The Effect of Recommendations on Network Structure

[WTF: What Happened To The Twitter Graph?]

Jessica Su
Stanford University
Computer Science

advised by

managed by

Sharad Goel
Stanford University
Management Science & Engineering

research buddies

Aneesh Sharma
Twitter
Recommendations

Recommendation systems

What effect do recommenders have on the overall marketplace?

"Yelp helps me find hole-in-the-wall restaurants that I'd never have seen otherwise"

"Yelp just boosts traffic for popular restaurants, making life harder for everyone else"

My project

My project

Who to follow · Refresh · View all
Vinay Mahagaokar @vinay...
Follow
Aneesh Sharma @aneeshs
Follow
Praveen B. @praveenbom
Follow

What effect do recommenders have on the structure of the Twitter graph?

"I'm so glad Who To Follow exists, because before I never got followers, and now I do!"

"Who To Follow is okay and all, but it always shows me people like Justin Bieber, who I already knew about"

What effect do recommenders have on the structure of the Twitter graph?

"I've made so many friends on Twitter! Who To Follow was where I met my first wife."

"I've been following all these people in Who To Follow, and none of them follow me back"

Causality

Correlational analyses are okay and all…

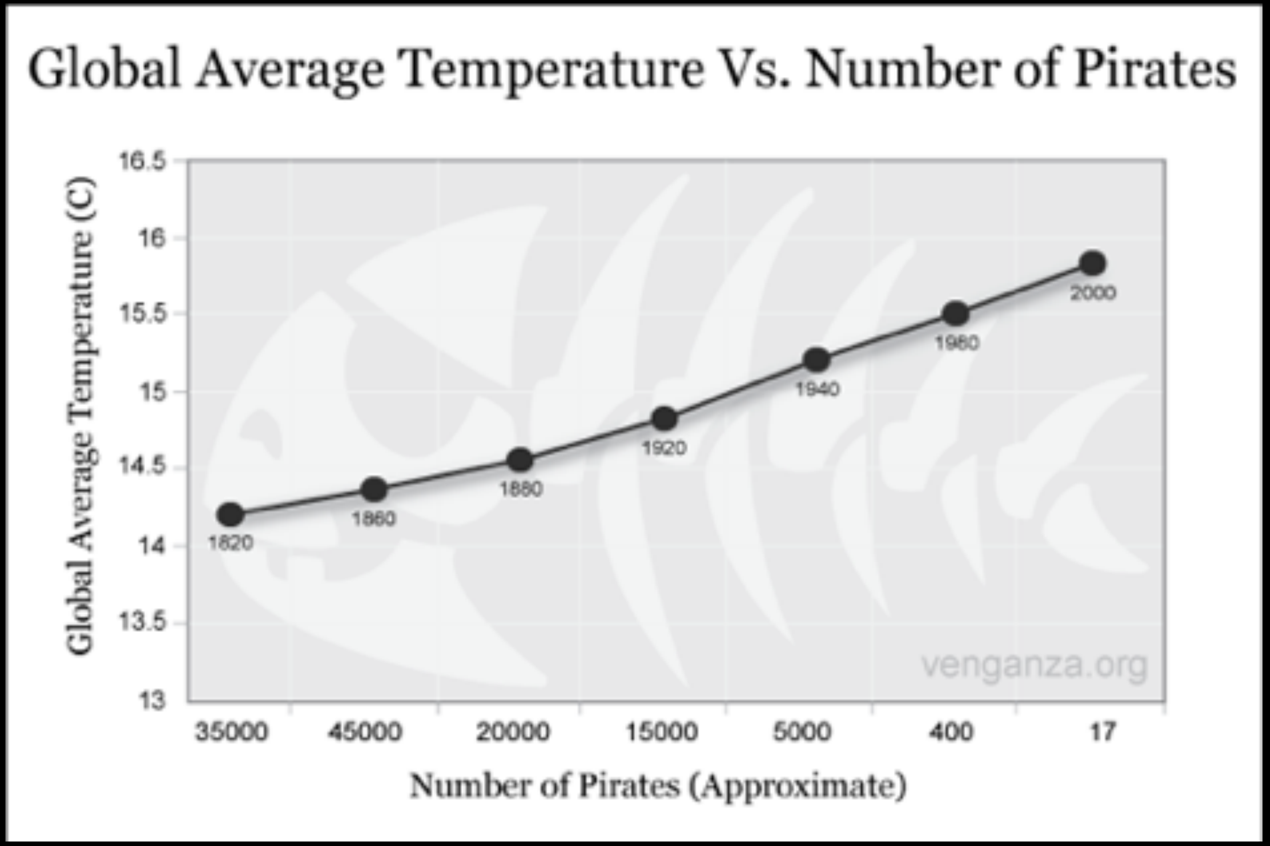

…but to really determine causation, you need to run an experiment

Running experiments on networks is tricky

“Stable unit treatment value assumption”: experimenting on one person doesn’t change what happens to the others

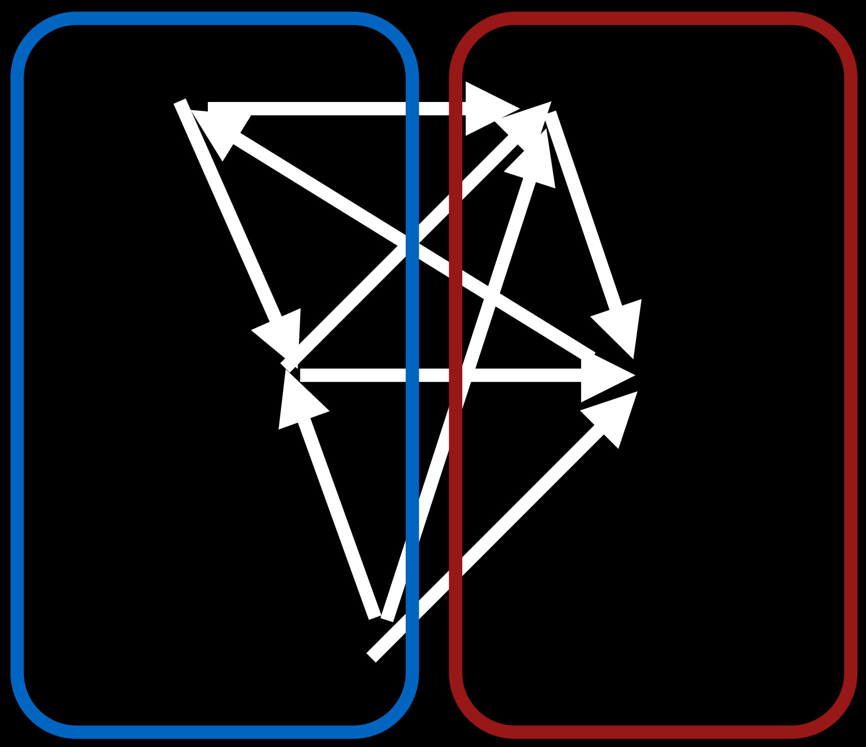

Control? Treatment?

If you encourage a user to follow someone, it affects other users too, making it hard to divide users into treatment and control groups

Effect of Who To Follow recommendations on the global structure of the Twitter graph

Look at the network at two points in time: (1) right before Who To Follow was introduced, and (2) right after.

Due to the short time window, any differences in network structure should be random, **except** for ones caused by Who To Follow.

Results

There is an effect

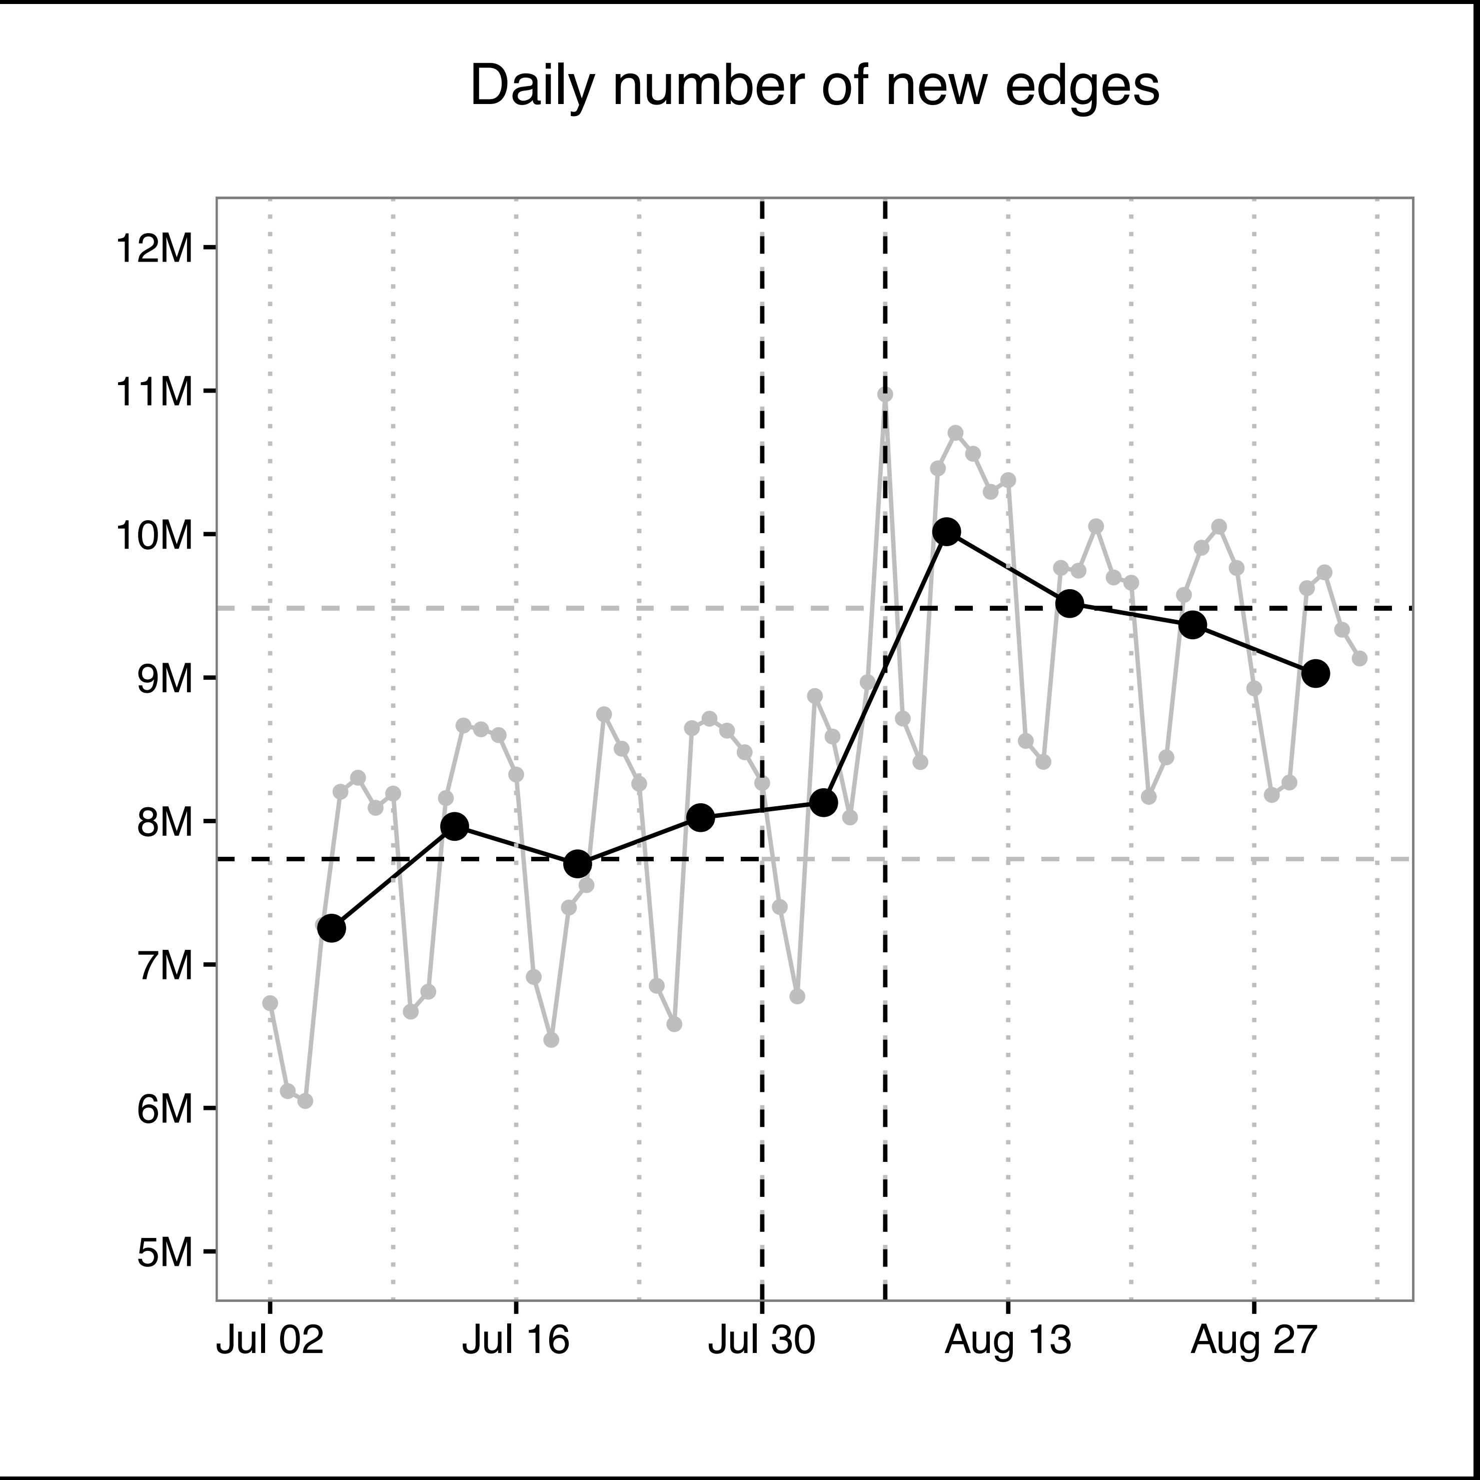

The rich get richer

(but a rising tide lifts all boats)

Everyone sees an increase in follower growth, but the popular people benefit the most

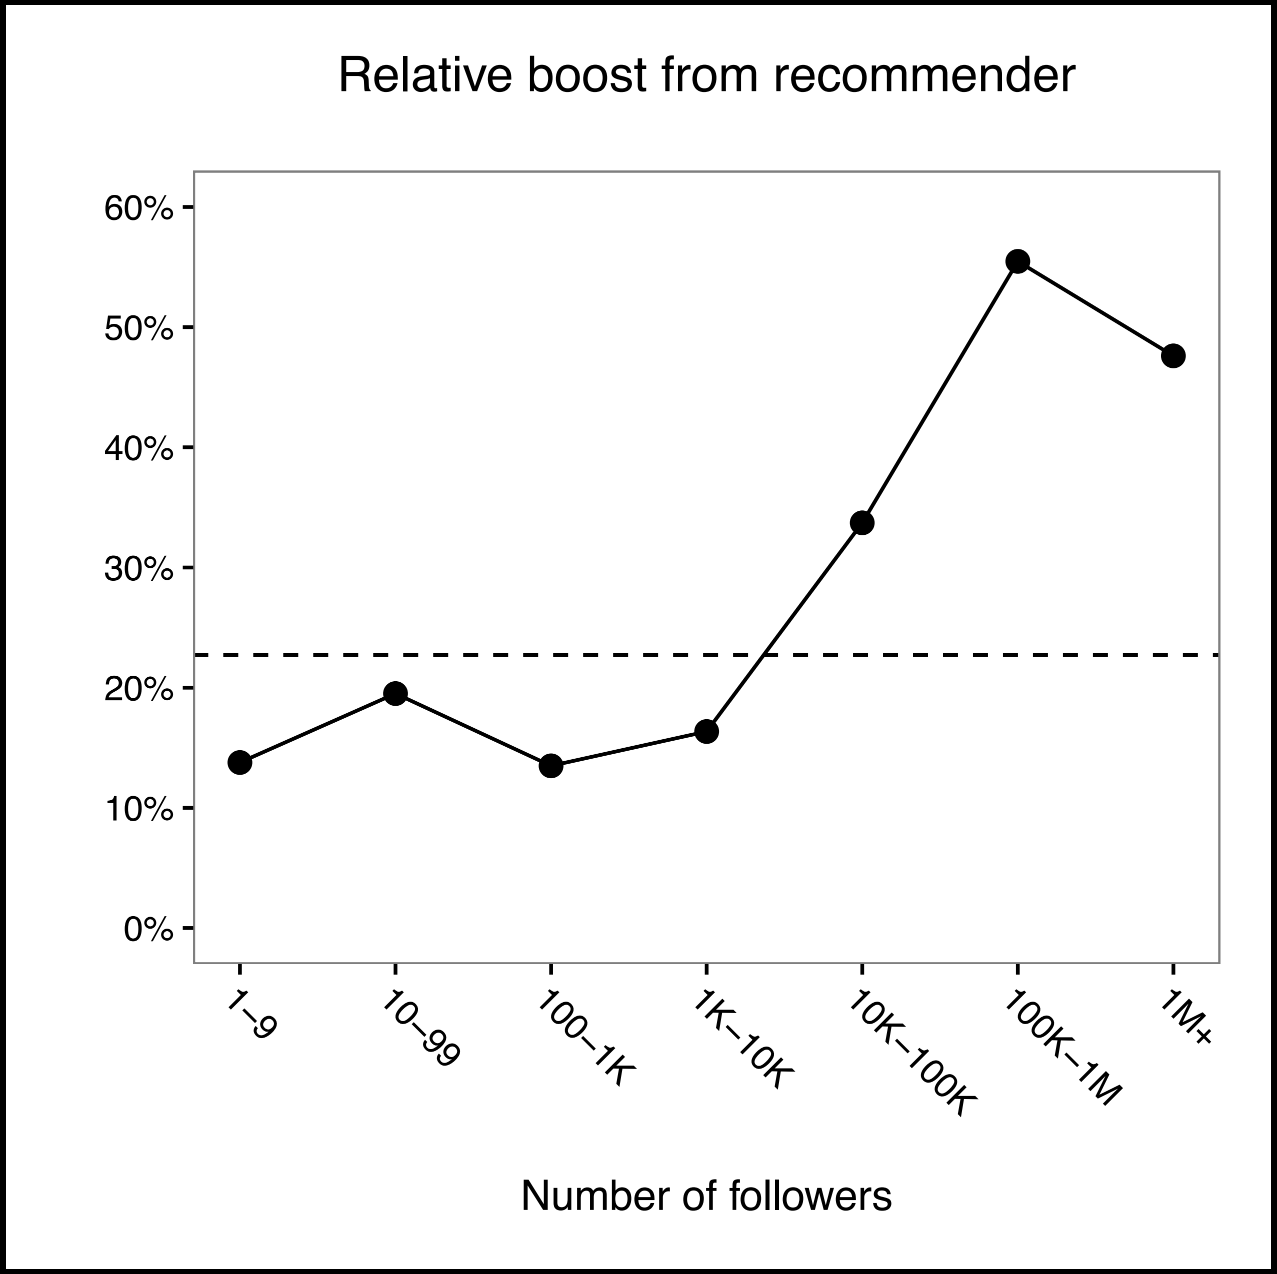

Why do the rich get richer?

Why do the rich get richer?

1. Mismatch between the dynamics of the recommender and the natural evolution of the network

2. Users respond better to popular users

Linear recommender boost

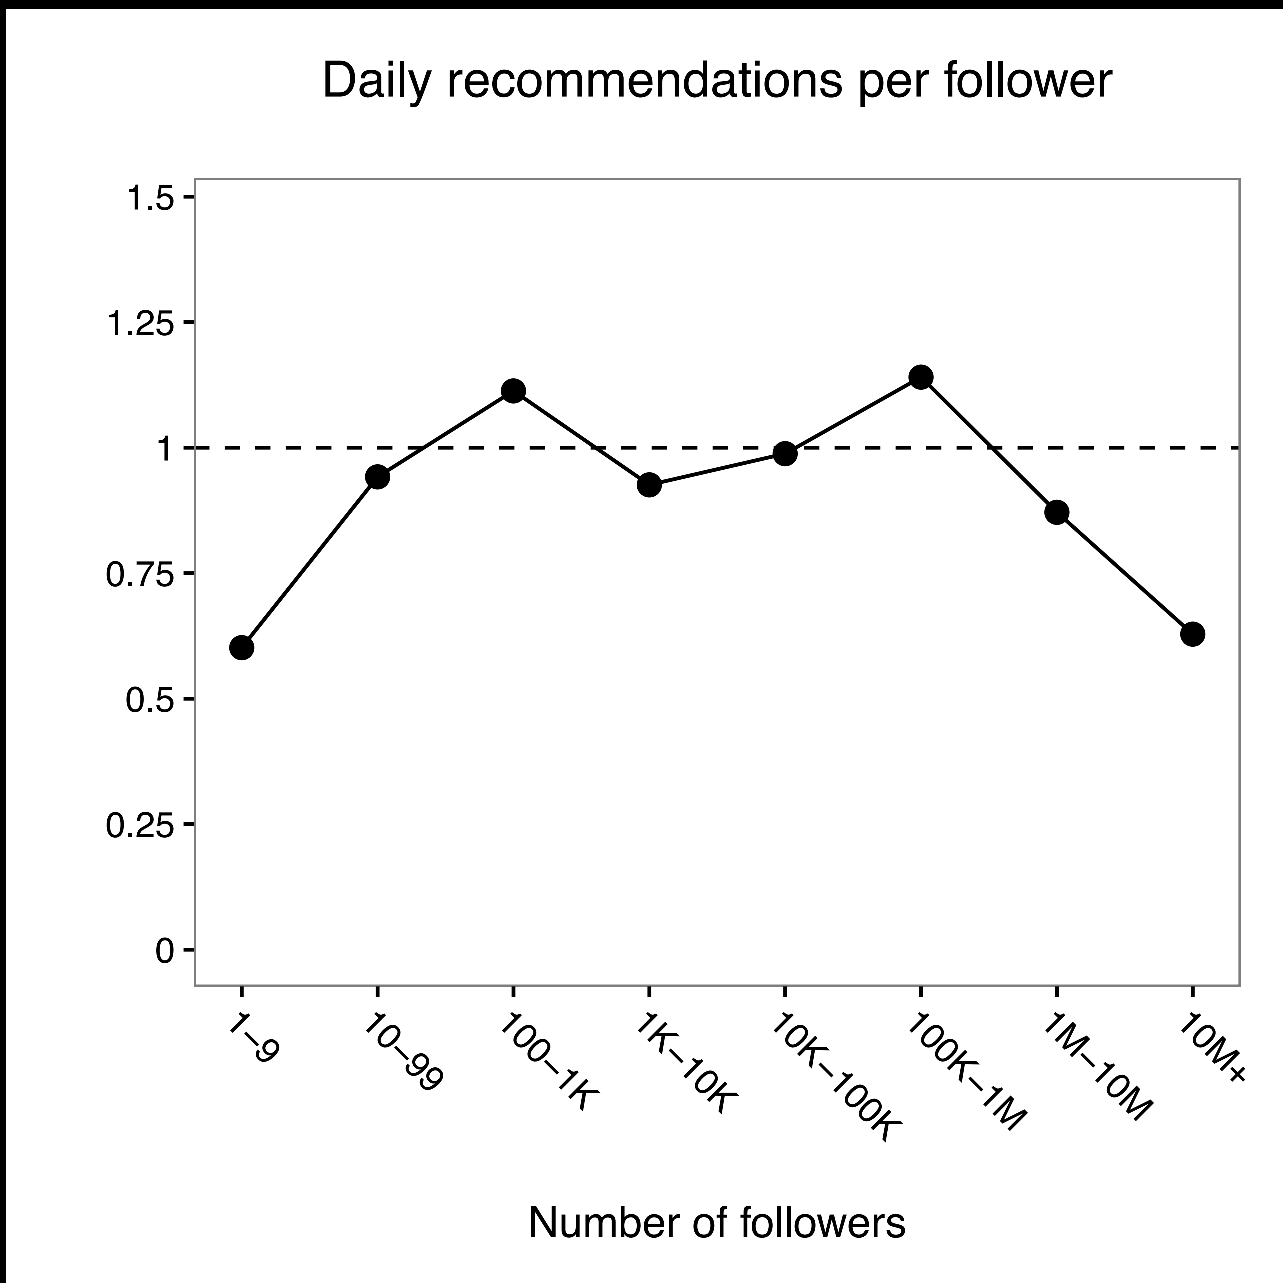

(y-axis obscured to protect confidentiality; 1 denotes the overall average)

The recommender displays users in a way that is directly proportional to their popularity

Under certain models, we can actually prove this

Linear recommender boost

Assume a "friend-of-friend" recommender, where suggestions come from people who your friends are following

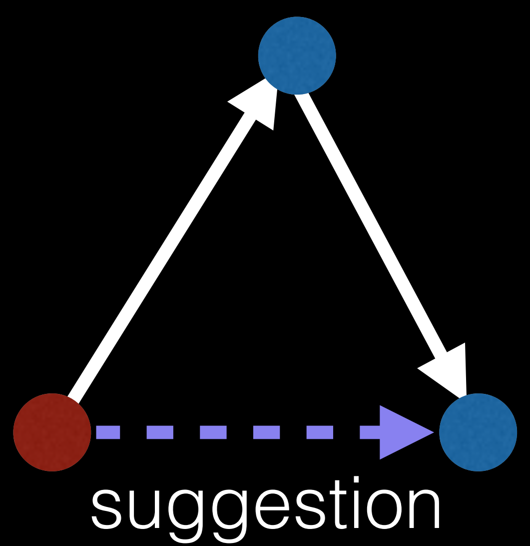

To create a recommendation, choose a friend at random, then choose one of their friends at random

Linear recommender boost

Assume a “friend-of-friend” recommender, where suggestions come from people who your friends are following

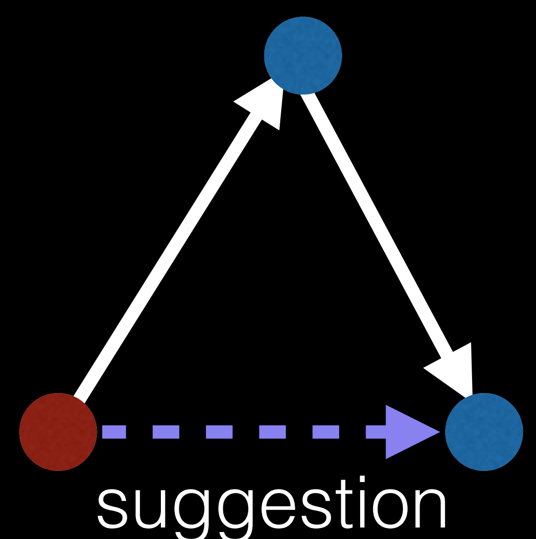

Theorem 1: If you run this friend-of-friend recommender on a random graph generated using the configuration model, then users are recommended in proportion to how many followers they have

Why the boost?

The recommender displays users in a way that is directly proportional to their popularity

Standard models of network growth predict that people gain followers at a rate proportional to their current follower count, so in theory the recommender shouldn't really change things

Sublinear growth

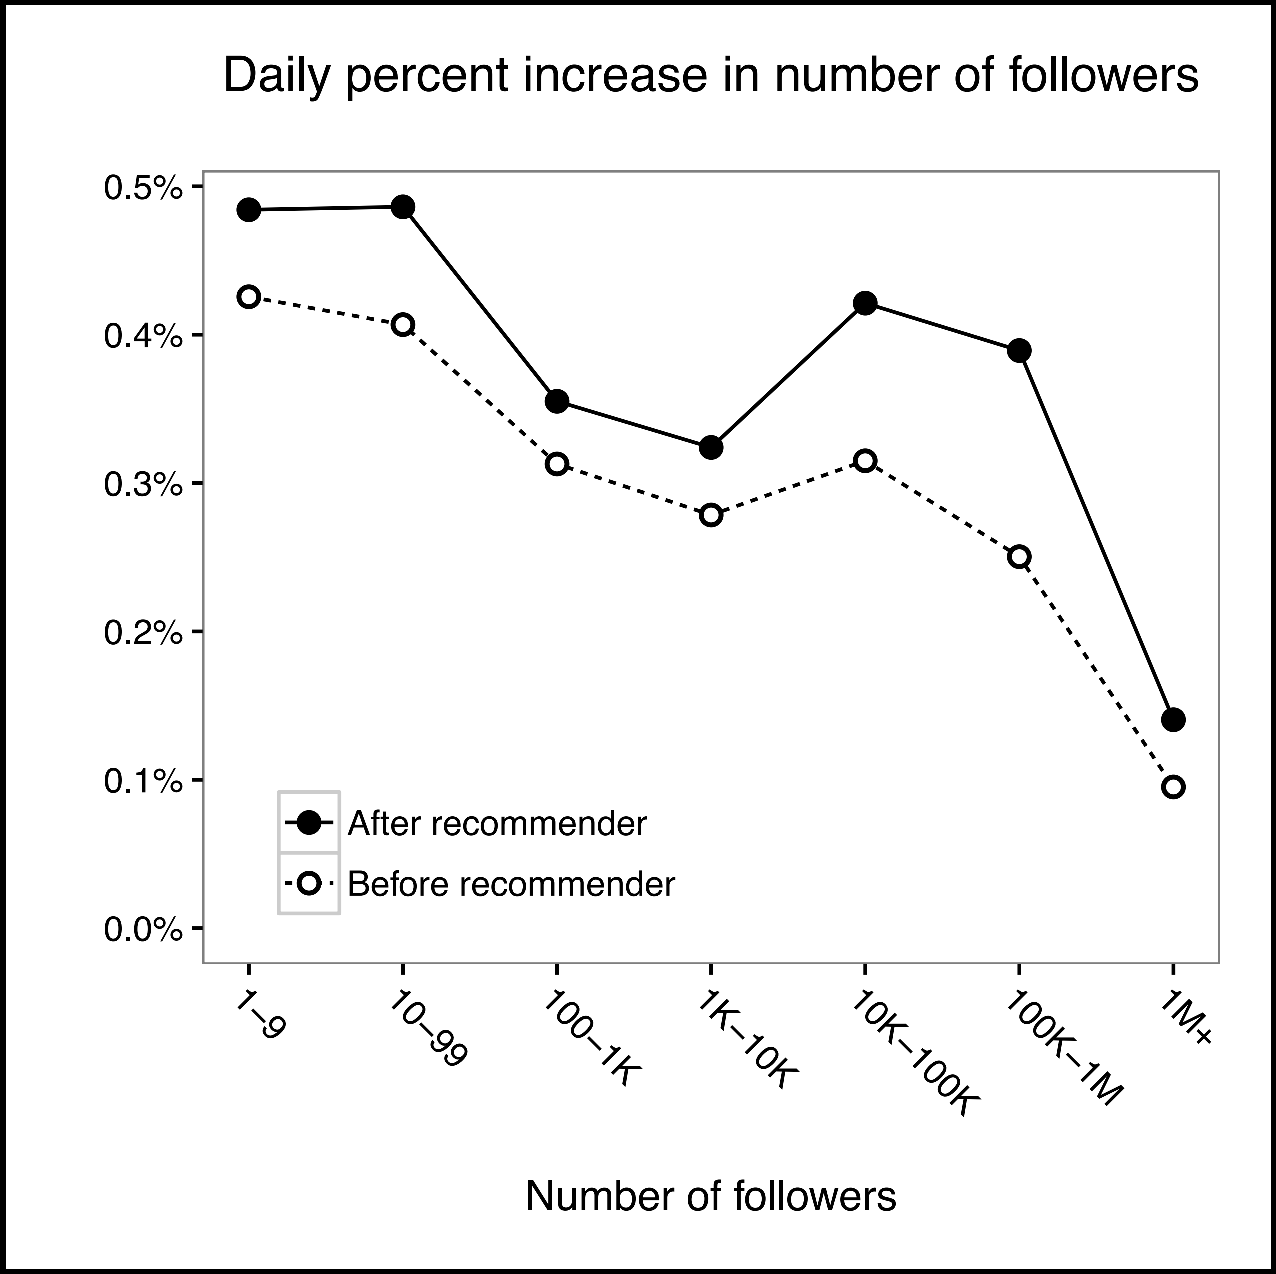

But people **don't** gain followers at a rate proportional to their current follower count

Follower gains grow as a sublinear function of user popularity

(If it were linear, the dotted line would be flat)

Sublinear growth

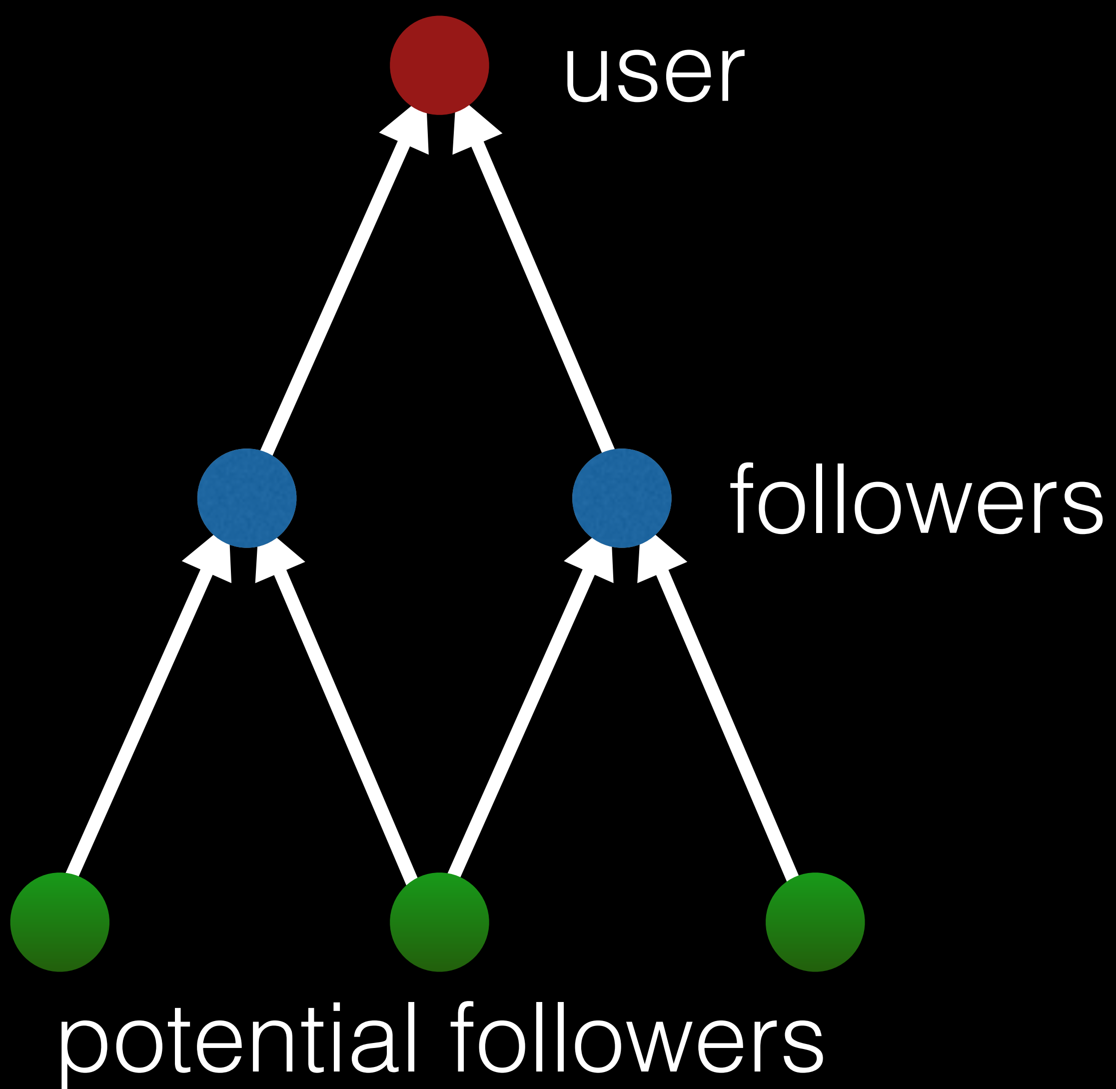

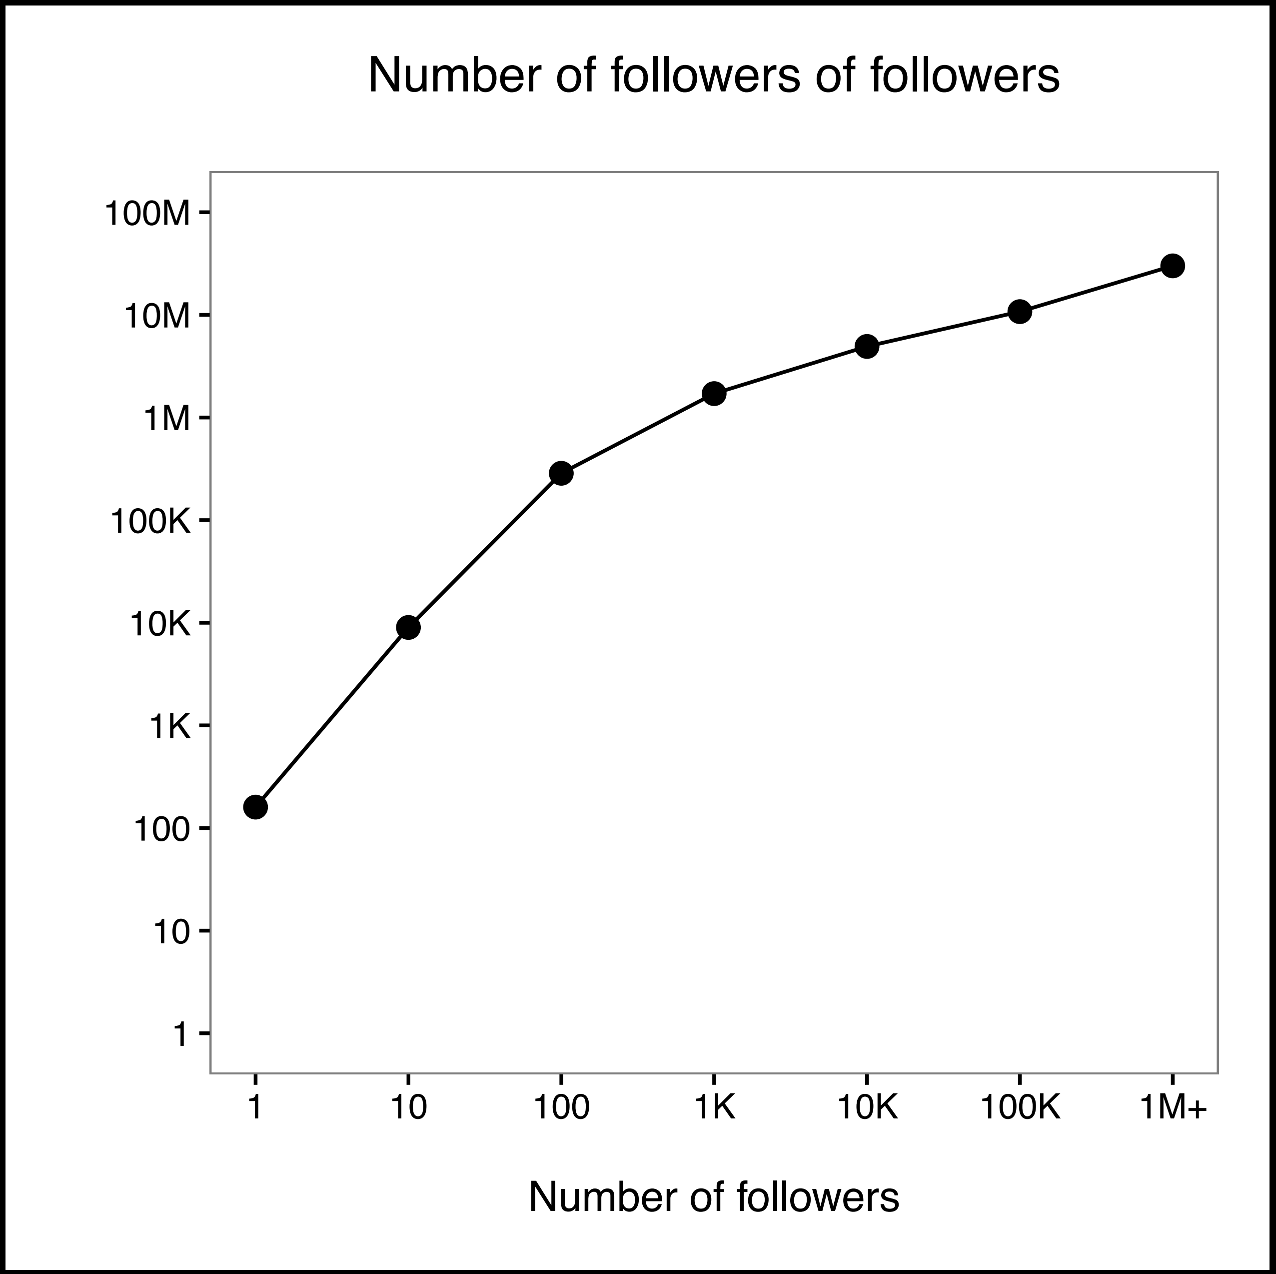

Popular people are penalized by network “overlap”

Sublinear growth / linear boost

How often the recommender displays users is linearly related to how many followers they have

Since the original rate of follower growth was sublinear as a function of follower count, a linear boost disproportionately favors people with more followers

Sublinear growth / linear boost

df/dt = sqrt(f)

f = 100 df/dt = 10

f = 10000 df/dt = 100

df/dt = sqrt(f) + f / 10

f = 100 df/dt = 20

f = 10000 df/dt = 1100

Since the original rate of follower growth was sublinear as a function of follower count, a linear boost disproportionately favors people with more followers

Sublinear growth / linear boost

df/dt = sqrt(f)

f = 100 $\boxed{\text{df/dt = 10}}$

f = 10000 df/dt = 100

df/dt = sqrt(f) + f / 10

f = 100 $\boxed{\text{df/dt = 20}}$

f = 10000 df/dt = 1100

Since the original rate of follower growth was sublinear as a function of follower count, a linear boost disproportionately favors people with more followers

Sublinear growth / linear boost

df/dt = sqrt(f)

f = 100 df/dt = 10

f = 10000 $\boxed{\text{df/dt = 100}}$

df/dt = sqrt(f) + f / 10

f = 100 df/dt = 20

f = 10000 $\boxed{\text{df/dt = 1100}}$

Since the original rate of follower growth was sublinear as a function of follower count, a linear boost disproportionately favors people with more followers

User response

People are more likely to click "Follow" for more popular users

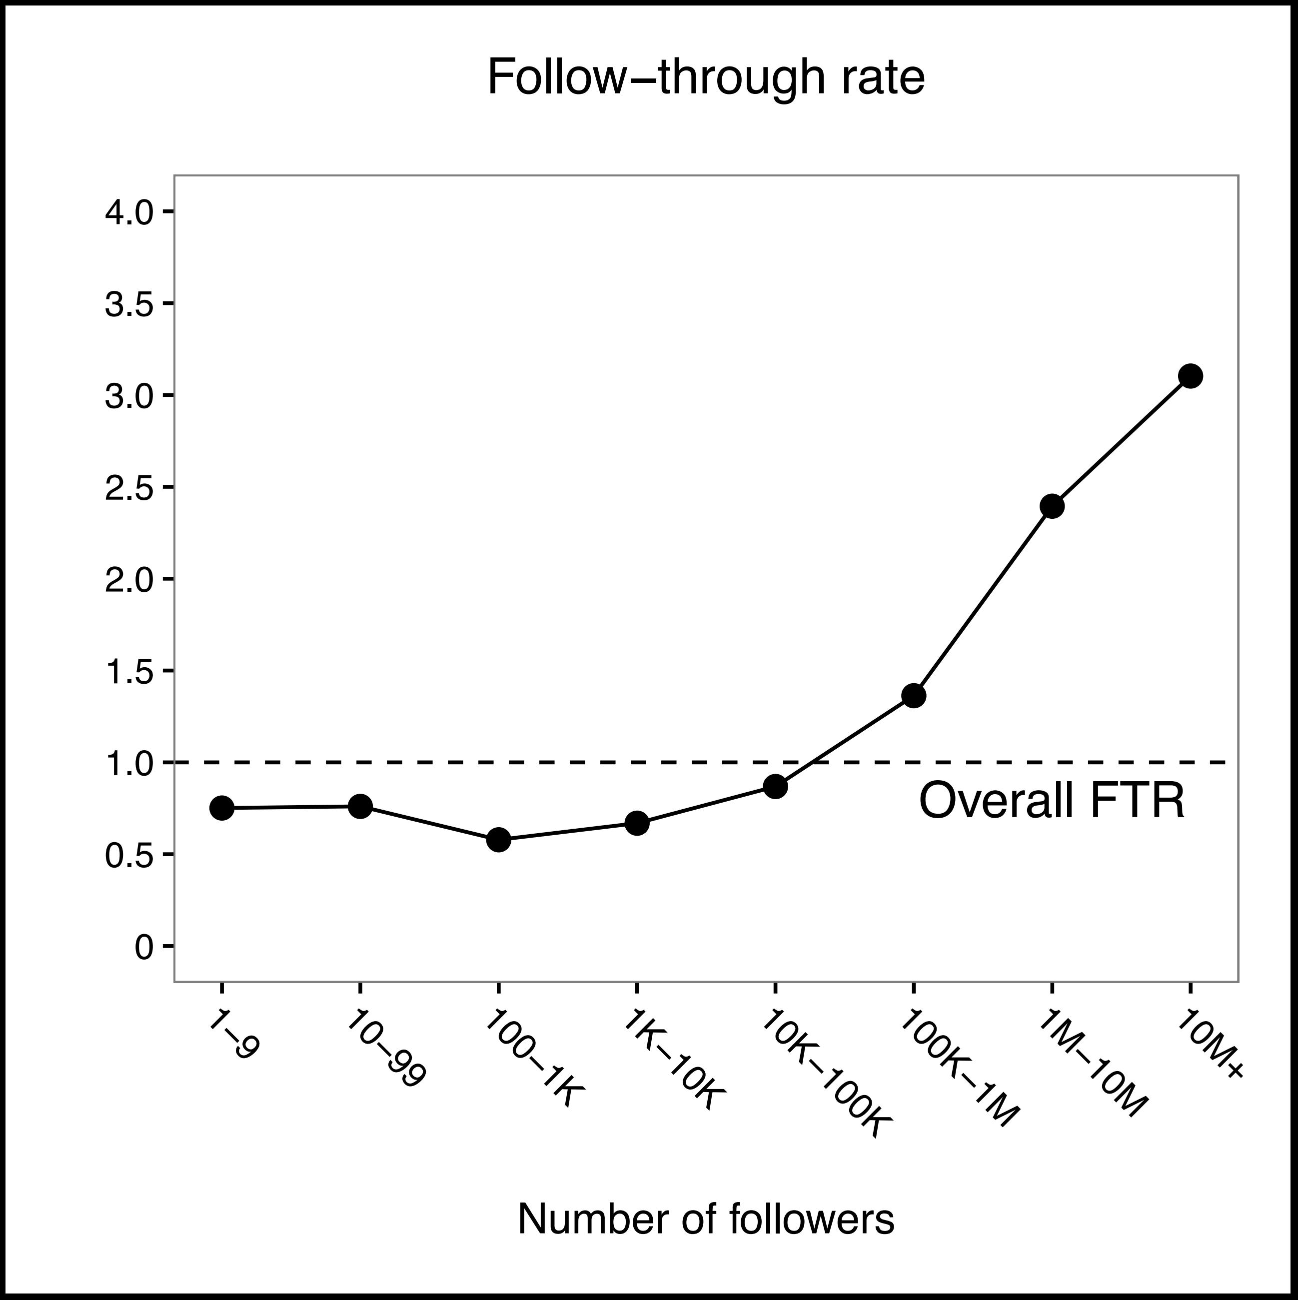

(y-axis obscured to protect confidentiality; 1 denotes the overall average)

Other structural changes

What kind of connections are being formed?

Do people follow outsiders, or people who are already "in their network?"

Do people make friends, or do they follow people who don't reciprocate their interest?

What kind of connections are being formed?

Do people follow outsiders, or people who are already "in their network?"

Do new edges close triangles?

Do people make friends, or do they follow people who don't reciprocate their interest?

What kind of connections are being formed?

Do people follow outsiders, or people who are already "in their network?"

Are new edges reciprocated?

Do new edges close triangles?

Do people make friends, or do they follow people who don't reciprocate their interest?

Graph becomes more “tightly connected”

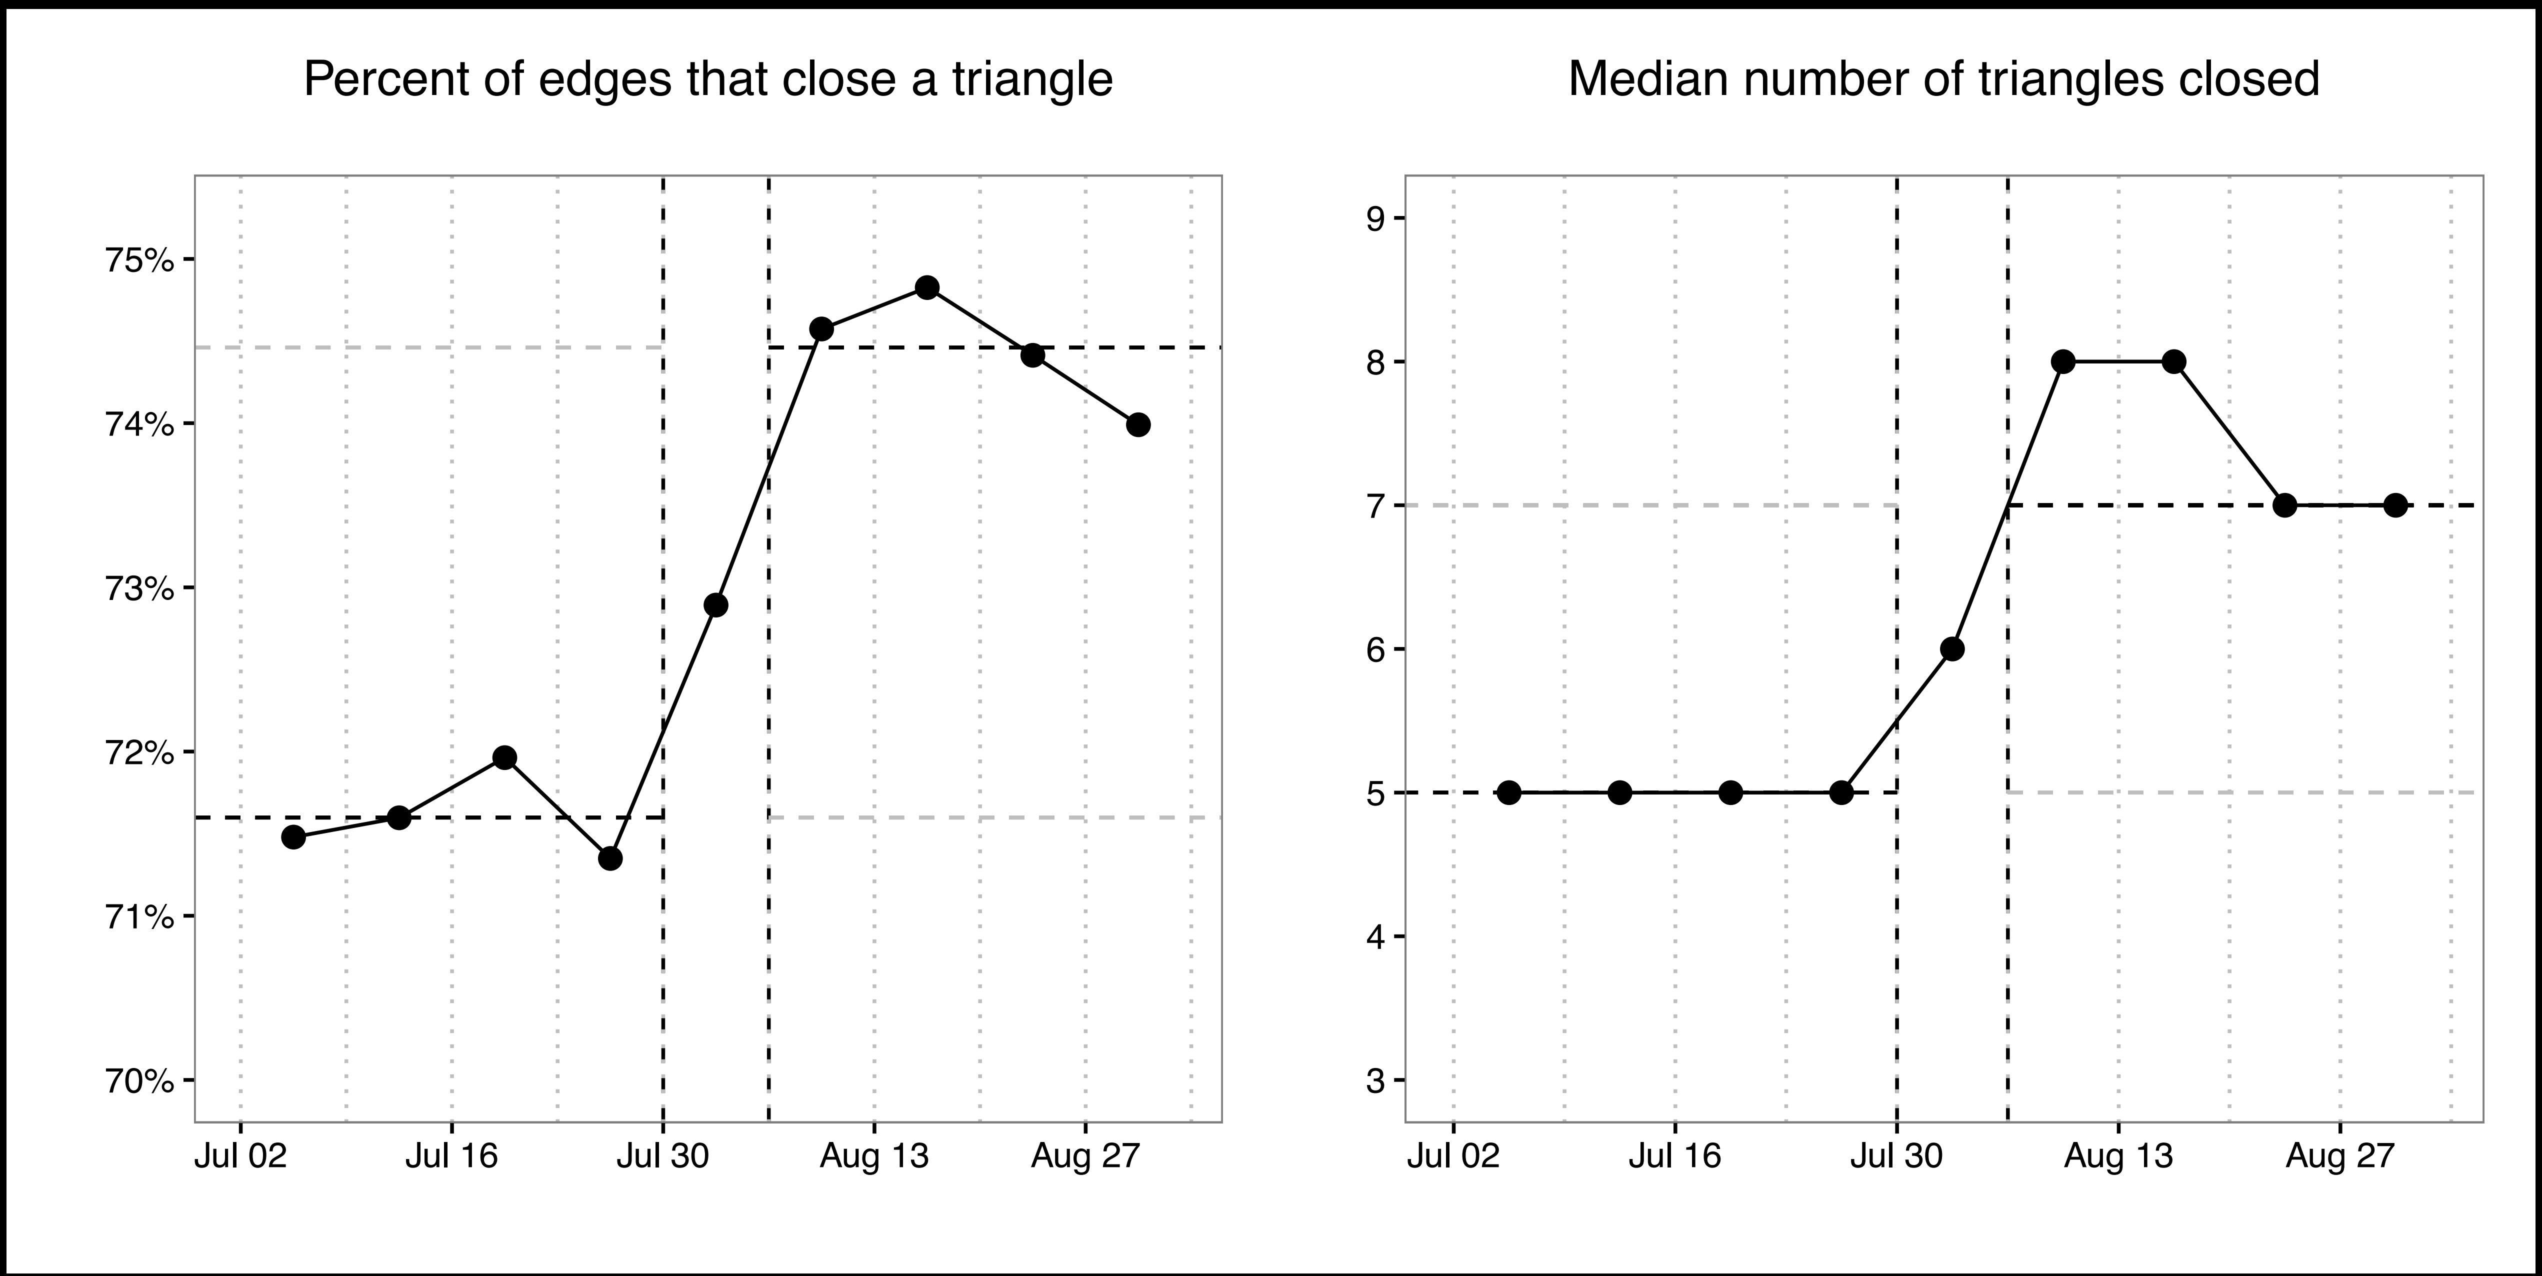

But people aren't really making friends

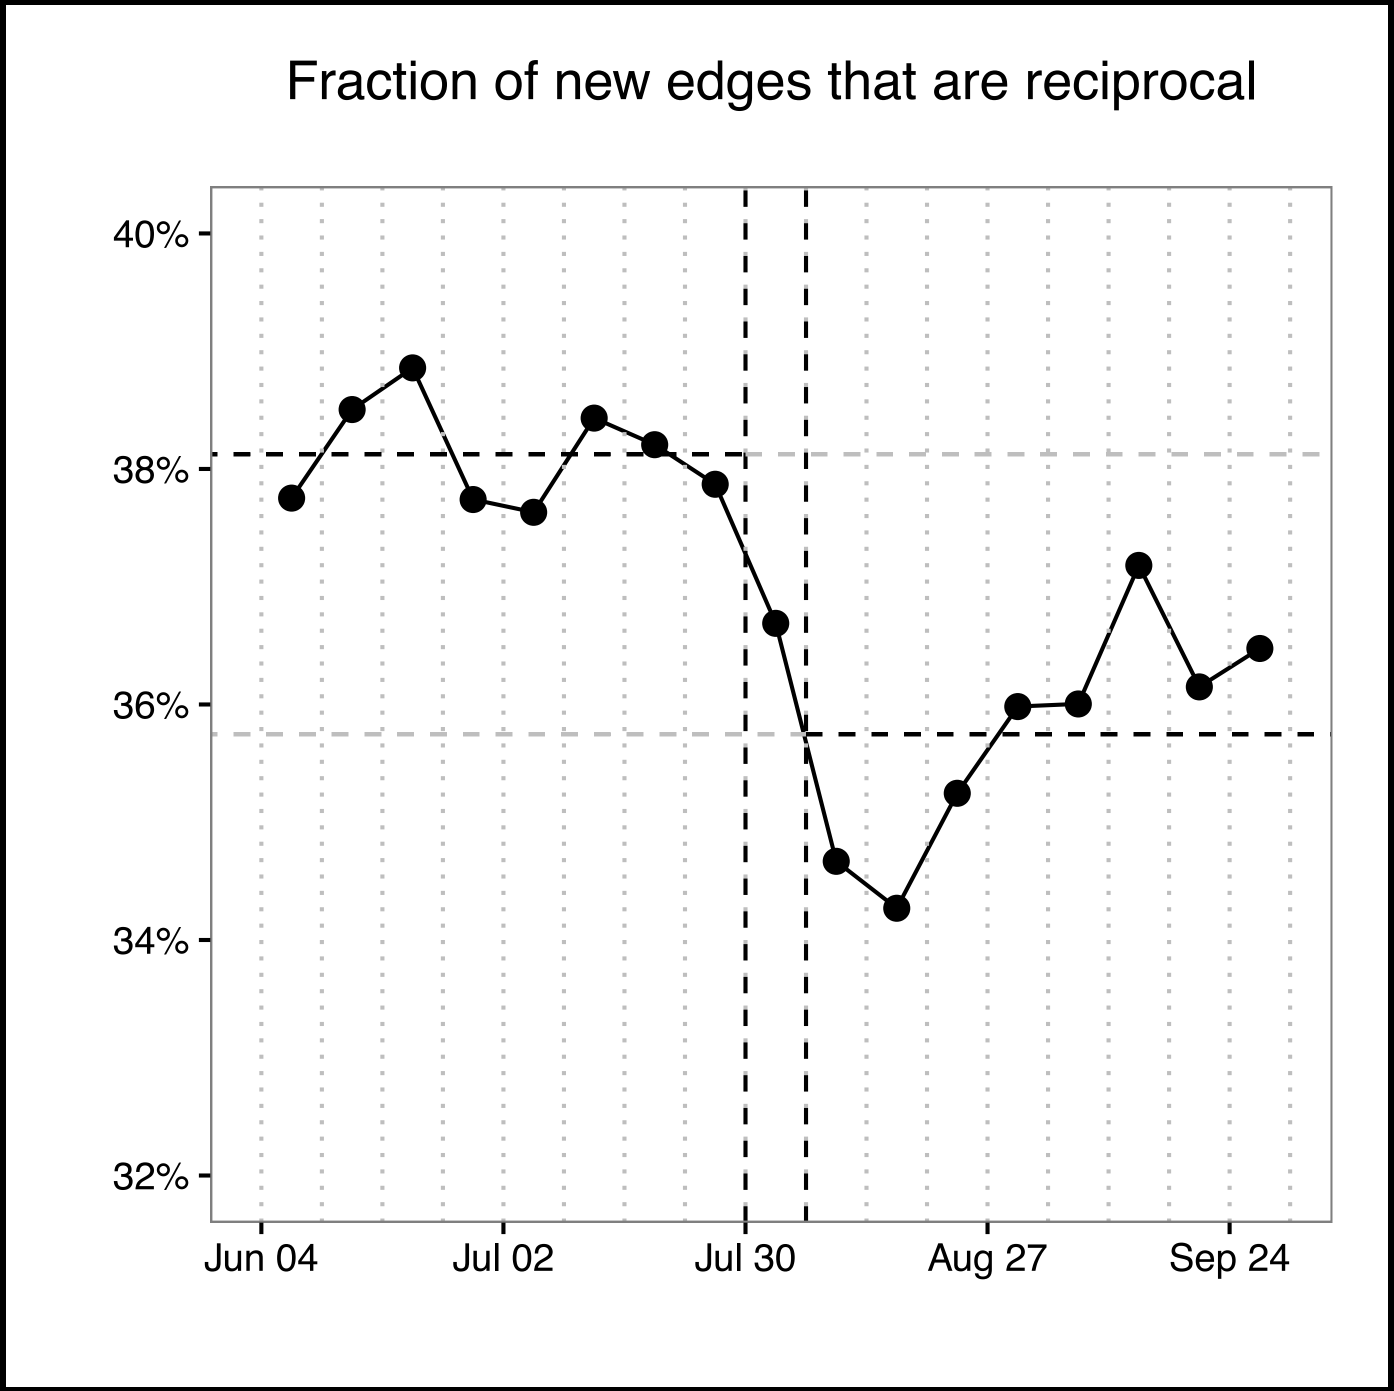

Thanks for listening

FAQ #1: What happens outside the window of observation?

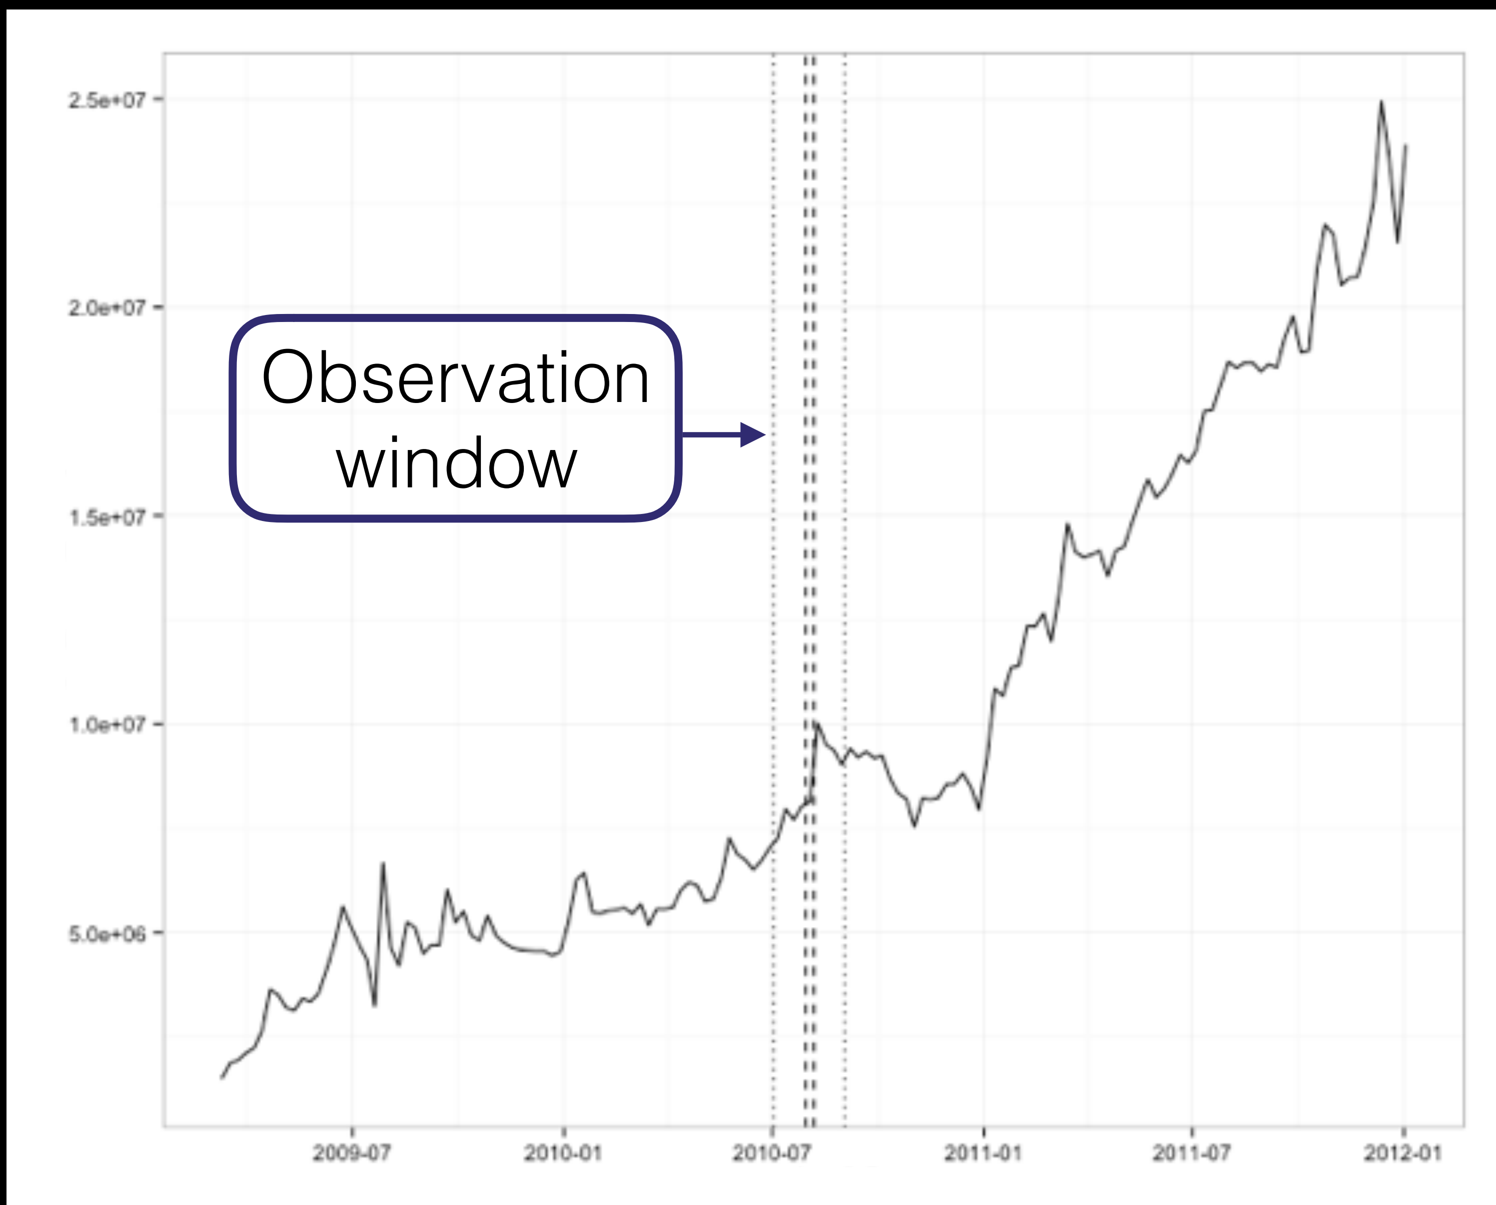

Number of edges added each day

Misleading

Natural experiment breaks down

Changes may no longer be caused by the recommendation system