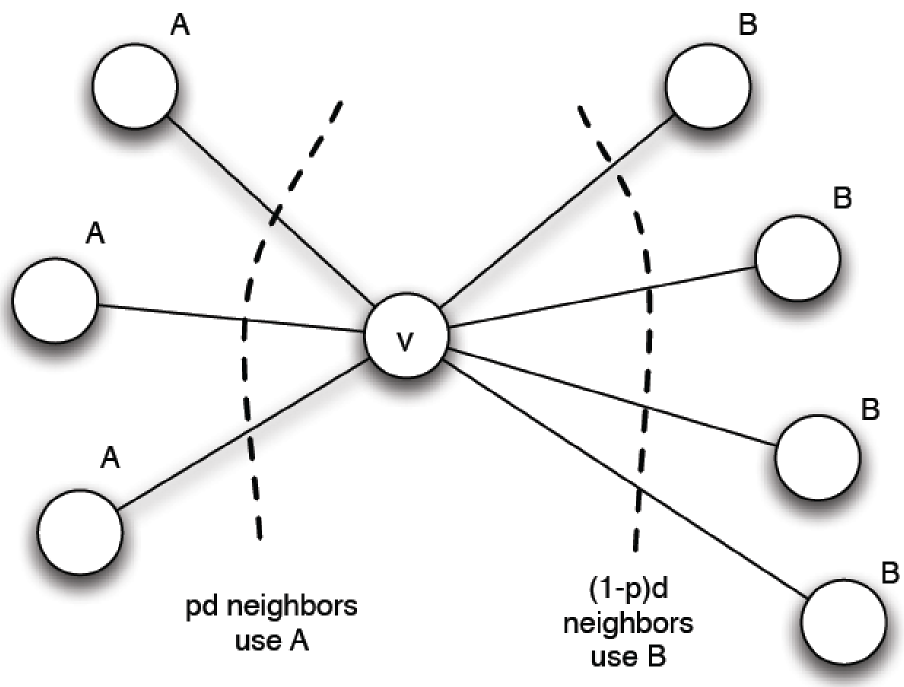

Threshold:

v choses *A* if *p*>*q*

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

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

Extending the Model:
Allow People to Adopt A and B

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 extra strategy "*A-B*"**
  - *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)

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*)

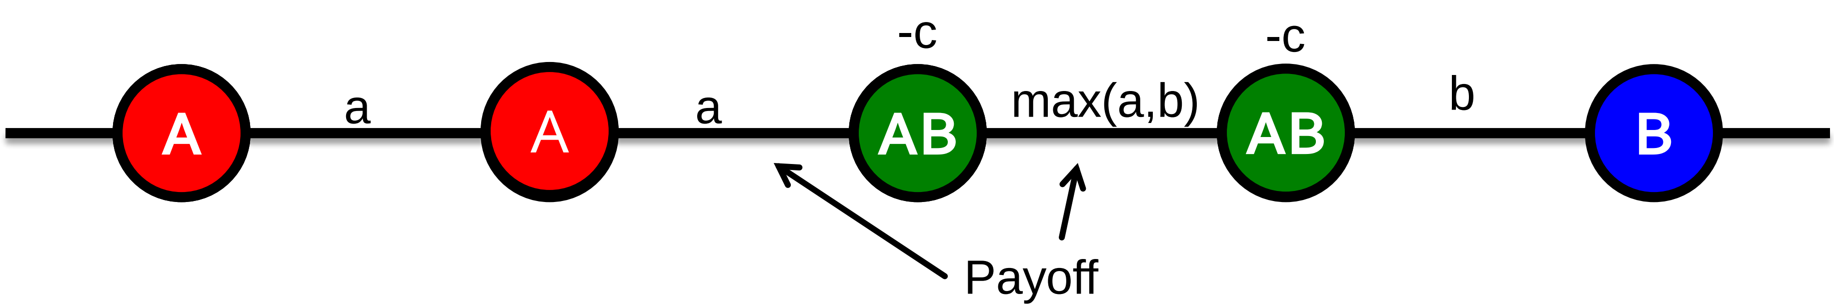

- **How will nodes switch from *B* to *A* or *AB*?**

Example

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

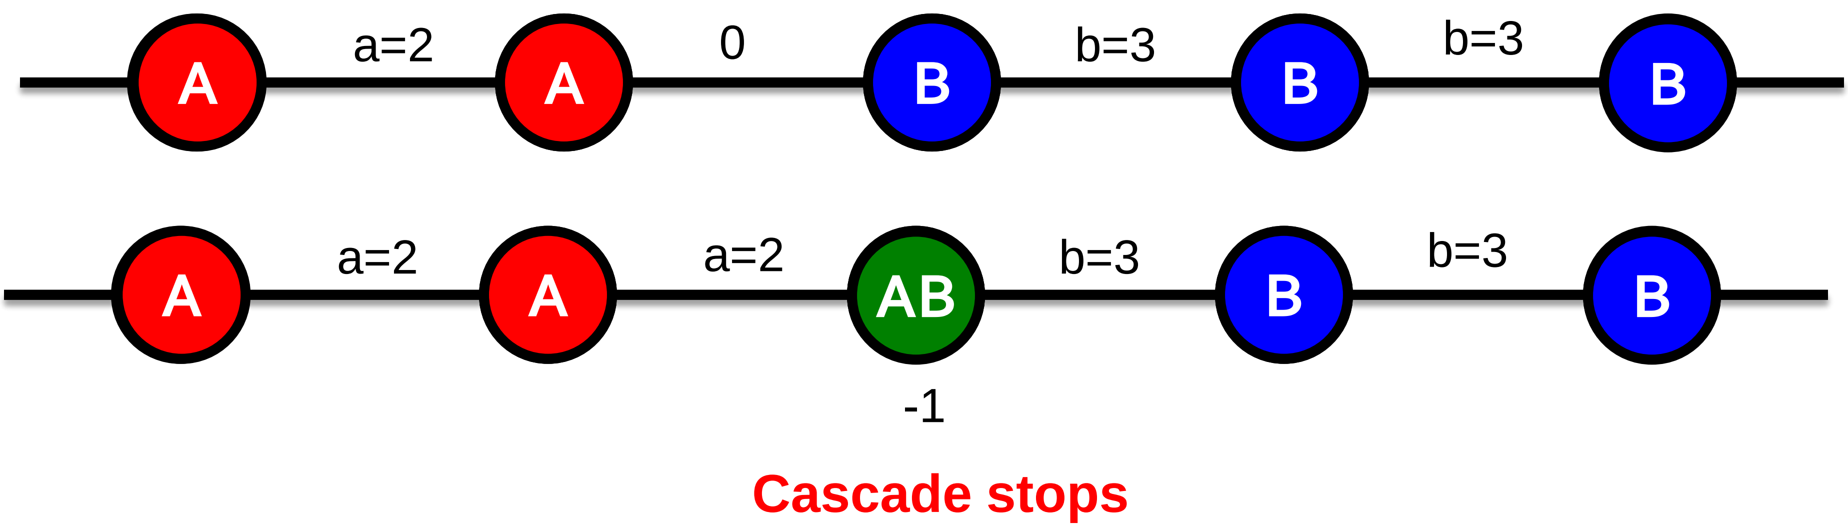

Example

- Let a=5, b=3, c=1

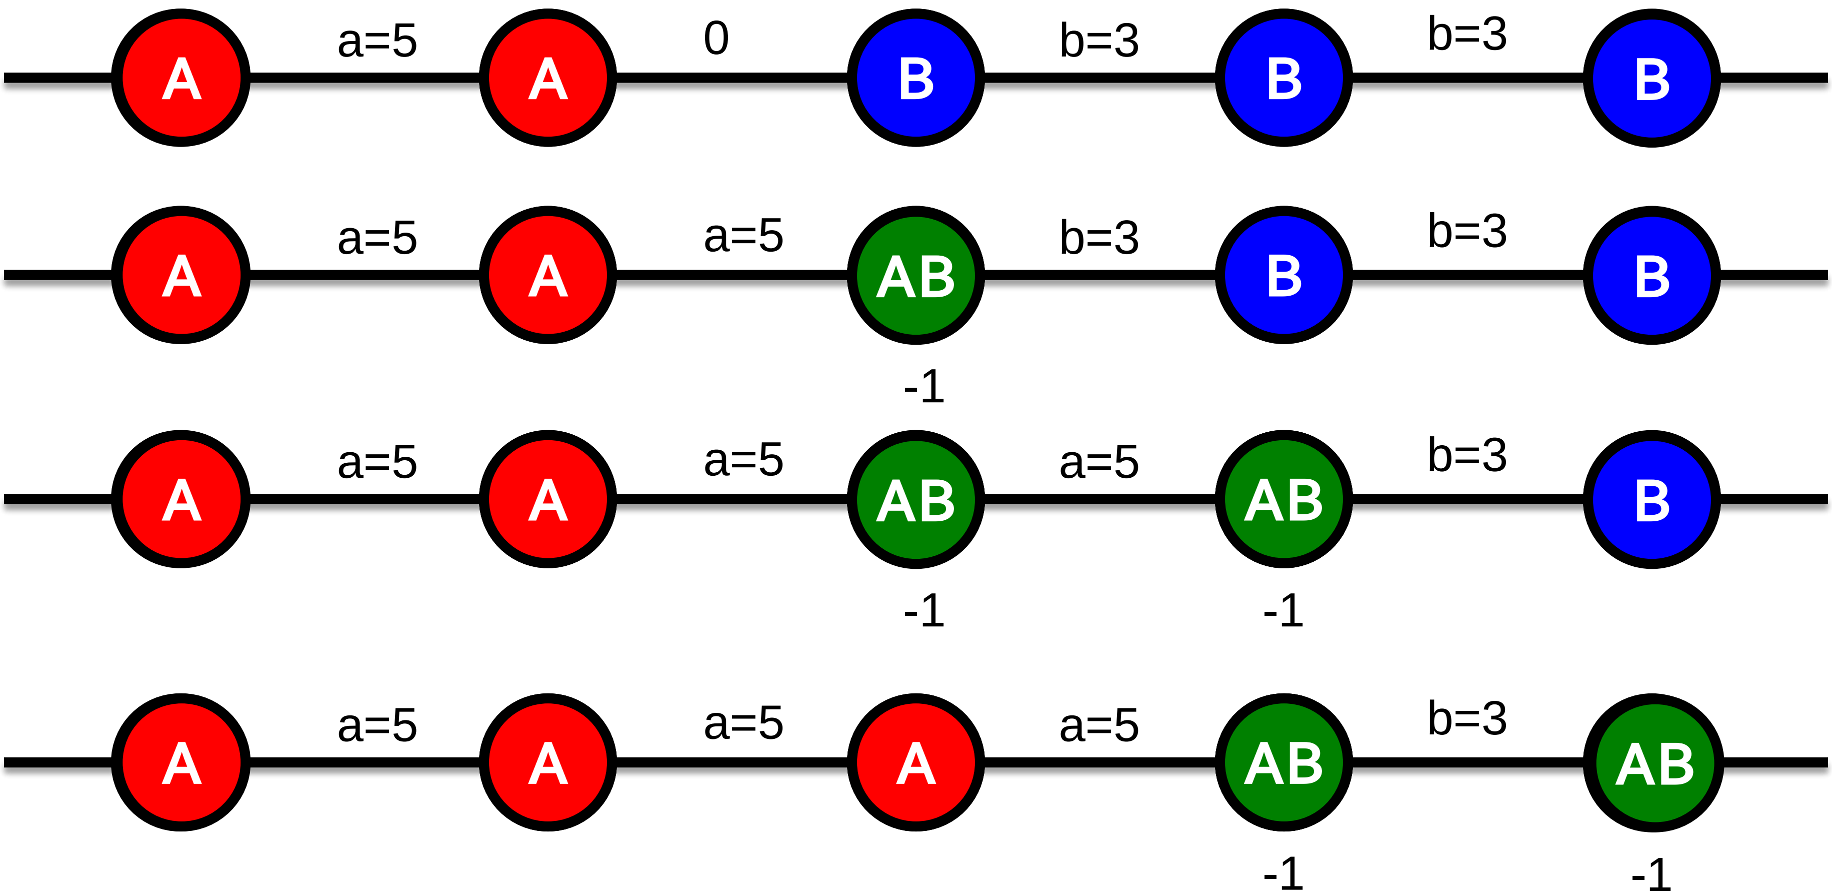

For what pairs (c,a) does A spread?

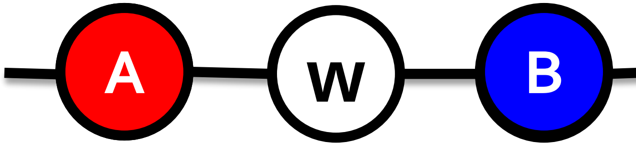

- **Infinite path, start with all Bs**
- **Payoffs for *w*:** A:a, B:1, AB:a+1-c
- What does node w in **A-w-B** do?

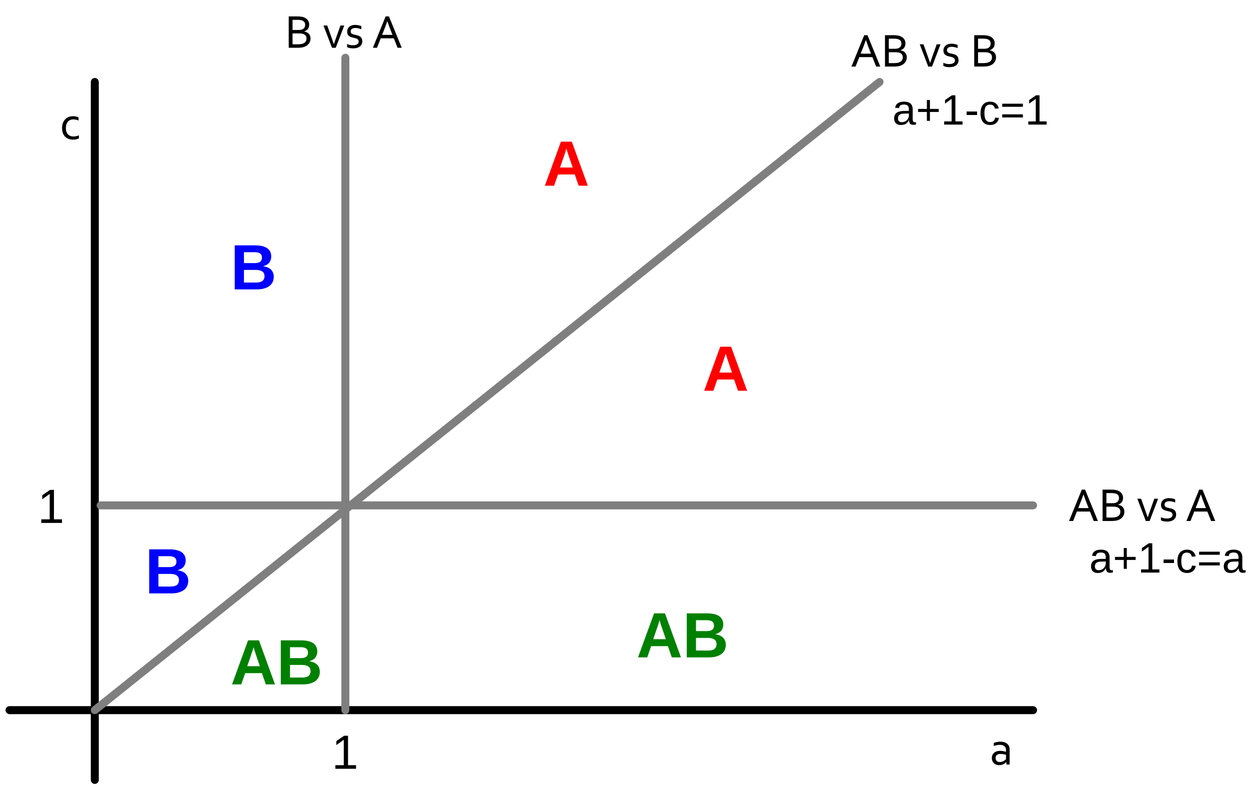

For what pairs (c,a) does A spread?

- **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?

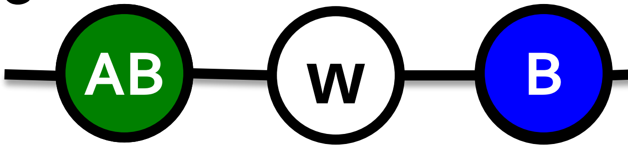

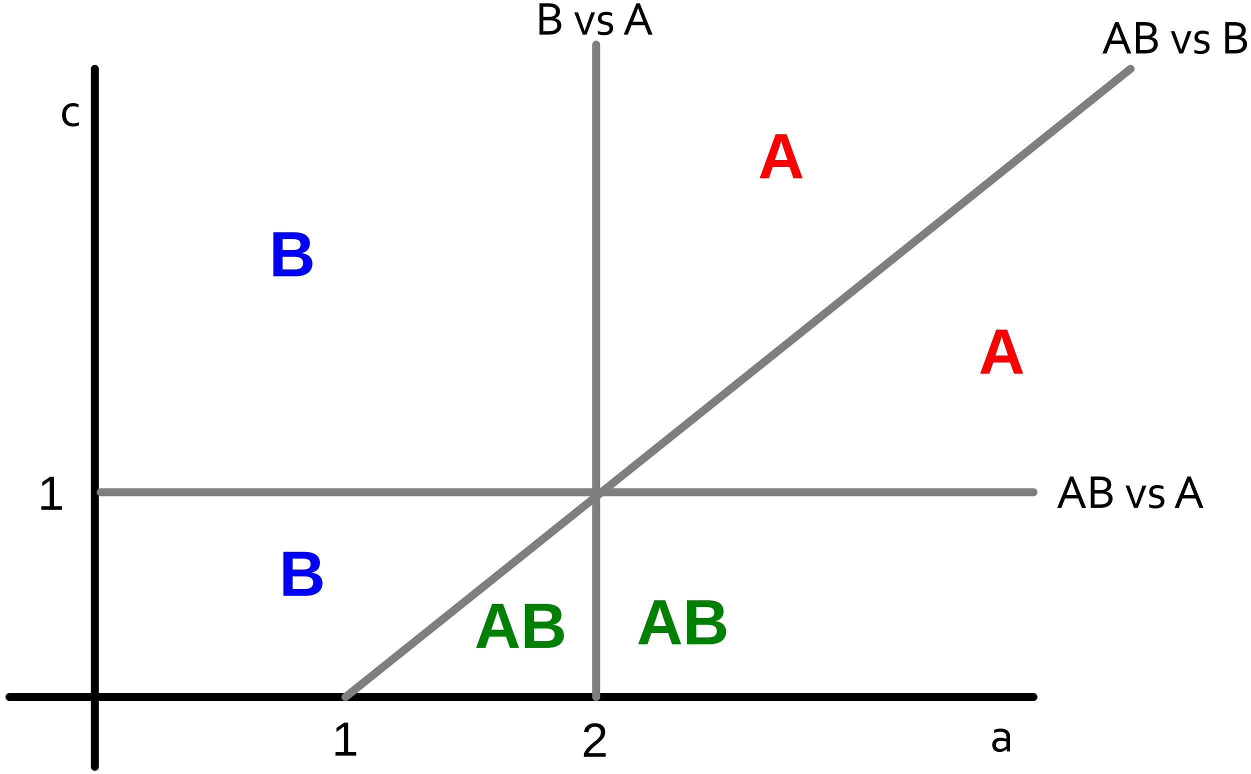

For what pairs (c,a) does A spread?

- Joining the two pictures:

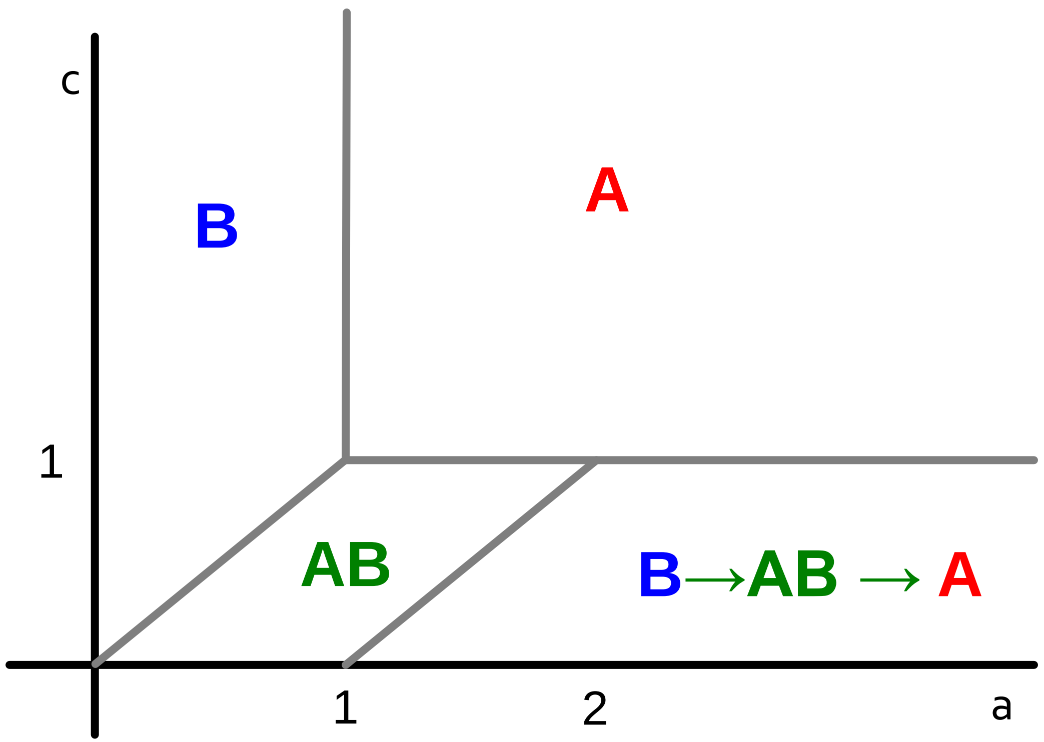

Lesson

- **You manufacture default B and new/better A comes along:**
  - **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

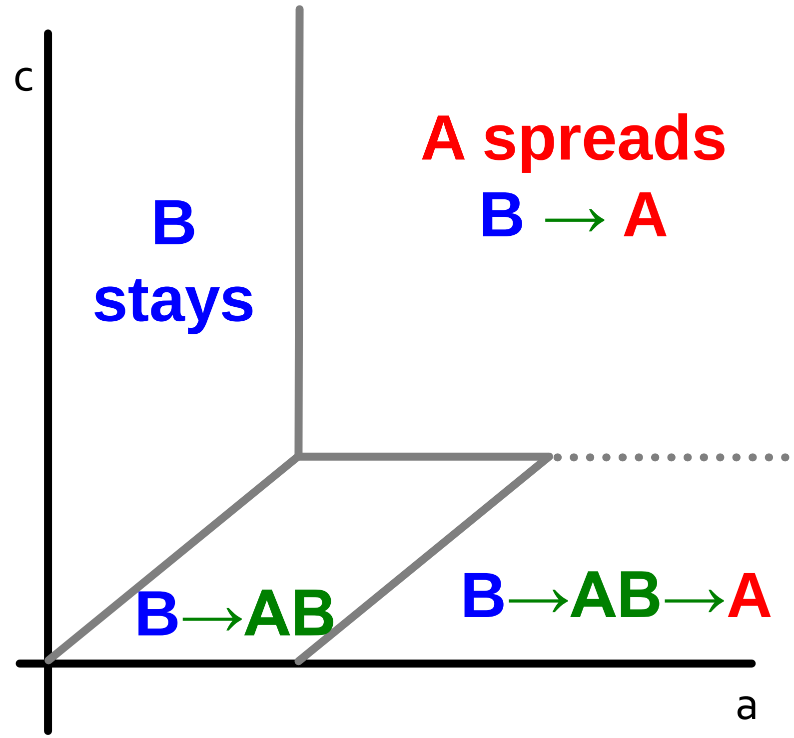

Decision Based Model: Herding

[Banerjee '92]

Decision Based Model: Herding

- **Influence of actions of others**
  - Model where everyone sees everyone else's behavior
- **Sequential decision making**
  - **Example:** Picking a restaurant
    - Consider you are choosing a restaurant in an unfamiliar town
    - Based on Yelp reviews you intend to go to restaurant A
    - But then you arrive there is no one eating at A but the next door restaurant B is nearly full
  - **What will you do?**
    - Information that you can infer from other's choices may be more powerful than your own

Herding: Structure

- **Herding:**
  - There is a decision to be made
  - People make the decision sequentially
  - Each person has some private information that helps guide the decision
  - You can't directly observe private information of the others but can see what they do
    - **You can make inferences about the private information of others**

Herding: Simple Experiment

- Consider an urn with 3 marbles. It can be either:
  - **Majority-blue:** 2 blue, 1 red, **or**
  - **Majority-red:** 1 blue, 2 red
- Each person wants to **best guess** whether the urn is majority-blue or majority-red
  - Guess **red** if P(**majority-red** | what she has seen or heard) > ½
- **Experiment:** One by one each person:
  - Draws a marble
  - **Privately** looks are the color and puts the marble back
  - **Publicly** guesses whether the urn is majority-red or majority-blue
- **You see all the guesses beforehand.**
  How should you make your guess?

[Banerjee '92]

Herding: What Happens?

See ch. 16 of Easley-Kleinberg for formal analysis

- **Informally, What happens?**
  - **#1 person:** Guess the color you draw from the urn.
  - **#2 person:** Guess the color you draw from the urn. **Why?**
    - If same color as 1st, then go with it
    - If different, break the tie by doing with your own color
  - **#3 person:**
    - If the two before made different guesses, go with your color
    - Else, go with their guess (**regardless** your color) – **cascade starts!**
  - **#4 person:**
    - Suppose the first two guesses were **R**, you go with **R**
      - Since 3rd person always guesses **R**
  - Everyone else guesses **R** (regardless of their draw)

Herding: Three Ingredients

- Three ingredients:
  - **State of the world:**
    - Whether the urn is **MR** or **MB**
  - **Payoffs:**
    - Utility of making a correct guess
  - **Signals:**
    - Models private information:
      - The color of the marble that you just draw
    - Models public information:
      - The **MR** vs **MB** guesses of people before you

Sequential Decision Making

- **Decision:** Guess **MR** if $P(\boldsymbol{MR}|past\ actions] > \frac{1}{2}$
- **Analysis (Bayes rule):**
  - **#1 follows her own color (private signal)!**
    - Why? $$P(MR\,|\,\mathrm{r}] = \frac{P(MR)P(r\,|\,MR)}{P(r)} = \frac{1/2\cdot 2/3}{1/2} = 2/3$$
      $$P(r) = P(r\,|\,MB)P(MB) + P(r\,|\,MR)P(MR) = \frac{1}{2}\frac{1}{3} + \frac{1}{2}\frac{2}{3} = 1/2$$
  - **#2 guesses her own color (private signal)!**
    - #2 knows #1 revealed her color. So, #2 gets 2 colors.
    - If they are the same, decision is easy.
    - If not, break the tie in favor of her own color

Sequential Decision Making

- **#3 follows majority signal!**
  - Knows #1, #2 acted on their colors. So, #3 gets 3 signals.
  - If #1 and #2 made opposite decisions, #3 goes with her own color. Future people will know #3 revealed its signal $P(MR \mid r, r, b] = 2/3$
  - If #1 and #2 made same choice, #3's decision conveyed no info. **Cascade has started!**
- **How does this unfold?** You are N-th person
  - **#MB** = **#MR** : you guess your color
  - |**#MB** - **#MR**|=1 : your color makes you indifferent, or reinforces you guess
  - |**#MB** - **#MR**| ≥ 2 : Ignore your signal. Go with majority.

Sequential Decision Making

- Cascade begins when the difference between the number of blue and red guesses reaches 2

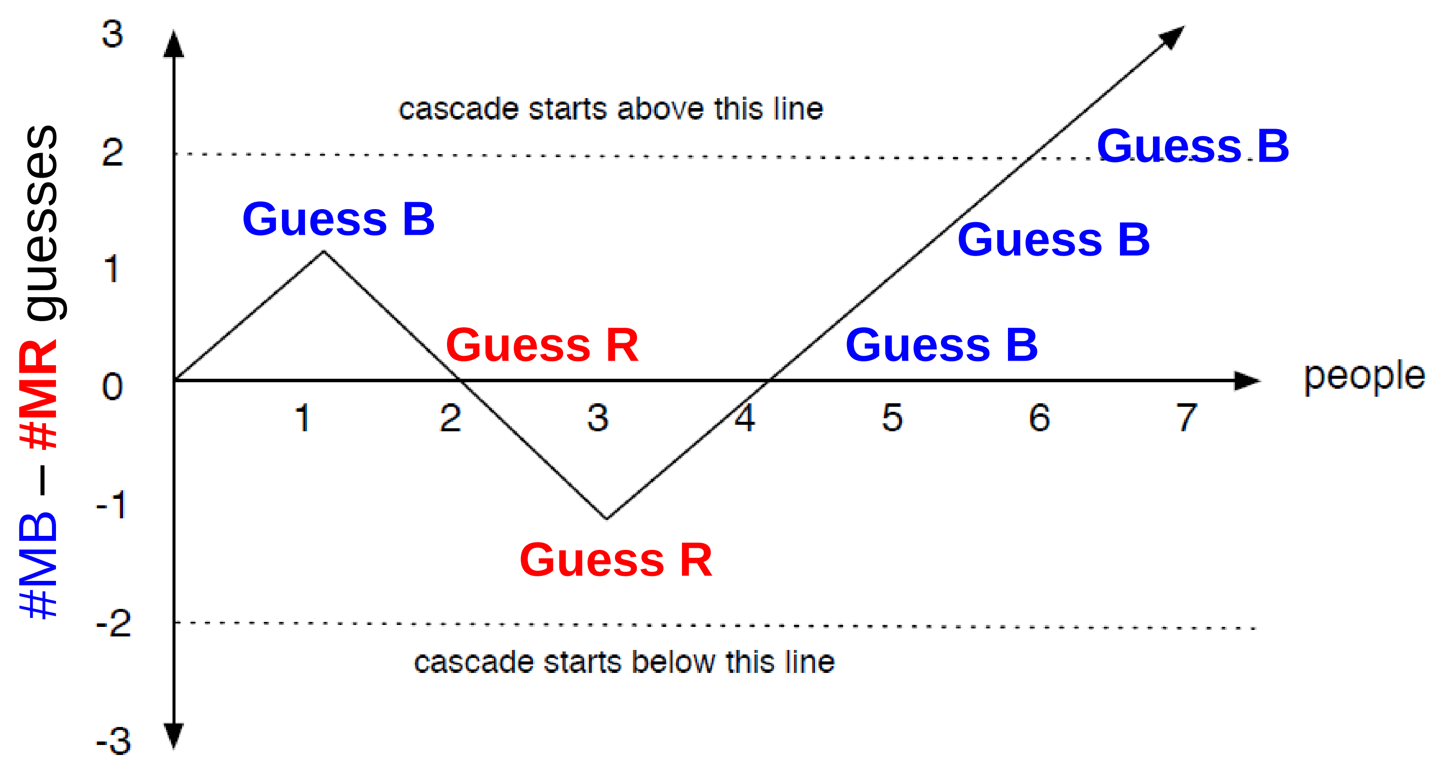

Herding: Observations

- Easy to occur given the right structural conditions
  - Can lead to bizarre patterns of decisions
- Non-optimal outcomes
  - With prob. $\frac{1}{3} \cdot \frac{1}{3} = \frac{1}{9}$ first two see the wrong color, from then on the whole population guesses wrong
- Can be very fragile
  - Suppose first two guess blue
  - People 100 and 101 draw red and cheat by showing their marbles
  - Person 102 now has 4 pieces of information, she guesses based on her own color
  - **Cascade is broken**

Empirical Studies of Cascading Behavior

Modeling Cascading Behavior

- **Basis for models:**
  - **Probability of adopting new behavior depends on the number of friends who have already adopted**
- **What's the dependence?**

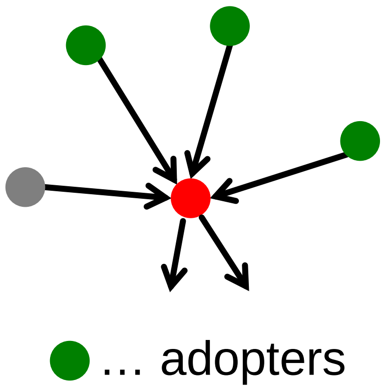

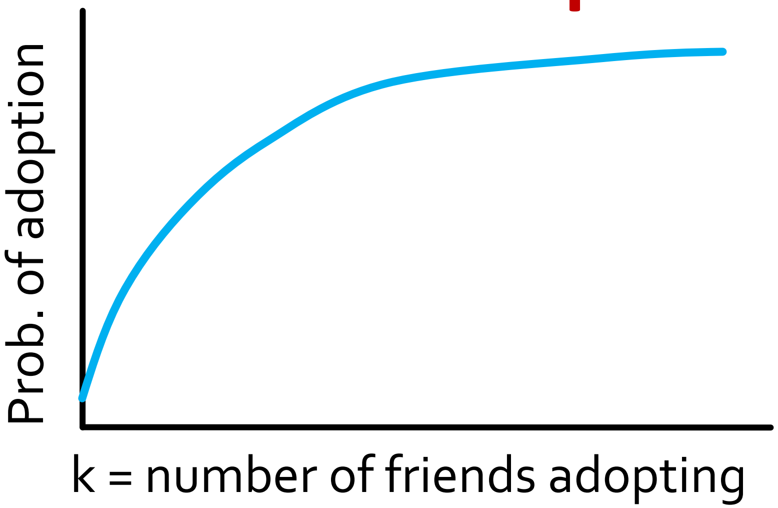

Diminishing returns:
Viruses, Information

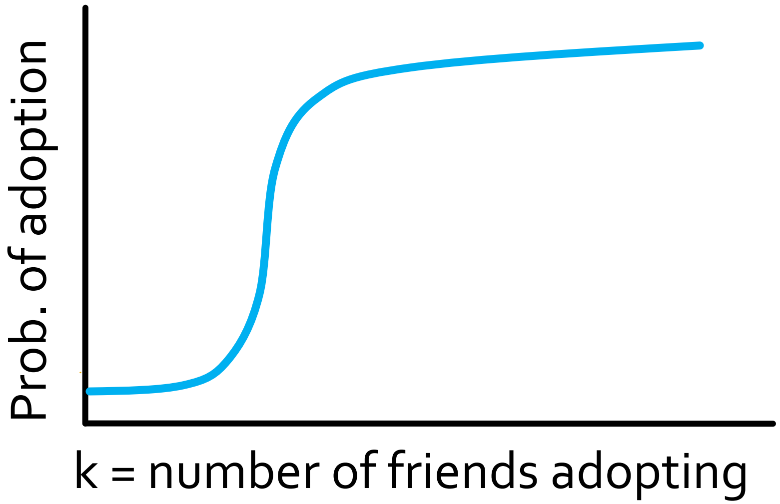

Critical mass:
Decision making

Adoption Curve: LiveJournal

- **Group memberships spread over the network:**
  - Red circles represent existing group members
  - Yellow squares may join
- **Question:**
  - How does prob. of joining a group depend on the number of friends already in the group?

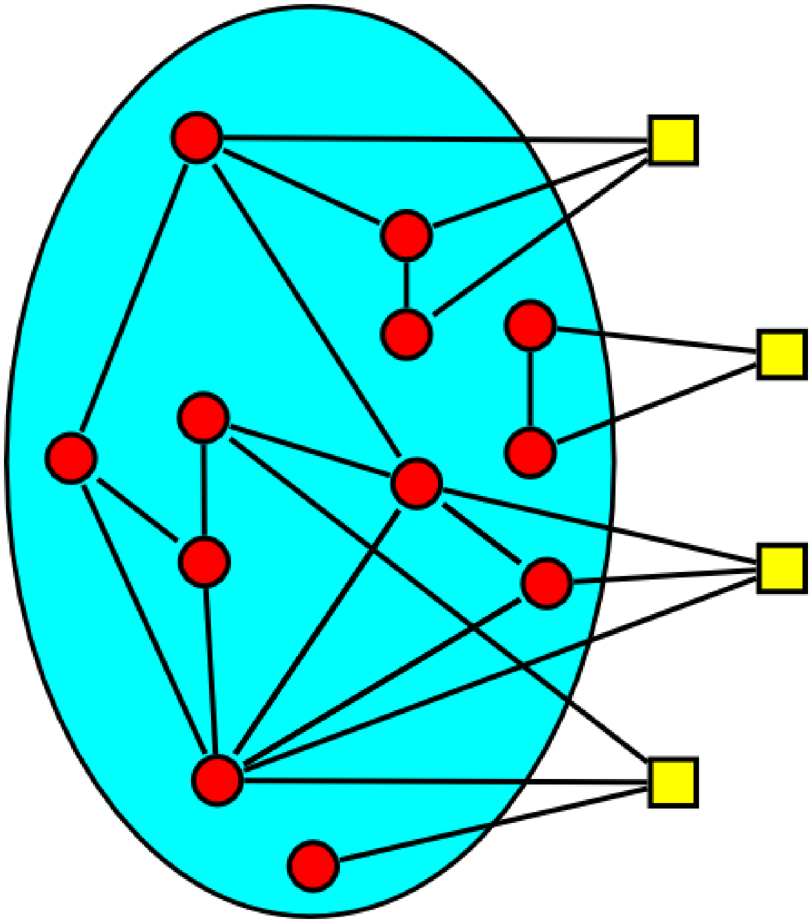

Adoption Curve: LiveJournal

- **LiveJournal group membership**

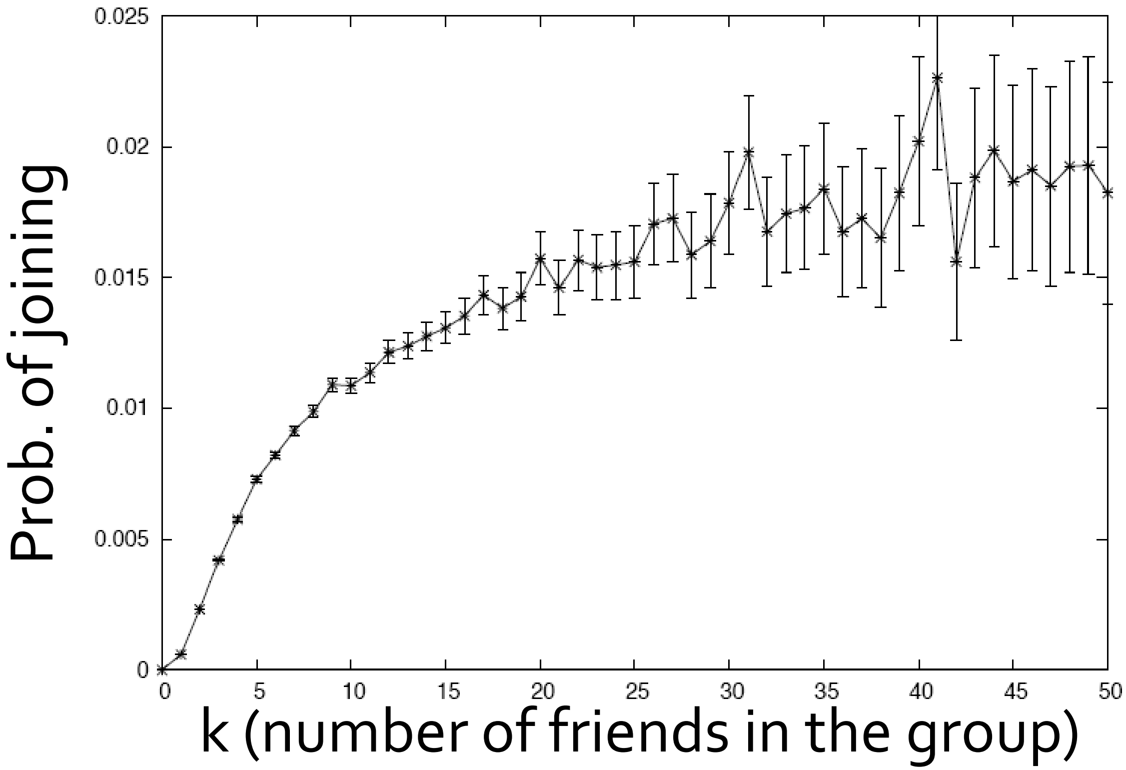

Diffusion in Viral Marketing

- **Senders and followers of recommendations receive discounts on products**

- **Data: Incentivized Viral Marketing program**
  - 16 million recommendations
  - 4 million people, 500k products

Adoption Curve: Validation

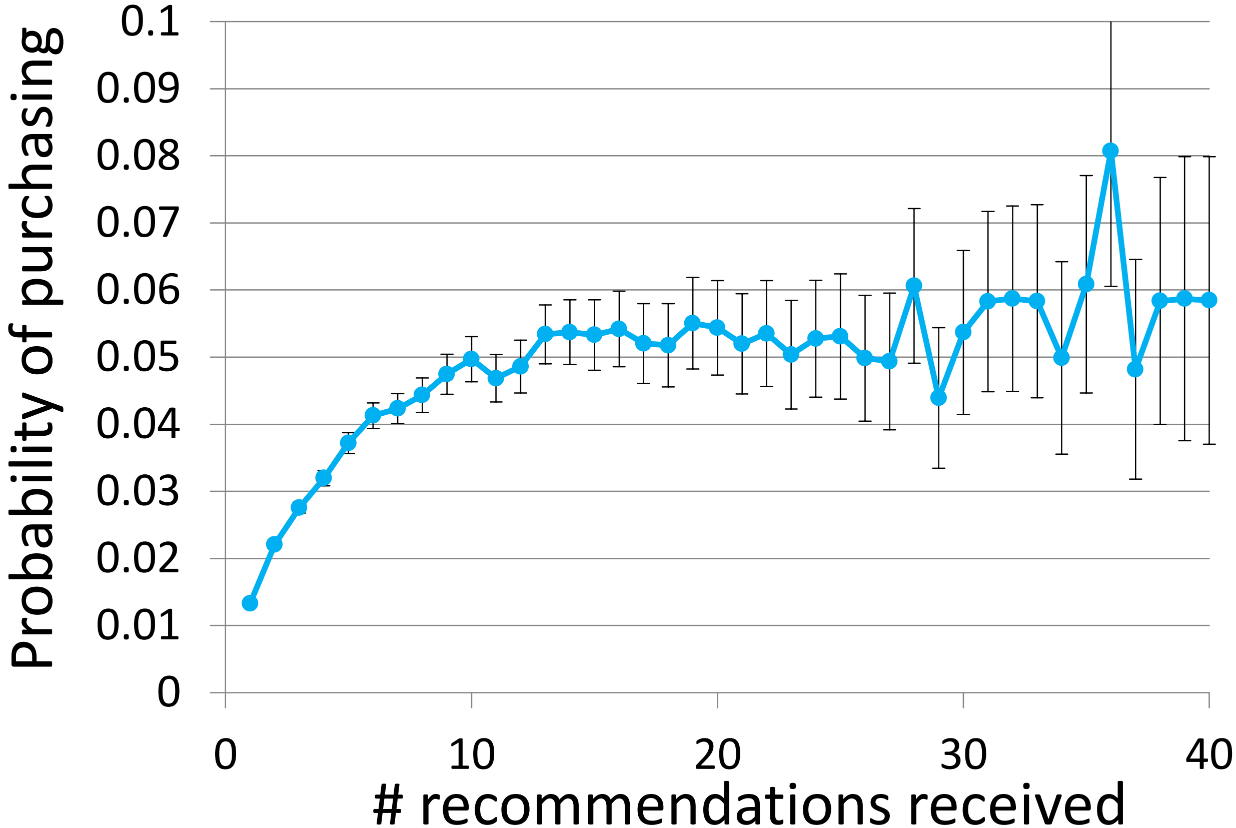

DVD recommendations
(8.2 million observations)

What are We Really Measuring?

- For viral marketing:
  - We see that node v receiving the *i*-th recommendation and then purchased the product
- For groups:
  - At time t we see the behavior of node *v*'s friends
- Good questions:
  - When did *v* become aware of recommendations or friends' behavior?
  - When did it translate into a decision by *v* to act?
  - How long after this decision did *v* act?

Cascading of Product Recommendations & Purchases

Diffusion in Viral Marketing

- Large Anonymous online retailer
(June 2001 to May 2003)
  - 15,646,121 recommendations
  - 3,943,084 distinct customers
  - 548,523 products recommended
  - Products belonging to 4 product groups:
    - Books, DVDs, music, VHS
- **Important:**
  - You can only make recommendations when you buy
  - Only the 1st person to respond to a recommendation gets 10% discount, recommender gets 10% credit

Viral Marketing: Subtle Features

- **What role does the product category play?**

	products	customers	recommenda-tions	edges	buy + get discount	buy + no discount
Book	103,161	**2,863,977**	5,741,611	2,097,809	**65,344**	**17,769**
DVD	**19,829**	805,285	**8,180,393**	962,341	**17,232**	**58,189**
Music	**393,598**	794,148	1,443,847	585,738	7,837	2,739
Video	26,131	239,583	280,270	160,683	909	467
Full	542,719	3,943,084	15,646,121	3,153,676	91,322	79,164

high
low

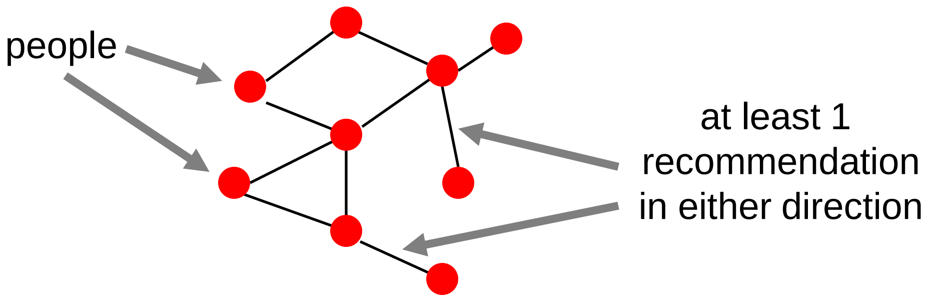

DVD Recommendation Network

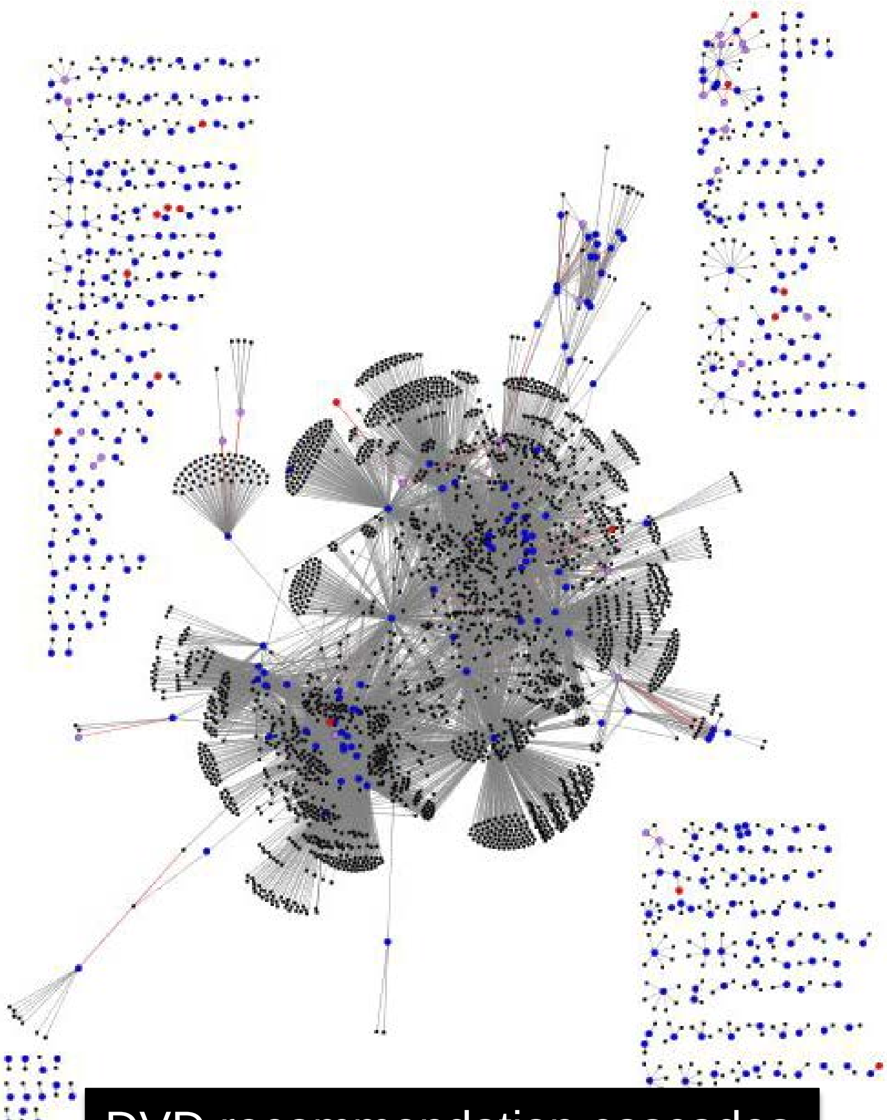

DVD recommendation cascades

- purchase following a recommendation
- customer recommending a product
- customer not buying a recommended product

Observations:

- Majority of recommendations do not cause purchases nor propagation
- Notice many star-like patterns
- Many disconnected components

Cascade Formation Process

- **Recommendations on a single product**
  - Time: $t_1 < t_2 < \ldots < t_n$

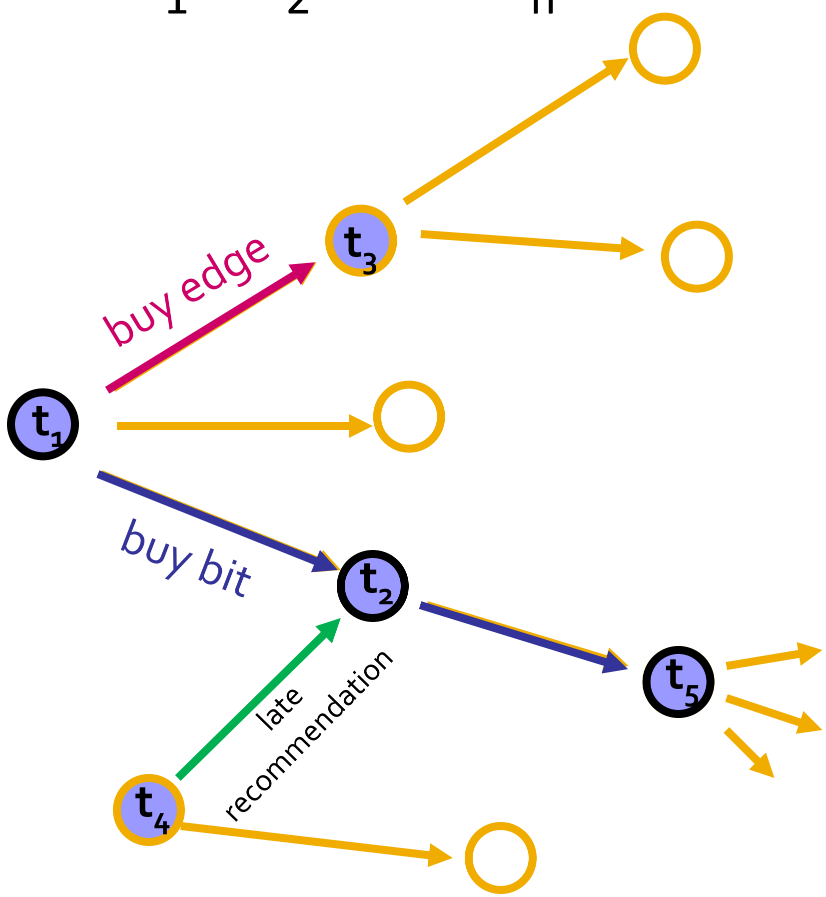

legend

- bought but didn't receive a discount
- bought and received a discount
- received a recommendation but didn't buy

How we know who purchased?

Buy-bit: receiver purchased first (got 10% credit)

Buy-edge: since t_1 recommended to t_3 and t_3 further recommended, t_3 must have purchased

Measuring Cascade Sizes

- **How big are cascades?**
  - Delete late recommendations
  - Count how many people are in a single cascade
  - Exclude nodes that did not buy

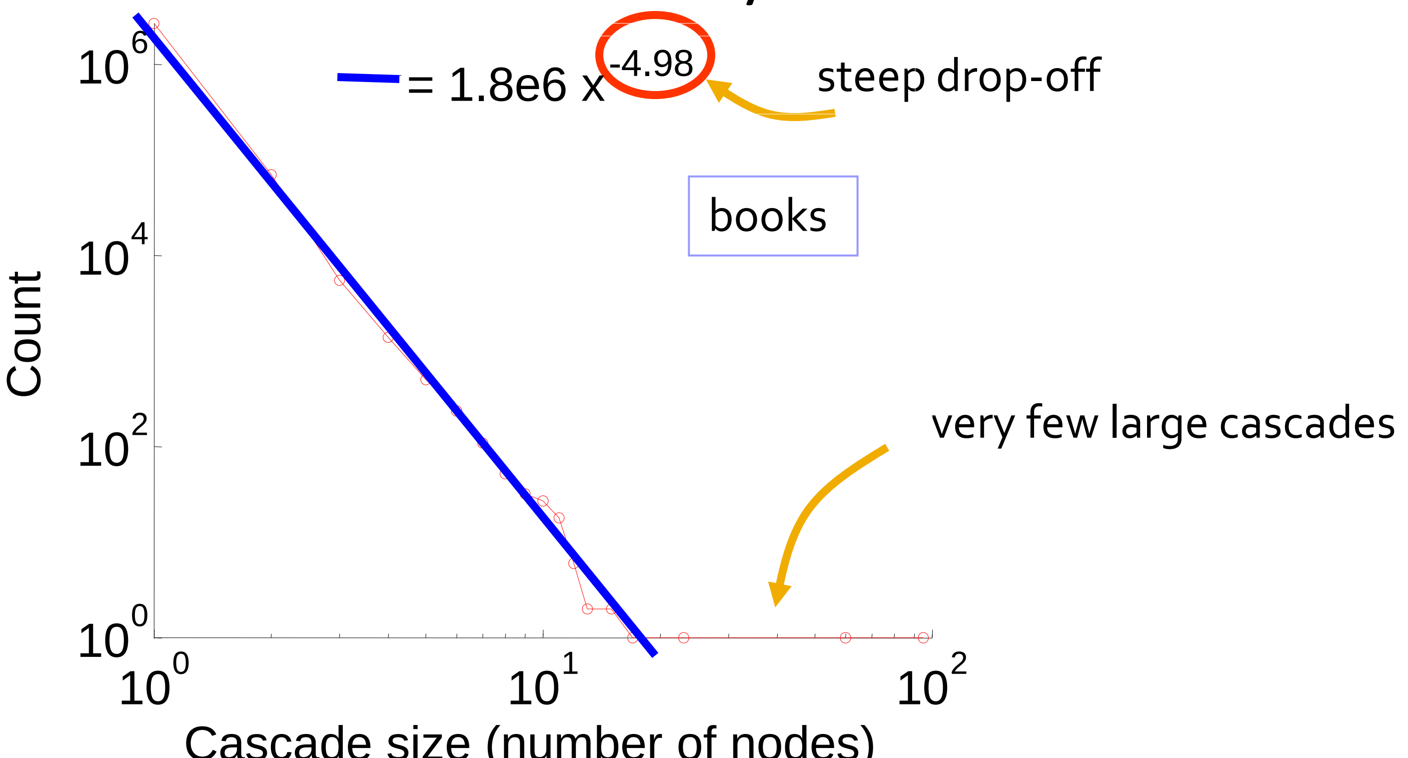

Cascade Size: DVDs

- DVD cascades can grow large
- Possibly as a result of websites where people sign up to exchange recommendations

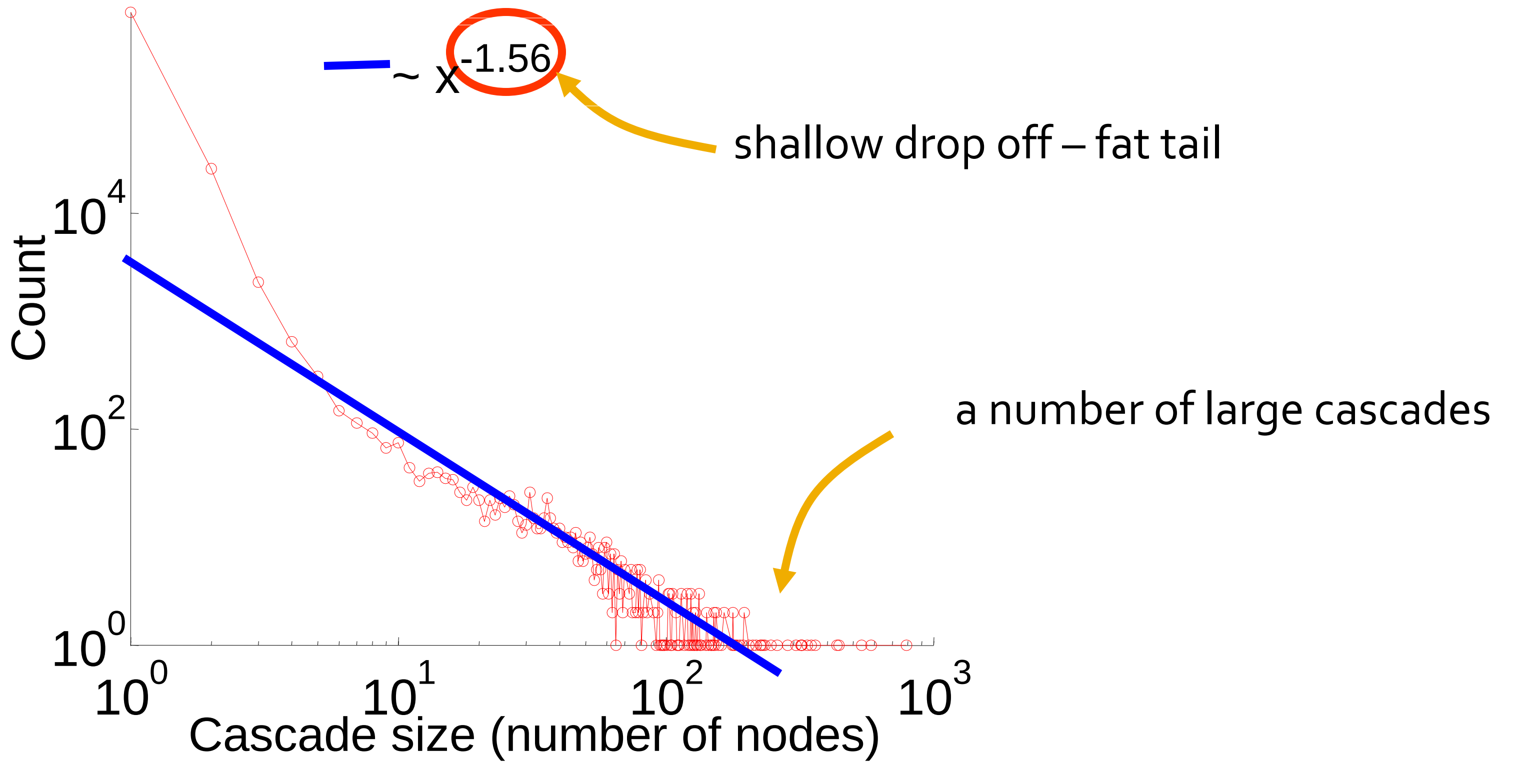

More Subtle Features

- **Does sending more recommendations influence more purchases?**

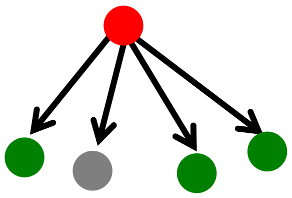

BOOKS

DVDs

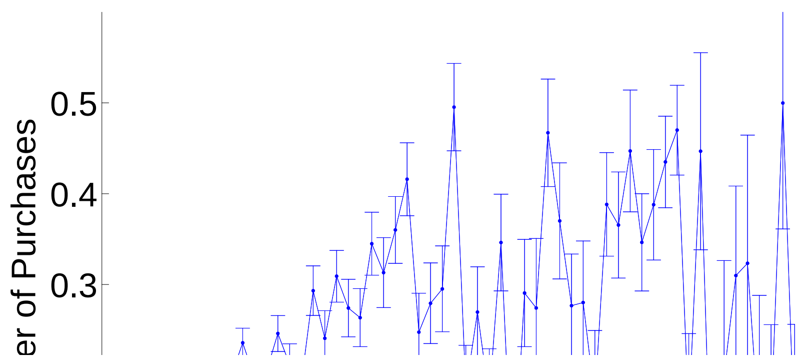

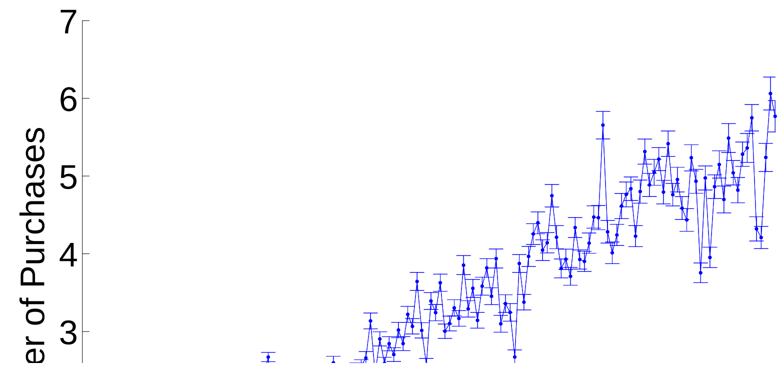

More Subtle Features

- **What is the effectiveness of subsequent recommendations?**

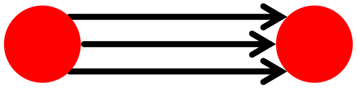

BOOKS

DVDs

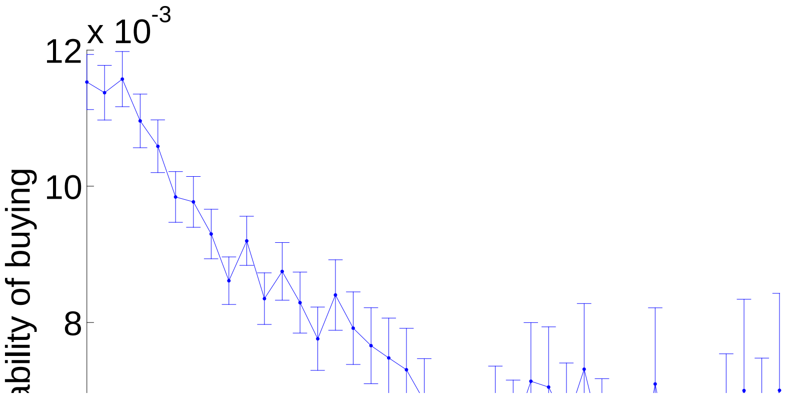

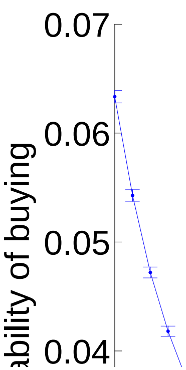

Observations on Product Groups

- We have relatively few DVD titles, but DVDs account for ~ 50% of all recommendations
- **Recommendations per person**
  - DVD: 10
  - books and music: 2
  - VHS: 1
- **Recommendations per purchase**
  - books: 69
  - DVDs: 108
  - music: 136
  - VHS: 203
- **Overall there are 3.69 recommendations per node on 3.85 different products**
- Music recommendations reached about the same number of people as DVDs but used only 20% as many recommendations
- Book recommendations reached by far the most people – 2.8 million
- All networks have a very small number of unique edges
  - For books, videos and music the number of unique edges is smaller than the number of nodes – the networks are highly disconnected

Product Characteristics

- consider successful recommendations in terms of
  - av. # senders of recommendations per book category
  - av. # of recommendations accepted
- books overall have a 3% success rate
  - (2% with discount, 1% without)
- lower than average success rate (significant at p=0.01 level)
  - fiction
    - romance (1.78), horror (1.81)
    - teen (1.94), children's books (2.06)
    - comics (2.30), sci-fi (2.34), mystery and thrillers (2.40)
  - nonfiction
    - sports (2.26)
    - home & garden (2.26)
    - travel (2.39)
- higher than average success rate (statistically significant)
  - professional & technical
    - medicine (5.68)
    - professional & technical (4.54)
    - engineering (4.10), science (3.90), computers & internet (3.61)
    - law (3.66), business & investing (3.62)

Anime DVDs

- 47,000 customers responsible for the 2.5 out of 16 million recommendations in the system
- 29% success rate per recommender of an anime DVD
- Giant component covers 19% of the nodes
- Overall, recommendations for DVDs are more likely to result in a purchase (7%), but the anime community stands out

Predicting Recommendation Success

Variable	transformation	Coefficient
const		-0.940 ***
# recommendations	$\ln(r)$	0.426 ***
# senders	$\ln(n_s)$	-0.782 ***
# recipients	$\ln(n_r)$	-1.307 ***
product price	$\ln(p)$	0.128 ***
# reviews	$\ln(v)$	-0.011 ***
avg. rating	$\ln(t)$	-0.027 *
R^2		**0.74**

significance at the 0.01 (***), 0.05 (**) and 0.1 (*) levels

Viral Marketing: Not spreading virally

- 94% of users make first recommendation without having received one previously
- Size of giant connected component increases from 1% to 2.5% of the network (100,420 users) – **small!**
- **Some sub-communities are better connected**
  - 24% out of 18,000 users for westerns on DVD
  - 26% of 25,000 for classics on DVD
  - 19% of 47,000 for anime (Japanese animated film) on DVD
- **Others are just as disconnected**
  - 3% of 180,000 home and gardening
  - 2-7% for children's and fitness DVDs

Viral Marketing: Consequences

Products suited for Viral Marketing:

- small and tightly knit community
  - few reviews, senders, and recipients
  - but sending more recommendations helps
- pricey products
- rating doesn't play as much of a role

Observations for future diffusion models:

- purchase decision more complex than threshold or simple infection
- influence saturates as the number of contacts expands
- links user effectiveness if they are overused

Conditions for successful recommendations:

- professional and organizational contexts
- discounts on expensive items
- small, tightly knit communities