Networks with Signed Edges

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

Announcement: Reaction paper

- **The idea of the reaction papers is:**
  - Familiarize yourselves more in depth with the class material
  - Do reading beyond what was covered
  - **You should be thinking beyond what you read, and not just take other people's work for granted**
  - **Think of the paper as a way to start thinking about the project**
- **Read at 2 to 3 papers:**
  - Anything from course site, last year's site, Easley-Kleinberg,…
- **Logistics:**
  - **Due in 1 week: Oct 20 in class!**
  - **Can be done in groups of 2-3 students**
  - **How to submit:**
    - Paper copy in a box **AND** upload to HW submission site
      - **Use the homework cover sheet**
  - See http://www.stanford.edu/class/cs224w/info.html for more info and examples of old reaction papers

Announcement: Reaction paper

- **On 3-5 pages answer the following questions:**
  - **1 page: Summary**
    - What is main technical content of the papers?
    - How do papers relate to the topics presented in the course?
    - What is the connection between the papers you are discussing?
  - **1 page: Critique**
    - What are strengths and weaknesses of the papers and how they be addressed?
    - What were the authors missing?
    - Was anything particularly unrealistic?
  - **1 page: Brainstorming**
    - What are promising further research questions in the direction of the papers?
    - How could they be pursued?
    - An idea of a **better model** for something? A **better algorithm**? A **test of a model or algorithm** on a dataset or simulated data?

Recap: Signed Networks

- Networks with positive and negative links
- Structure of **signed triangles**
  - **Structural balance:**

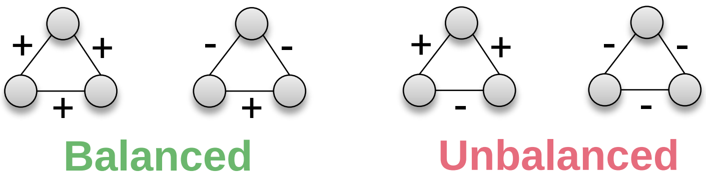

  - **Status theory:**
    - $A \xrightarrow{+} B$ :: B has higher status than A
    - $A \xrightarrow{-} B$ :: B has lower status than A
- How to compare the two theories?
  - Triads provide **context**
  - **Surprise:** Change in behavior of A/B when we know the context

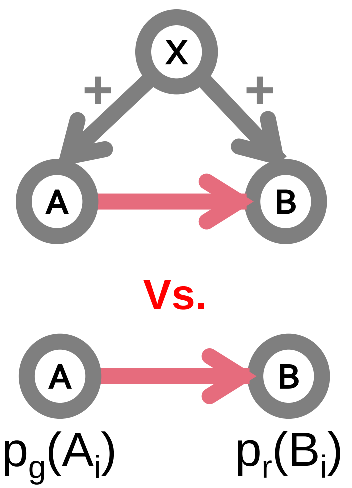

$$s_g(t) = \frac{k - \sum_{i=1}^{n} p_g(A_i)}{\sqrt{\sum_{i}^{n} p_g(A_i)(1 - p_g(A_i))}}$$

$p_g(A_i)$… generative baseline of A_i
$p_r(B_i)$… receptive baseline of B_i

Status: Two Examples

- **Two basic examples:**

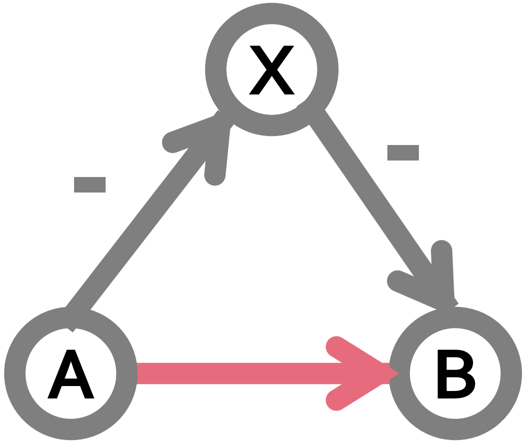

Gen. surprise of A: —
Rec. surprise of B: —

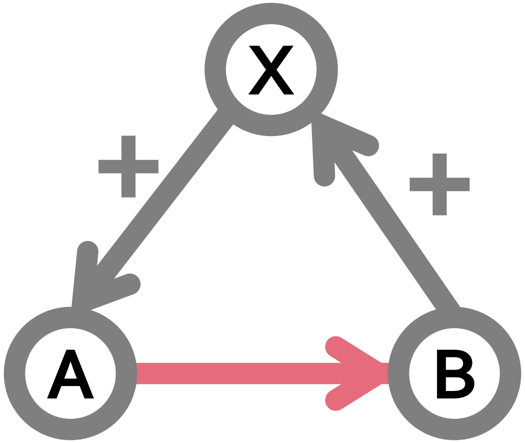

Gen. surprise of A: —
Rec. surprise of B: —

Joint Positive Endorsement

- X positively endorses A and B
- Now A links to B

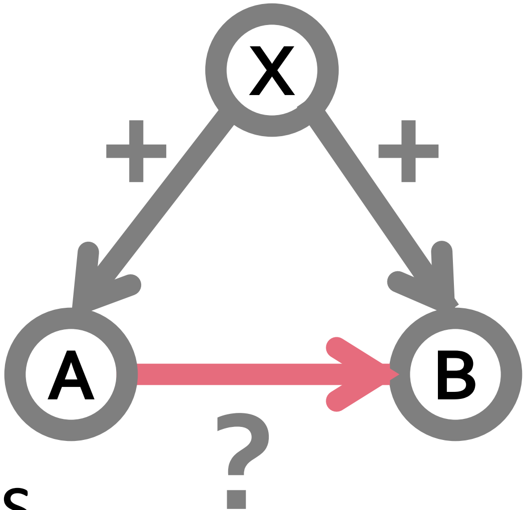

A puzzle:

- In our data we observe:
  Fraction of positive links deviates
  - Above generative baseline of A
  - Below receptive baseline of B

- Why?

A Story: Soccer Team

- Ask every node: How does skill of B compare to yours?
  - Build a signed directed network
- We haven't asked A about B
- But we know that X thinks A and B are both better than him
- What can we infer about A's answer?

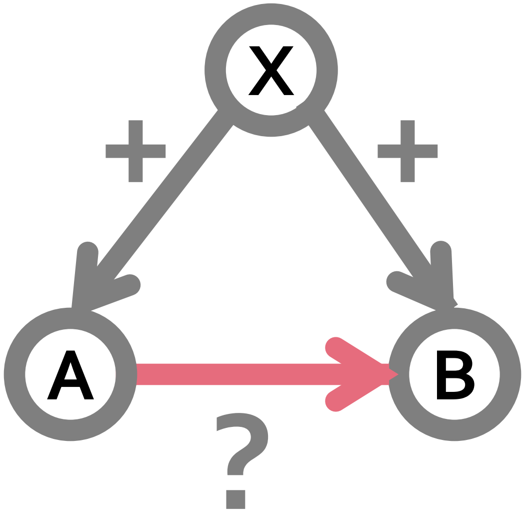

A Story: Soccer Team

- A's viewpoint:
  - Since B has positive evaluation, B is high status
  - Thus, evaluation A gives is more likely to be positive than the baseline

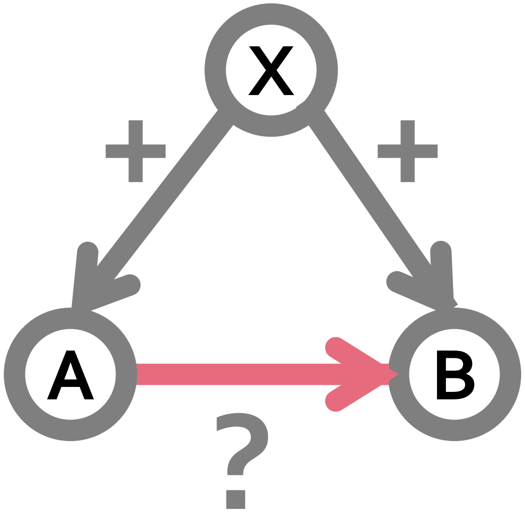

How does A evaluate B?

A is evaluating someone who is better than avg.
→ A is more positive than average

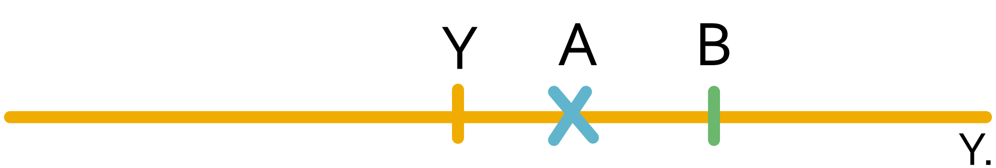

Y… average node

A Story: Soccer Team

- B's viewpoint:
  - Since A has positive evaluation, A is high status
  - Thus, evaluation B receives is less likely to be positive than the baseline

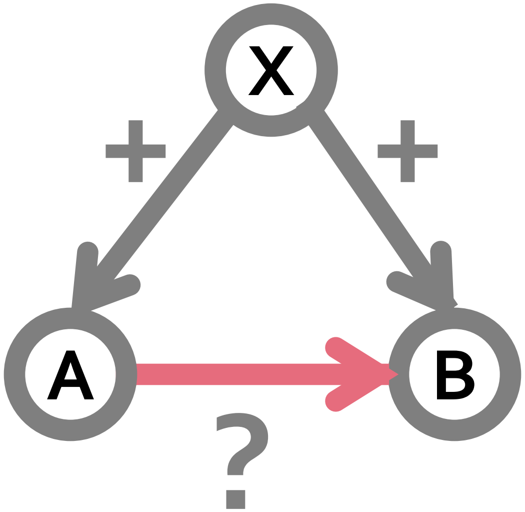

How is B evaluated by A?

B is evaluated by someone better than average.
→ They will be more negative to B than average

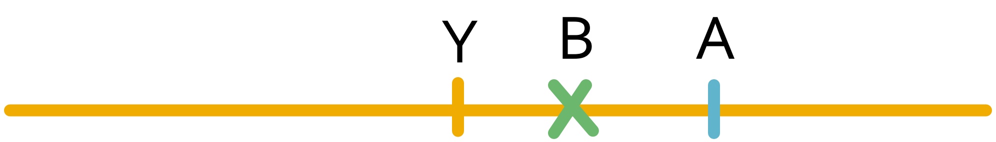

Y… average node

Consistency with Status

- **Determine node status:**
  - Assign X status 0
  - Based on signs and directions of edges set status of A and B
- Surprise is **status**-consistent, if:
  - Gen. surprise is status-consistent if it has **same** sign as status of B
  - Rec. surprise is status-consistent if it has the **opposite** sign from the status of A
- Surprise is **balance**-consistent, if:
  - If it completes a balanced triad

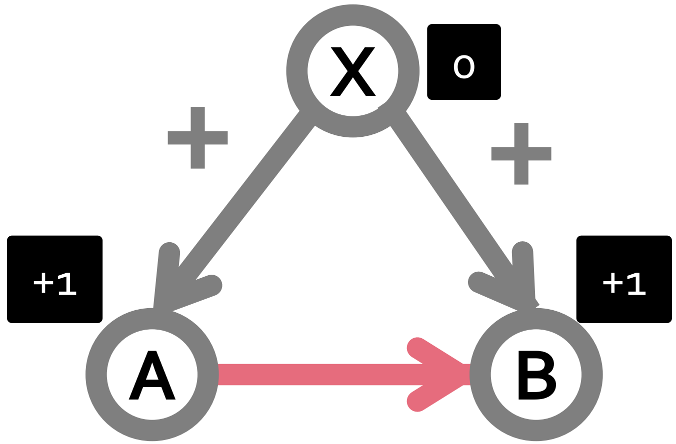

Status-consistent if:
Gen. surprise > 0
Rec. surprise < 0

Status vs. Balance (Epinions)

- **Predictions:**

t_i	count	$P(+)$	$S_g(t_i)$	$S_r(t_i)$	B_g	B_r	S_g	S_r
t_1	178,051	0.97	95.9	197.8	✓	✓	✓	✓
t_2	45,797	0.54	-151.3	-229.9	✓	✓	✓	●
t_3	246,371	0.94	89.9	195.9	✓	✓	●	✓
t_4	25,384	0.89	1.8	44.9	○	○	✓	✓
t_5	45,925	0.30	18.1	-333.7	○	✓	✓	✓
t_6	11,215	0.23	-15.5	-193.6	○	○	✓	✓
t_7	36,184	0.14	-53.1	-357.3	✓	✓	✓	✓
t_8	61,519	0.63	124.1	-225.6	✓	○	✓	✓
t_9	338,238	0.82	207.0	-239.5	✓	○	✓	✓
t_{10}	27,089	0.20	-110.7	-449.6	✓	✓	✓	✓
t_{11}	35,093	0.53	-7.4	-260.1	○	○	✓	✓
t_{12}	20,933	0.71	17.2	-113.4	○	✓	✓	✓
t_{13}	14,305	0.79	23.5	24.0	○	○	✓	✓
t_{14}	30,235	0.69	-12.8	-53.6	○	○	✓	●
t_{15}	17,189	0.76	6.4	24.0	○	○	●	✓
t_{16}	4,133	0.77	11.9	-2.6	✓	○	✓	●
Number of correct predictions					8	7	14	13

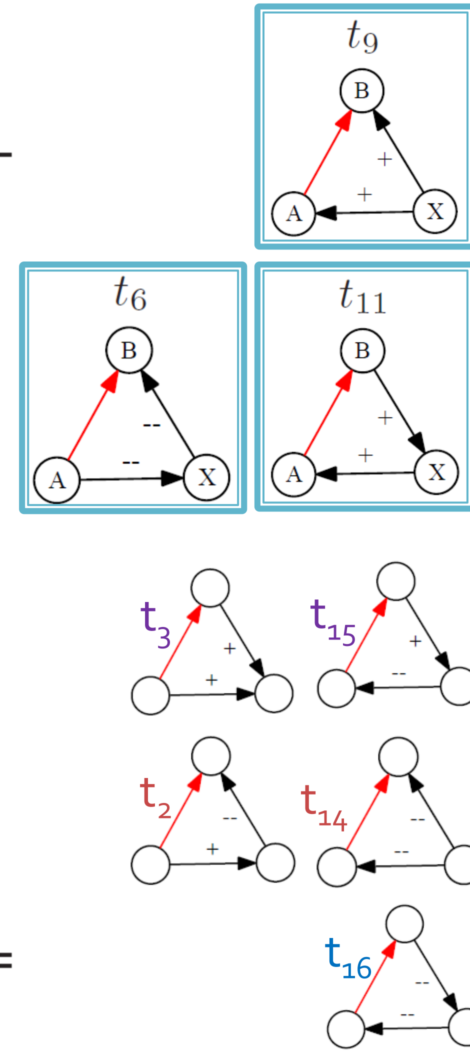

From Local to Global Structure

[WWW '10]

- **Both theories make predictions about the global structure of the network**

- Structural balance – Factions
  - Find coalitions
- Status theory – Global Status
  - Flip direction and sign of minus edges
  - Assign each node a unique status so that edges point from low to high

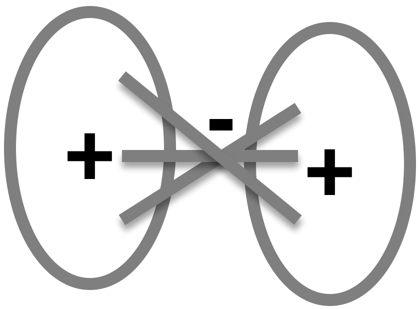

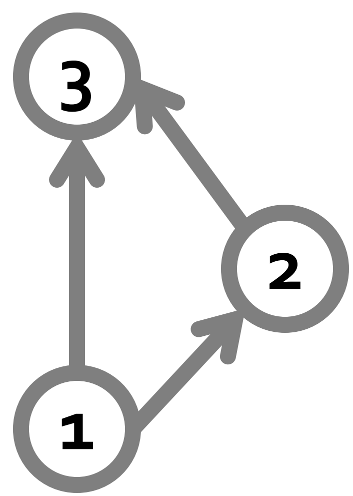

[WWW '10]

From Local to Global Structure

- Fraction of edges of the network that satisfy Balance and Status?

- Observations:
  - No evidence for global balance beyond the random baselines
    - Real data is 80% consistent vs. 80% consistency under random baseline
  - Evidence for global status beyond the random baselines
    - Real data is 80% consistent, but 50% consistency under random baseline

Predicting Edge Signs

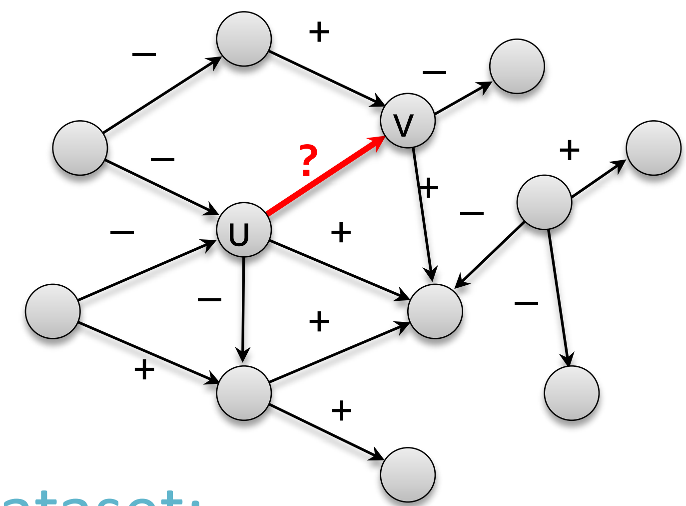

Edge sign prediction problem

- Given a network and signs on all but one edge, predict the missing sign

Machine Learning Formulation:

- Predict sign of edge (u,v)
- Class label:
  - +1: positive edge
  - -1: negative edge
- Learning method:
  - Logistic regression

$$P(+|x) = \frac{1}{1+e^{-(b_0+\sum_i^n b_i x_i)}}$$

- Dataset:
  - Original: 80% **+**edges
  - **Balanced**: 50% **+**edges
- Evaluation:
  - Accuracy
- Features for learning:
  - Next slide

Features for Learning

For each edge (u,v) create features:

- Triad counts (16):
  - Counts of signed triads edge $u \rightarrow v$ takes part in
- Node degree (7 features):
  - Signed degree:
    - $d^+_{out}(u)$, $d^-_{out}(u)$, $d^+_{in}(v)$, $d^-_{in}(v)$
  - Total degree:
    - $d_{out}(u)$, $d_{in}(v)$
  - Embeddedness of edge (u,v)

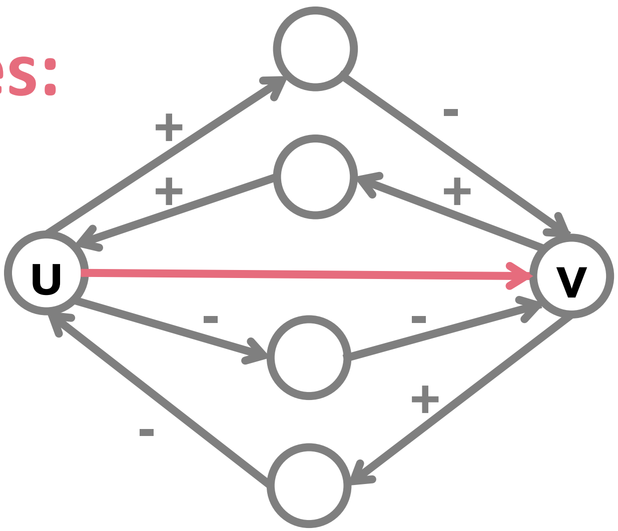

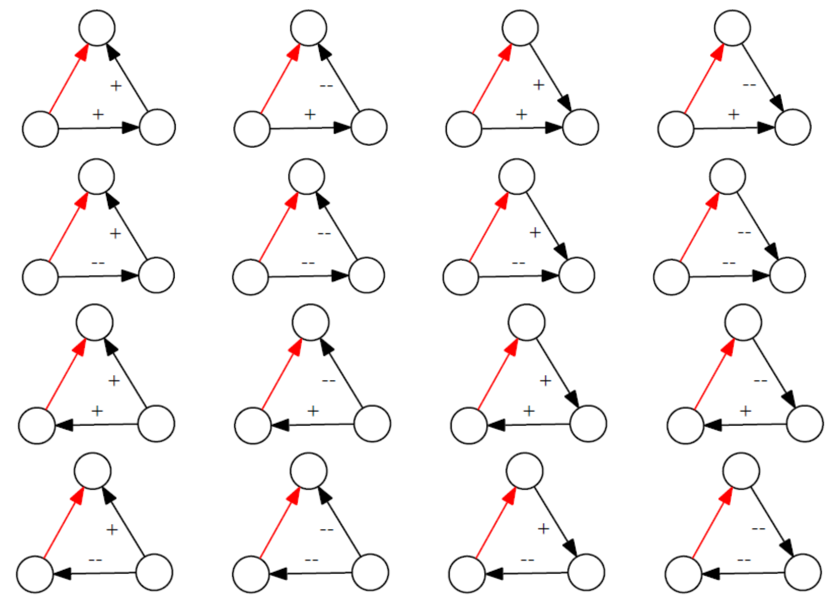

Edge Sign Prediction

[WWW '10]

- **Classification Accuracy:**
  - Epinions: 93.5%
  - Slashdot: 94.4%
  - Wikipedia: 81%
- **Signs can be modeled from local network structure alone**
  - Trust propagation model of [Guha et al. '04] has 14% error on Epinions
- **Triad features perform less well for less embedded edges**
- Wikipedia is harder to model:
  - Votes are publicly visible

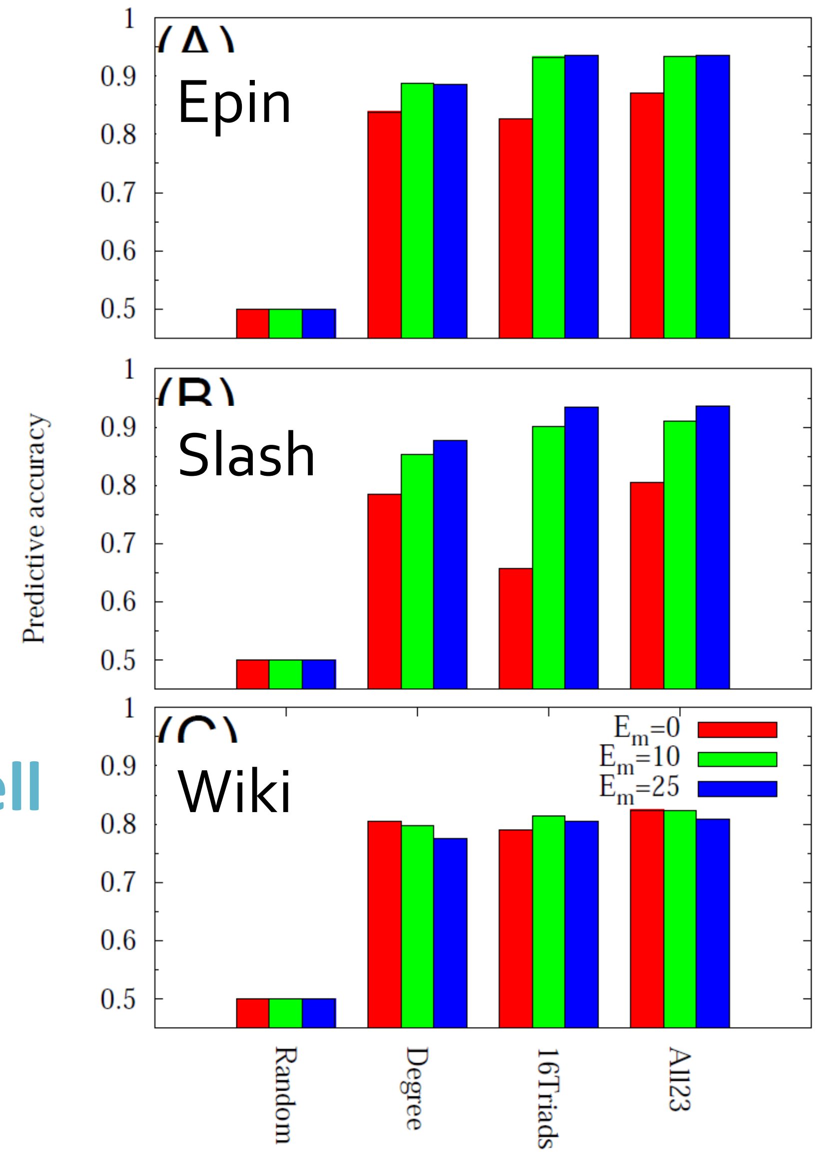

Balance and Status: Complete Model

Feature	Bal	Stat	**Epin**	**Slashd**	**Wikip**
const			-0.2	0.02	-0.2
$\bullet \xrightarrow{+} \bullet \xrightarrow{+} \bullet$	1	1	**0.5**	**0.9**	**0.3**
$\bullet \xrightarrow{+} \bullet \xrightarrow{-} \bullet$	-1	0	**-0.5**	**-0.9**	**-0.4**
$\bullet \xrightarrow{-} \bullet \xrightarrow{+} \bullet$	-1	0	**-0.4**	**-1.1**	**-0.3**
$\bullet \xrightarrow{-} \bullet \xrightarrow{-} \bullet$	1	-1	**-0.7**	**-0.6**	**-0.8**
$\circ \xrightarrow{+} \circ \xleftarrow{+} \circ$	1	0	0.3	0.4	0.05
$\circ \xrightarrow{+} \circ \xleftarrow{-} \circ$	-1	1	-0.01	-0.1	-0.01
$\bullet \xrightarrow{-} \bullet \xleftarrow{+} \bullet$	-1	-1	**-0.9**	**-1.2**	**-0.2**
$\circ \xrightarrow{-} \circ \xleftarrow{-} \circ$	1	0	0.04	-0.07	-0.03
$\circ \xleftarrow{+} \circ \xrightarrow{+} \circ$	1	0	0.08	0.4	0.1
$\bullet \xleftarrow{+} \bullet \xrightarrow{-} \bullet$	-1	-1	**-1.3**	**-1.1**	**-0.4**
$\circ \xleftarrow{-} \circ \xrightarrow{+} \circ$	-1	1	-0.1	-0.2	0.05
$\circ \xleftarrow{-} \circ \xrightarrow{-} \circ$	1	0	0.08	-0.02	-0.1
$\circ \xleftarrow{+} \circ \xleftarrow{+} \circ$	1	-1	-0.09	-0.09	-0.01
$\circ \xleftarrow{+} \circ \xleftarrow{-} \circ$	-1	0	-0.05	-0.3	-0.02
$\circ \xleftarrow{-} \circ \xleftarrow{+} \circ$	-1	0	-0.04	-0.3	0.05
$\circ \xleftarrow{-} \circ \xleftarrow{-} \circ$	1	1	-0.02	0.2	-0.2

Generalization

- Do people use these very different linking systems by obeying the same principles?
  - How generalizable are the results across the datasets?
    - Train on row "dataset", predict on "column"

All23	Epinions	Slashdot	Wikipedia
Epinions	0.9342	0.9289	0.7722
Slashdot	0.9249	0.9351	0.7717
Wikipedia	0.9272	0.9260	0.8021

- Nearly **perfect generalization** of the models even though networks come from very different applications

Important Points

- **Signed networks provide insight into how social computing systems are used:**
  - Status vs. Balance
  - Role of embeddedness and public display
- Sign of relationship can be reliably predicted from the local network context
  - ~90% accuracy sign of the edge

Important Points

- More evidence that networks are globally organized based on status
- People use signed edges consistently regardless of particular application
  - Near perfect generalization of models across datasets

How do People Evaluate Others?

People Express Opinions

People express positive and negative attitudes/opinions:

- Through actions:
  - Rating a product
  - Pressing "like" button
- Through text:
  Sentiment analysis
  [Pang-Lee '08]
  - Writing a comment, a review

People Express Opinions

- **About items:**
  - Movie and product reviews

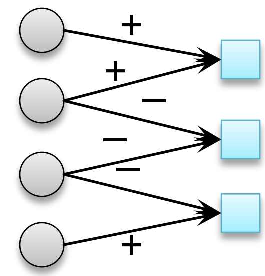

- **About other users:**
  - Online communities

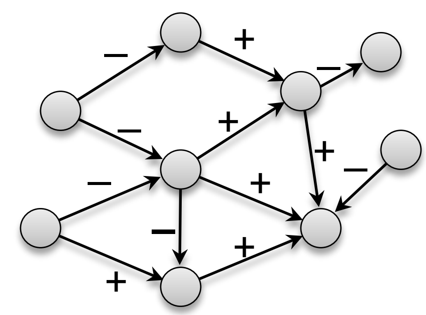

- **About items created by others:**
  - Q&A websites

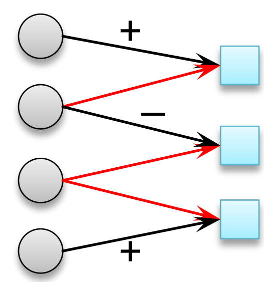

Users Evaluating Others

- Any user A can evaluate any user B:

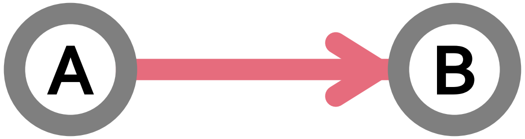

  - Positive (+) vs. negative (–) evaluation
- **Data:**
  - **Users to users:**
    - Epinions: Does A trust B's product reviews?
    - Wikipedia: Does A support B to become Wiki admin?
  - **Users to items:**
    - StackOverflow: Up/down vote (6M votes):
      Does A think B contributed a good answer?

Relative vs. Absolute Assessment

- How do properties of **evaluator A** and **target B** affect A's vote?

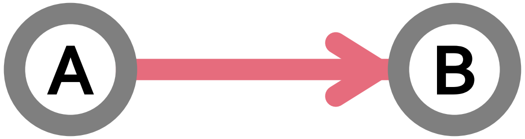

- Two natural (but competing) hypotheses:
  - (1) Prob. that B receives a positive evaluation depends primarily on the characteristics of B
    - There is some objective criteria for a user to receive a positive evaluation

Relative vs. Absolute Assessment

- How do properties of **evaluator A** and **target B** affect A's vote?

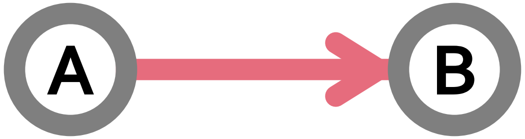

- Two natural (but competing) hypotheses:
  - (2) Prob. that B receives a positive evaluation depends on relationship between characteristics of A and B
    - **Similarity:** Prior interaction between A and B
    - **Status:** A compares status of B to her own status

Status (Level of Contribution)

Ways to quantify status (seniority, merit) of a user:

- Total number of edits of a user:
  - The more edits the user made the higher status she has
- Total number of answers of a user:
  - The more answers given by the user the higher status she has

Status: How to Model?

- **How does the prob. of A evaluating positively depend on the status of A and status of B?**

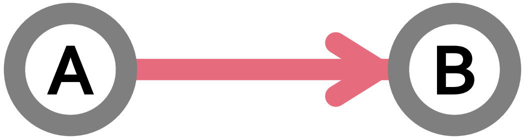

  - Model it as a function of status S_A of A and S_B of B separately?
  - Model as the status difference S_A-S_B?
  - Model as the status ratio S_A/S_B?

Status: Relative Assessment (1)

- **How does status of B affect A's evaluation?**
  - Each curve is fixed status difference: $\Delta = S_A - S_B$
- **Observations:**
  - Flat curves: Prob. of positive evaluation doesn't depend on B's status
  - Different levels: Different values of Δ result in different behavior

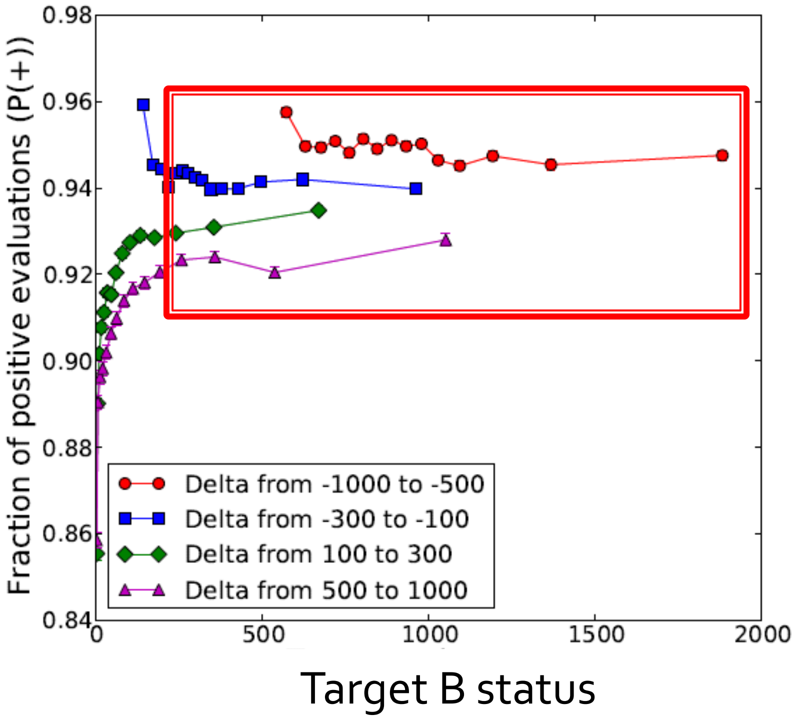

Status difference remains salient even as A and B acquire more status

Status: Relative Assessment (2)

- How does status of B affect A's evaluation?
  - Each curve is fixed status difference: $\Delta = S_A - S_B$
- Observations:
  - Below some threshold targets are judged based on their absolute status
    - And independently of evaluator's status

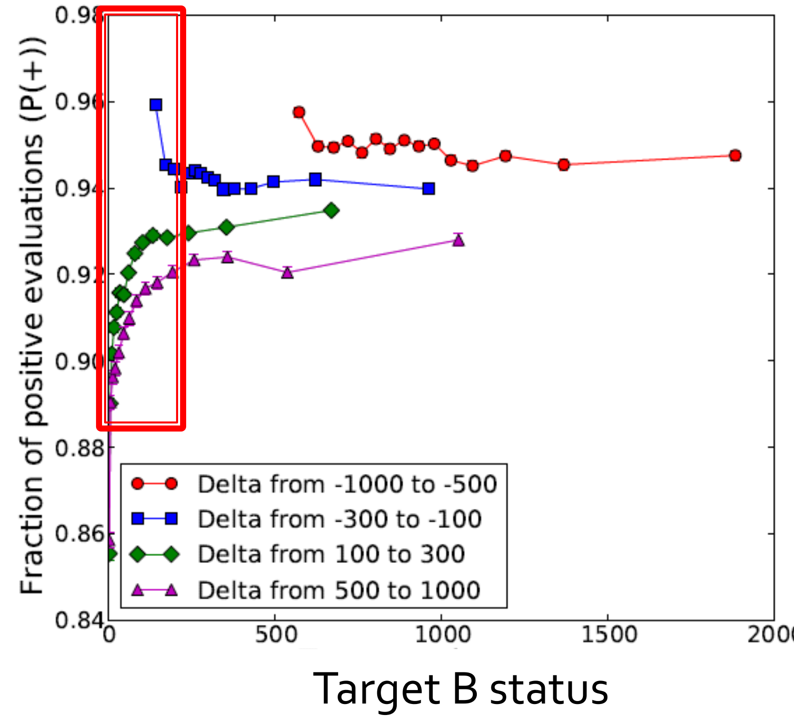

Low-status targets are evaluated based on absolute status

Effects of Similarity

- **How does prior interaction shape evaluations?**
  - (1) Evaluators are more supportive of targets in their area
  - (2) More familiar evaluators know weaknesses and are more harsh
- **Observation:**
  - Prior interaction/similarity increases prob. of a positive evaluation

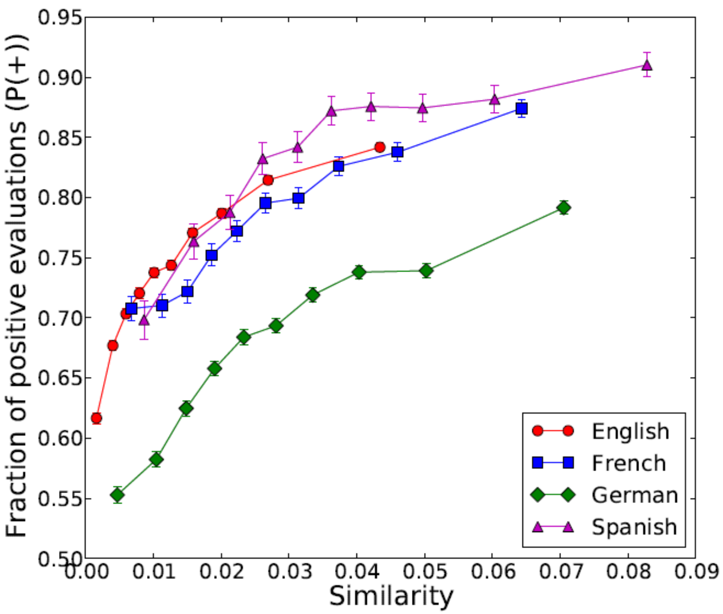

Prior interaction/ similarity boosts positive evaluations

Relating Status and Similarity (1)

- **Observation:**
  - Evaluation depends less on status when evaluator A is more informed
- **Consequence:**
  - Evaluators use status as proxy for quality in the absence of direct knowledge of B

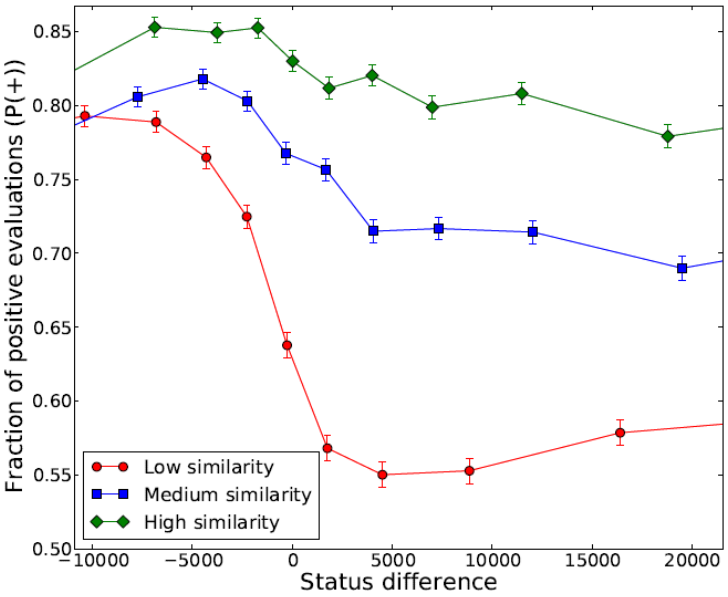

Status is a proxy for quality when evaluator does not know the target

Relating Status and Similarity (2)

- **Observation:**
  - Evaluators with higher status than the target are more similar to the target

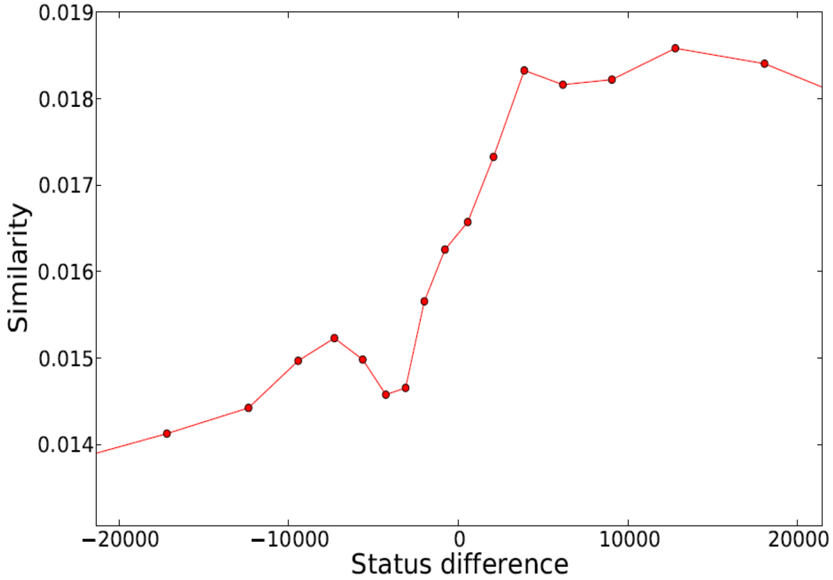

- **Selection bias:**
  - High-status evaluators are more similar to the target

Elite evaluators vote on targets in their area of expertise

Puzzle: Status

- **Evaluator A evaluates target B**
- Prob. of positive evaluation of A as a function of status difference: $\Delta = S_A - S_B$
  - Hypothesis: **Monotonically decreases**

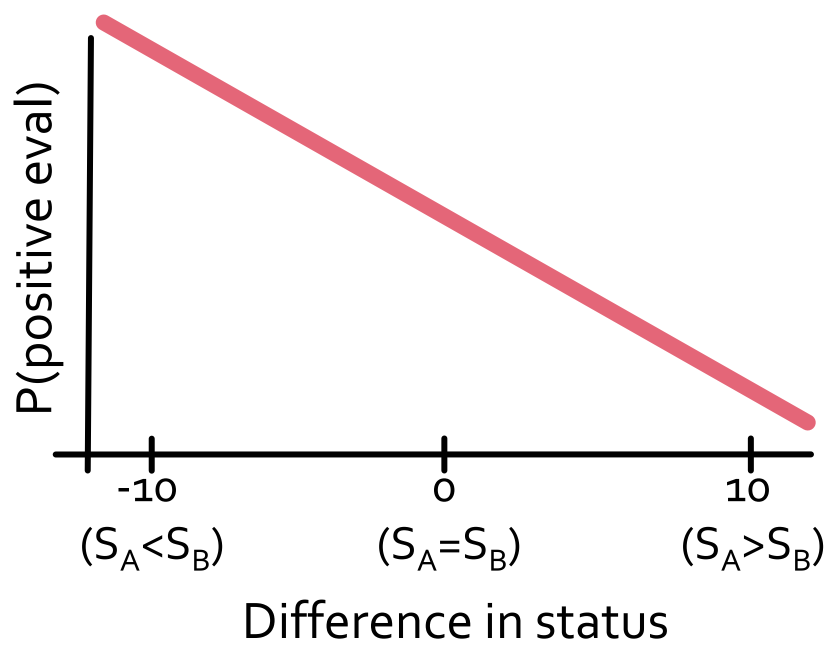

Puzzle: The Mercy Bounce

- Prob. of positive evaluation of B as a function of status difference: $\Delta = S_A - S_B$
- **Observations:**
  - A is especially negative when status equals: $S_A = S_B$
  - "Mercy bounce" for $S_A > S_B$

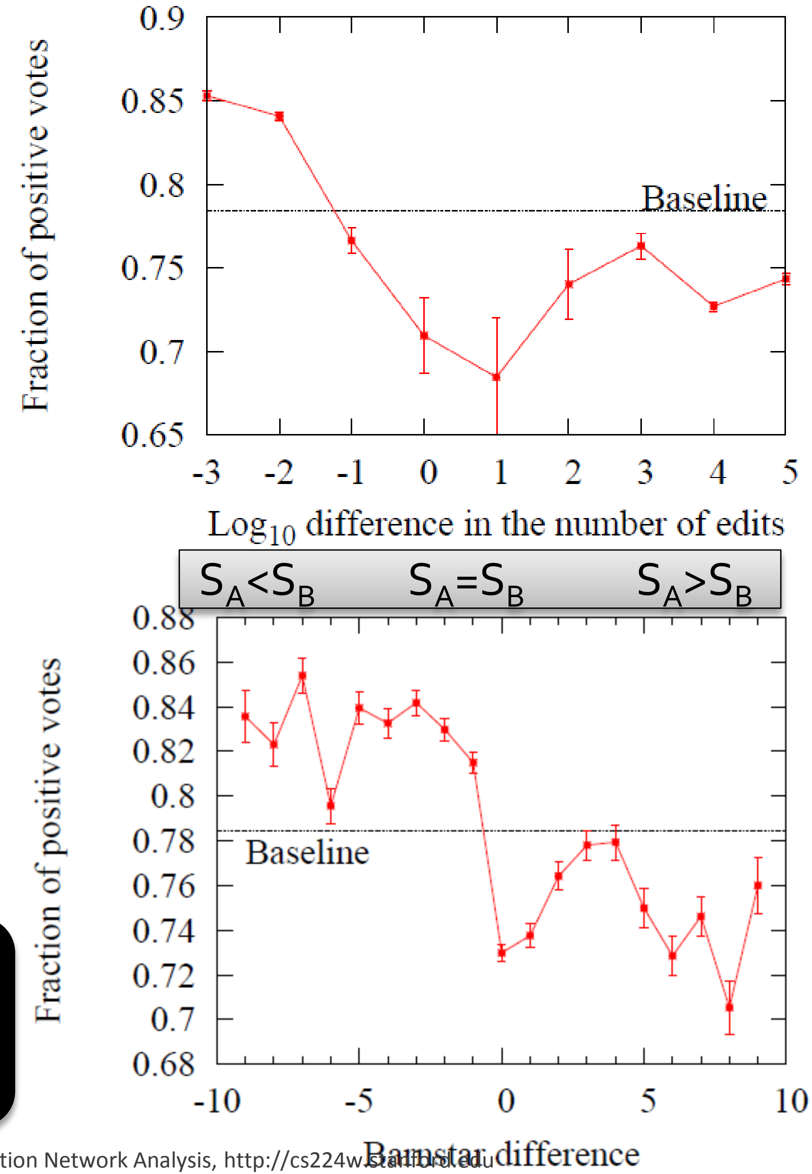

How to explain the **mercy bounce?**

Why Most Harsh at 0 Difference?

How to explain low aggregate evaluations given by users to others of same status?

- Not due to users being tough on each other
  - Similarity increases the positivity of evaluations

Possible explanation:

- Most targets have low status (small $\Delta > 0$)
- Low-status targets are judged on abs. status
  - The rebound persists even for high-status targets

Important Points

- Social media sites are governed by (often implicit) user evaluations
- Wikipedia voting process has an explicit, public and recorded process of evaluation
- Main characteristics:
  - Importance of relative assessment: **Status**
  - Importance of prior interaction: **Similarity**
  - Diversity of individuals' response functions
- **Application:** Ballot-blind prediction

Ballot-blind Prediction

- **Predict Wikipedia election results without seeing the votes**
  - Observe identities of the first k(=5) people voting (but *not* how they voted)
  - Want to predict the election outcome (promotion/no promotion)
- **Why is it hard?**
  - Don’t see the votes (just voters)
  - Only see first 5 voters (10% of the election)

Ballot-blind: the Model

- **Idea:** Split the status-similarity space (s,Δ) in to 4 quadrants
- Model deviation in voter's behavior when they evaluate a candidate from a particular quadrant:
  - d(s,Δ) … avg. deviation in fraction of positive votes
    - When voters evaluate a candidate C from a particular (s,Δ) quadrant, how does this change their behavior

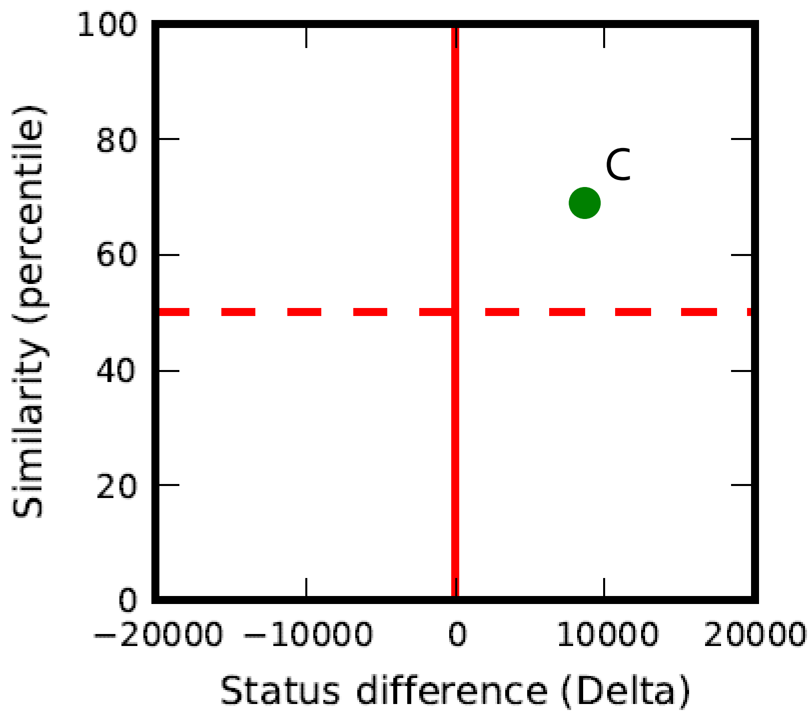

Ballot-blind: the Model

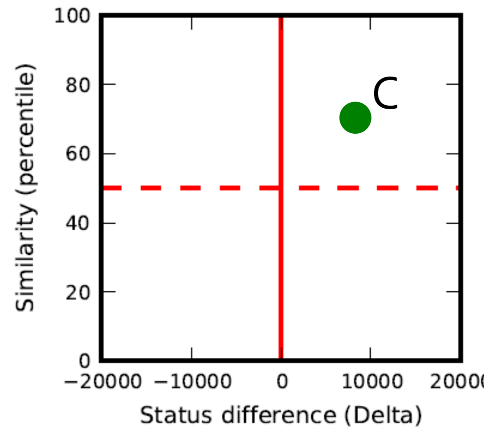

- d(s,Δ) … signed deviation in the fraction of positive votes when *E* evaluates *C* of similarity s and status difference Δ
  - *P(E_i=1)* … prob. evaluator *E* votes + in election *i*
- The models:
  - Global **M1**: $P(E_i = 1) = P_i + d(\Delta_i, s_i)$
  - Personal **M2**:

$$P(E_i = 1) = \alpha \cdot P_i(\Delta_i, s_i) + (1 - \alpha) \cdot d(\Delta_i, s_i)$$

where P_i is empirical frac. of + votes of E

Results: Wikipedia

- Predictive accuracy of baselines:
  - Guessing: 52%
  - If we know votes: 85%
  - Bag-of-features **B1**: 69%
- **Model based on status and similarity:**
    - Does not see votes
    - Sees only first 5 votes (10% of the lection)
  - Global model **M1**: 76%
  - Personal model **M2**: 75%

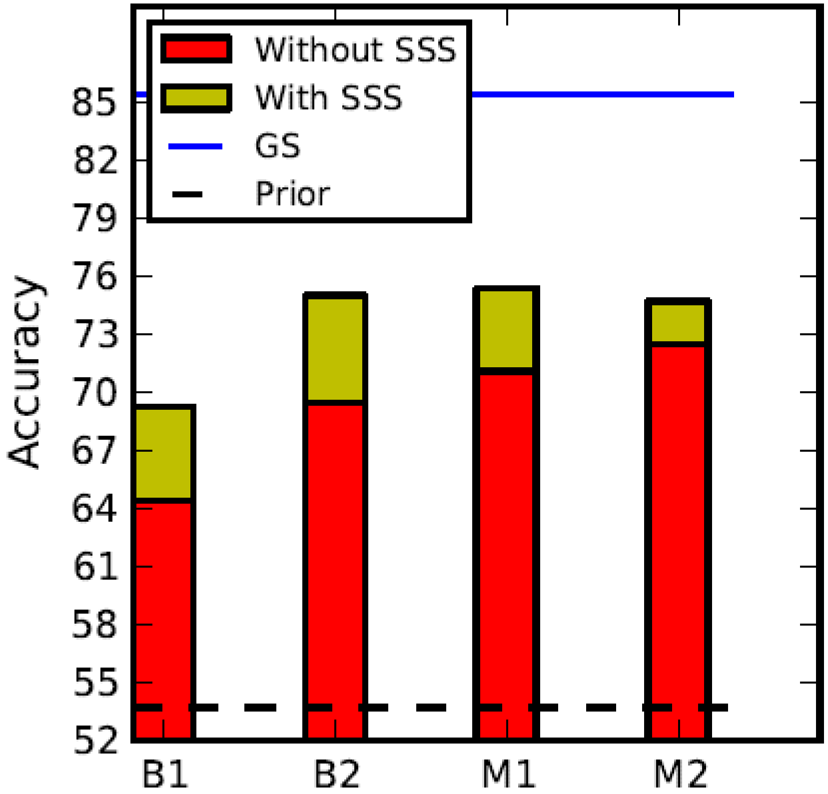

Audience composition predict audience's reaction

Important Points

- Online social systems are globally organized based on **status**
- **Similarity** plays important role
- Audience composition helps predict audience's reaction
- **What kinds of opinions do people find helpful?**

What do people find helpful?

- **What do people think about our recommendations and opinions?**

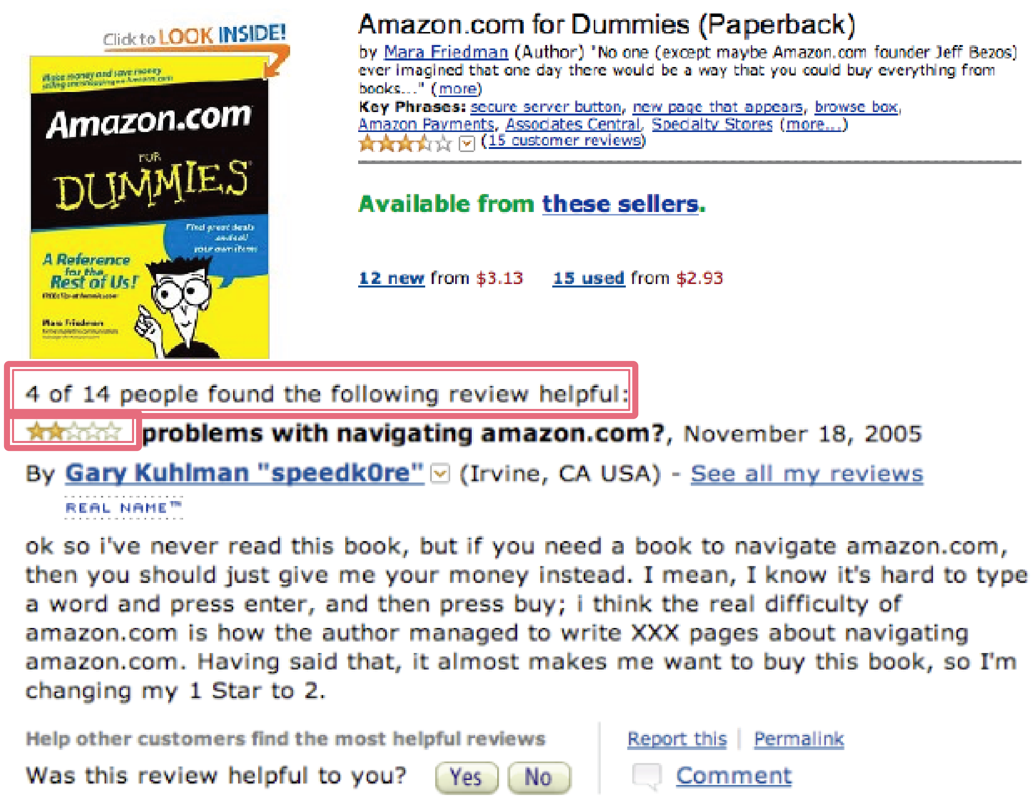

Review helpfulness: Conformity

- People find **conforming** opinions more helpful

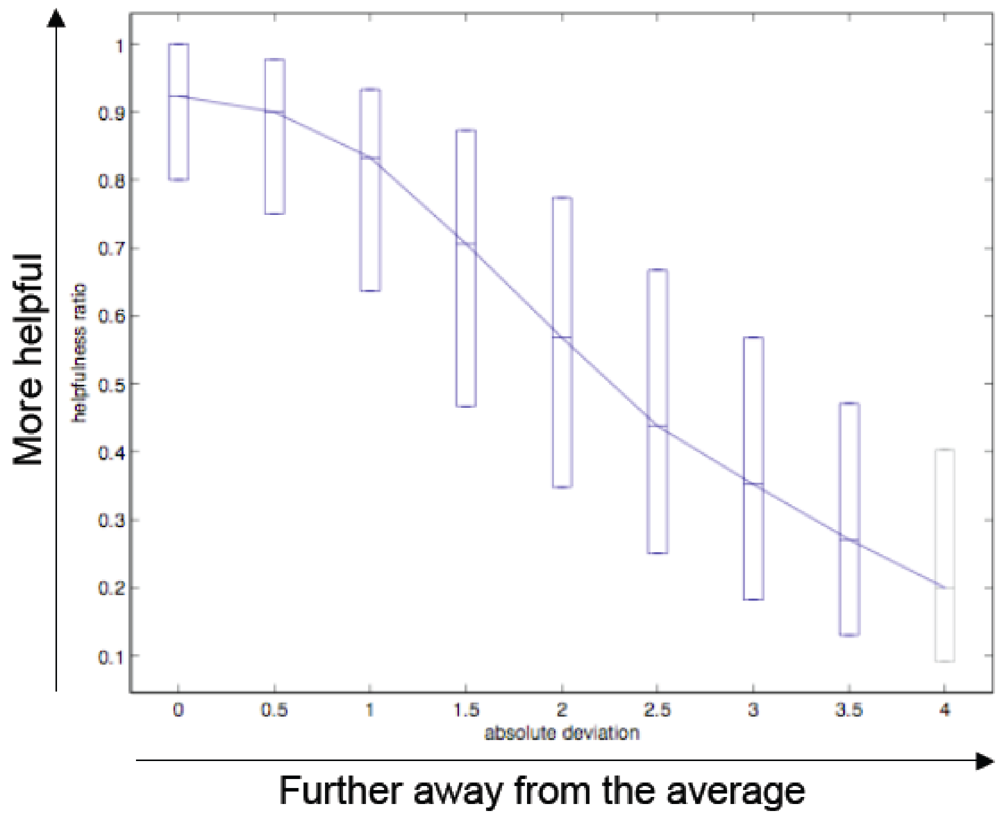

Review helpfulness: Deviation

- **Positive** reviews are more helpful