Diffusion and Viral Marketing in Networks

3-31-2010

Theory - review

Diffusion through social networks: *why* things spread

- Fun: i.e., why do things get popular?
  - Fashion, fads, internet memes, research ideas, …
  - First-order approximation: *preferential attachment* in graphs
- Rational decisions:
  - Decisions made publically with limited information
    - Specifically, decisions where *choice* is **public** but some *evidence* used in the choice is **private**
  - **Decisions made about products** (or behaviors, etc) **that have "network effects"** (aka "externalities")
    - Specifically, the benefits and costs of the behavior are **not completely local** to the decision-maker

Start with some simple cases in a non-networked world

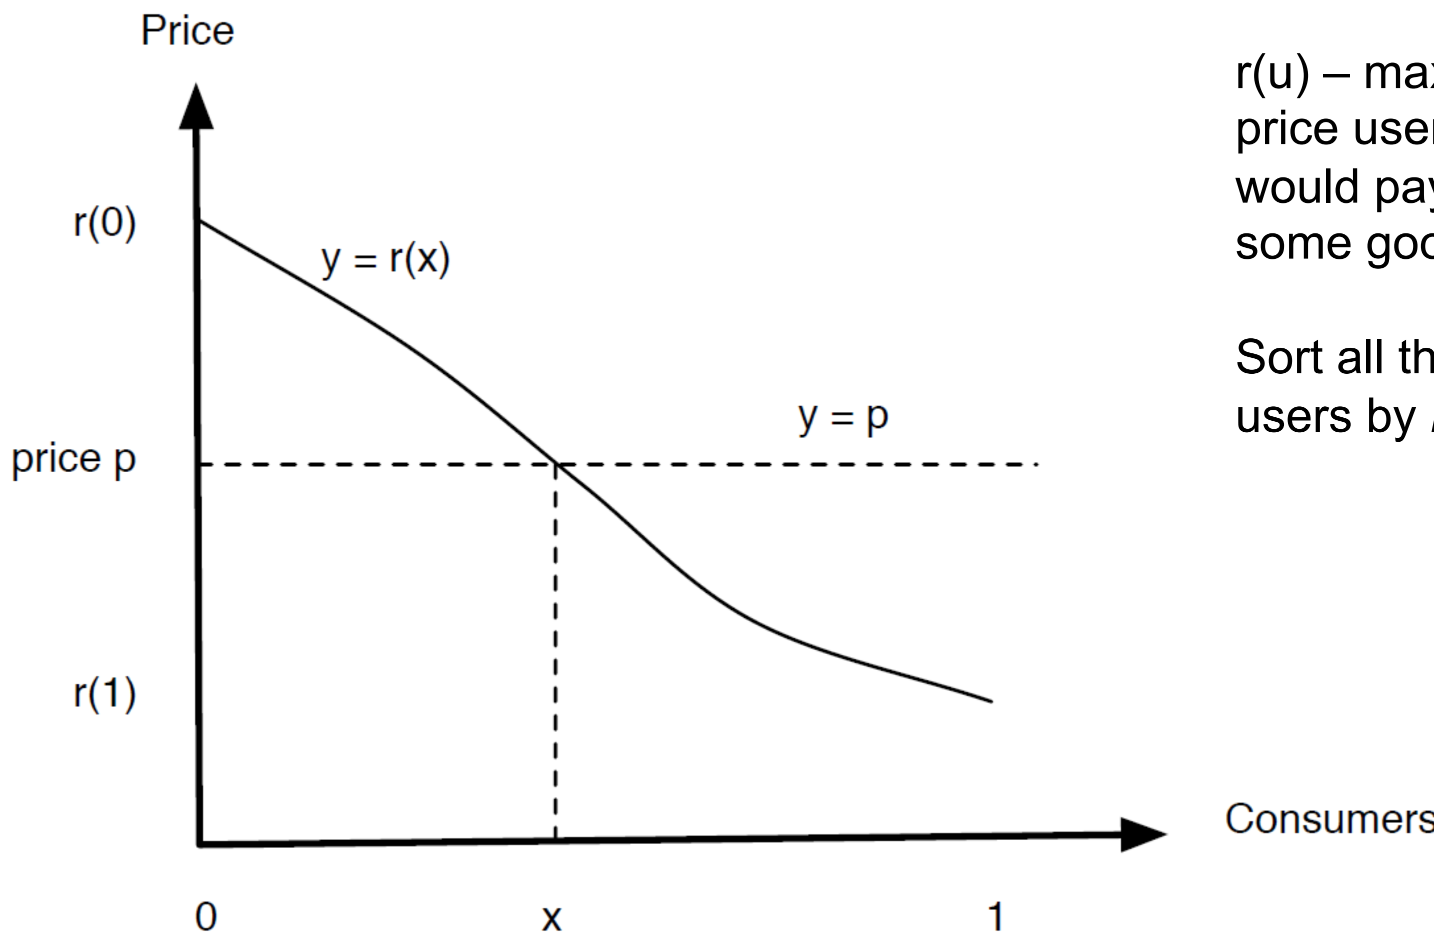

r(u) – max price user *u* would pay for some good.

Sort all the users by *r(u).*

Figure 17.1: When there are no network efforts, the demand for a product at a fixed market price p can be found by locating the point where the curve $y = r(x)$ intersects the horizontal line $y = p$.

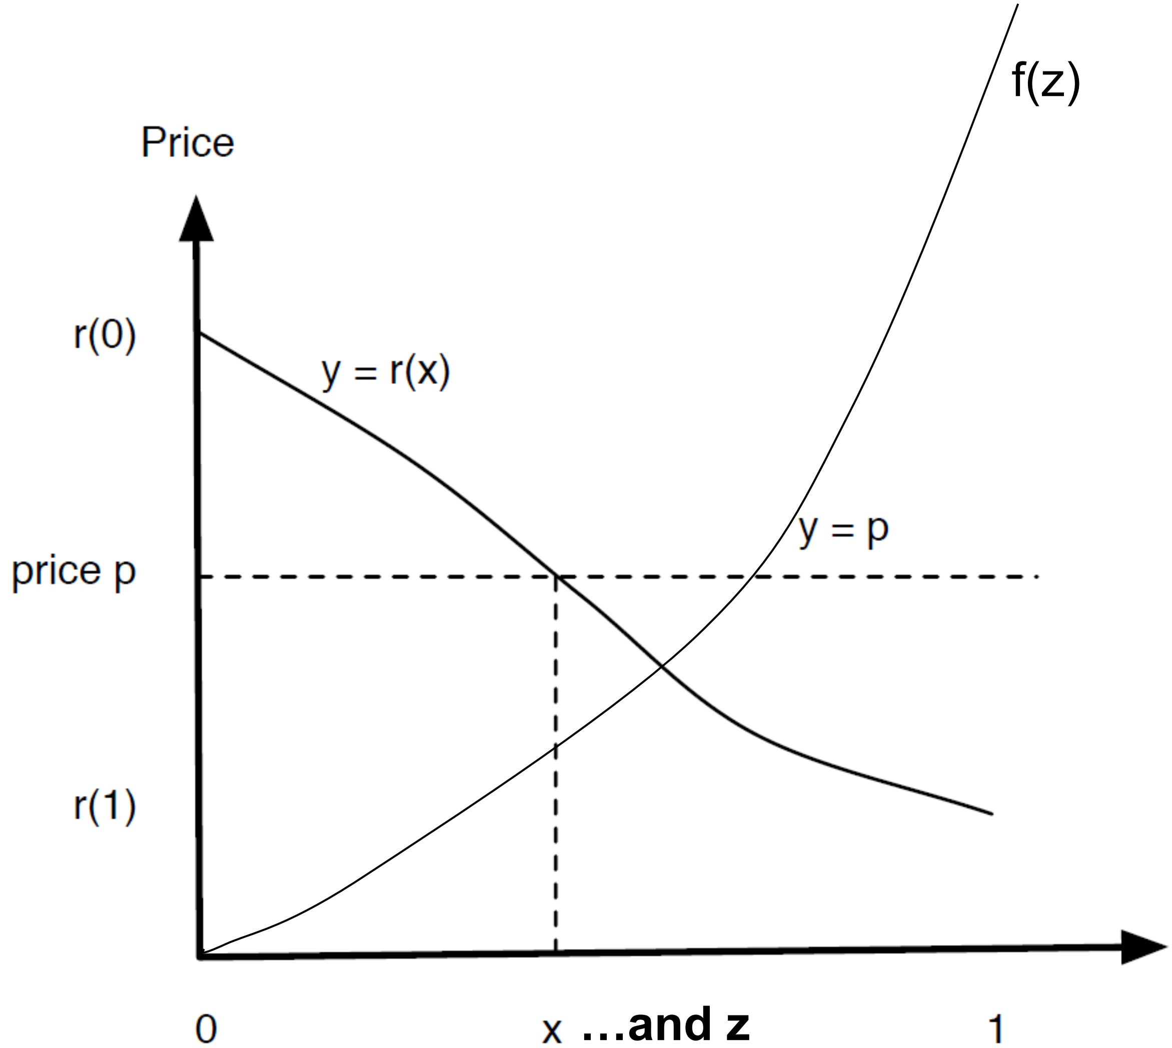

r(u) – *intrinsic* value
f(z) – *network* “value inflation factor” if fraction *z* of users are purchasers

Claim: r(z)*f(z) is max price user z would pay for some good, if fraction *z* of all users buy the good.

Figure 17.1: When there are no network efforts, the demand for a product at a fixed market price p can be found by locating the point where the curve $y = r(x)$ intersects the horizontal line $y = p$.

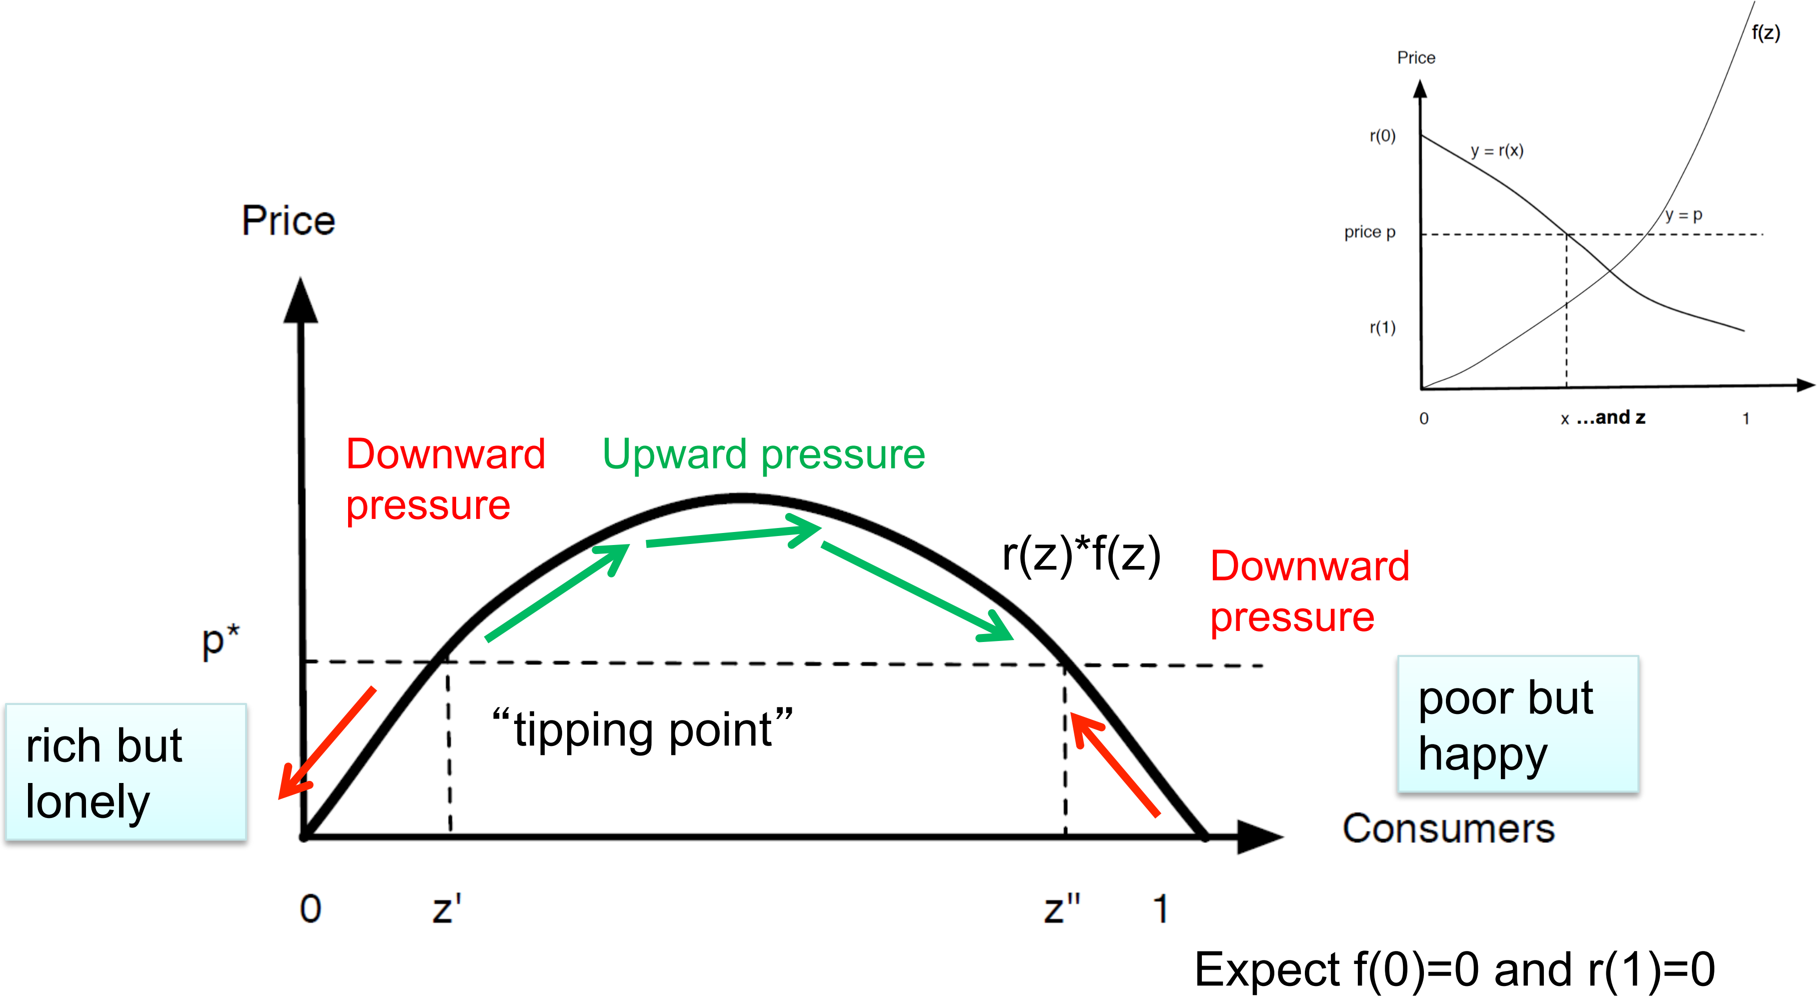

Figure 17.3: Suppose there are network effects and $f(0) = 0$, so that the good has no value to people when no one is using it. In this case, there can be multiple self-fulfilling expectations equilibria: at $z = 0$, and also at the points where the curve $r(z)f(z)$ crosses the horizontal line at height p^*.

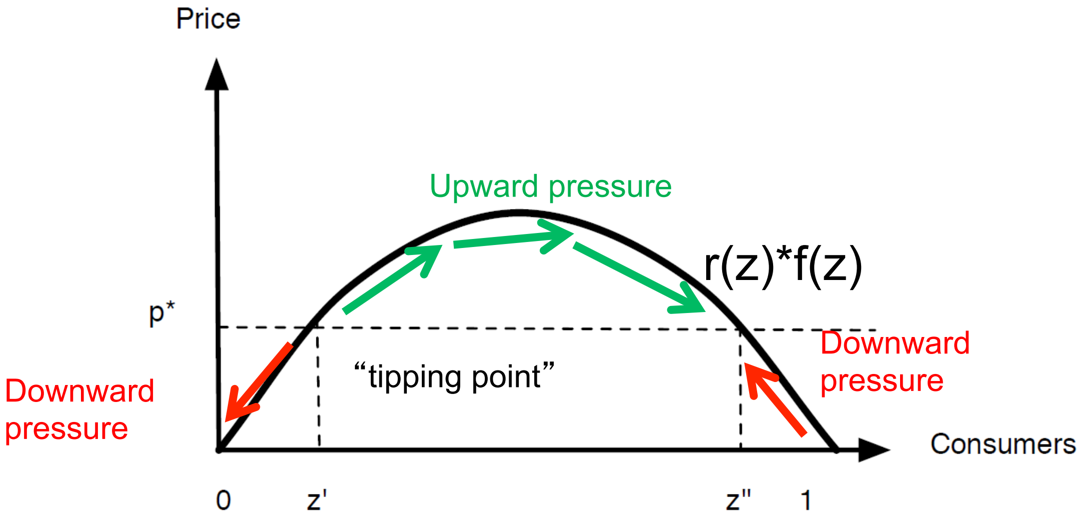
Price
Upward pressure
r(z)*f(z)
p*
Downward pressure
"tipping point"
Downward pressure
Consumers
0
z'
z"
1

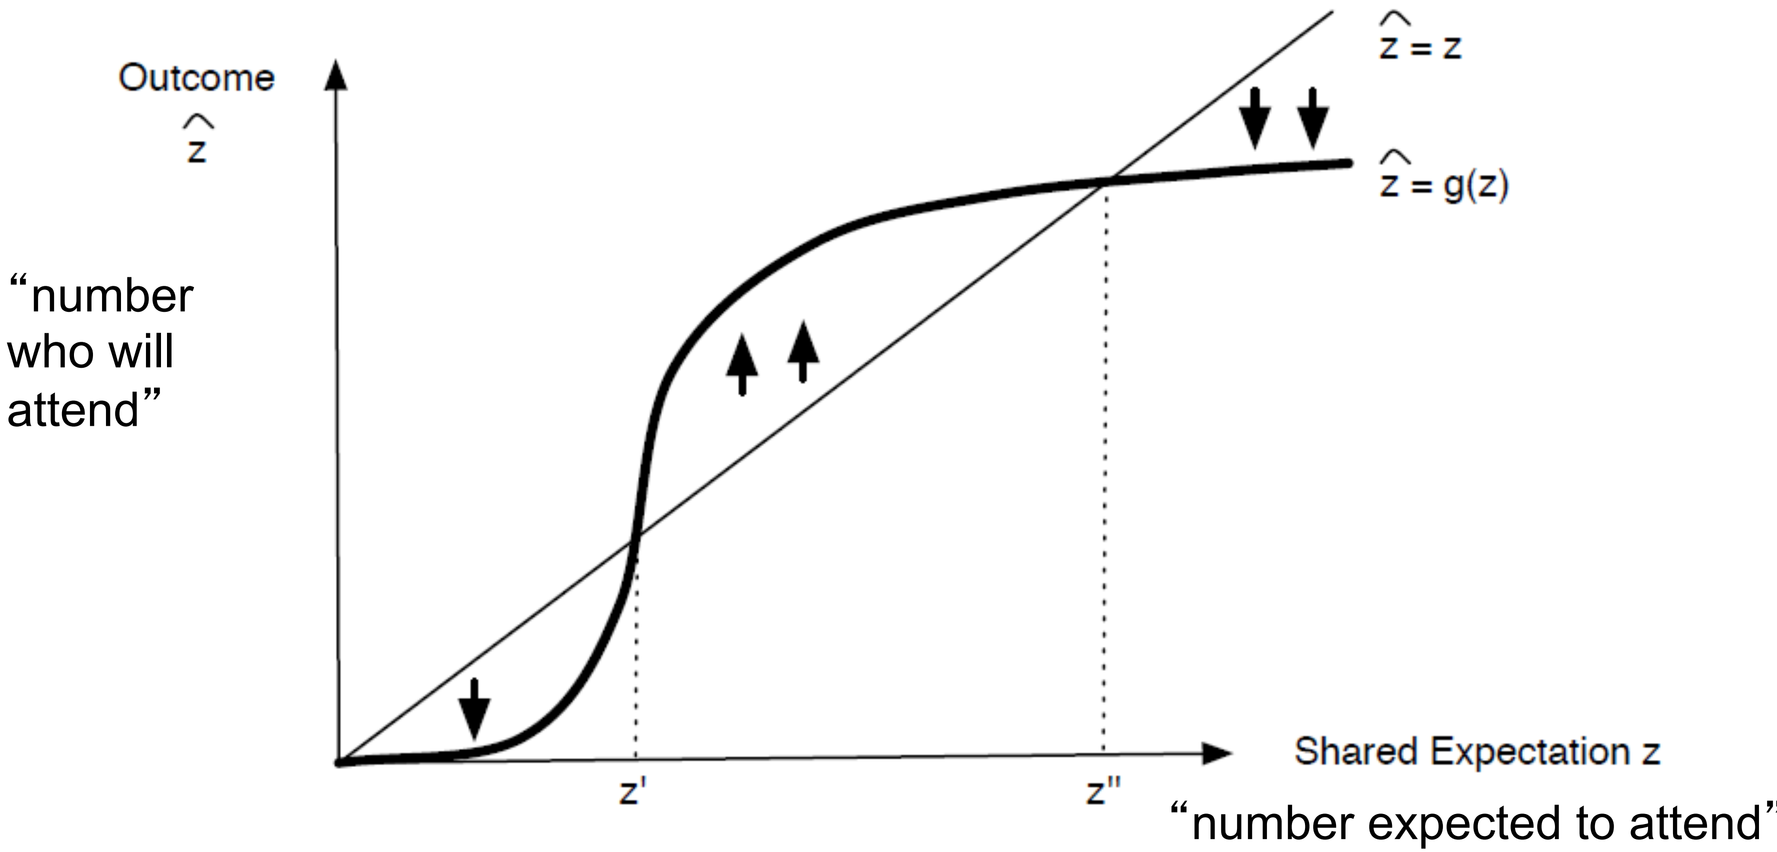
Outcome
$\hat{z}$
$\hat{z} = z$
$\hat{z} = g(z)$
"number who will attend"
Shared Expectation z
z'
z"
"number expected to attend"

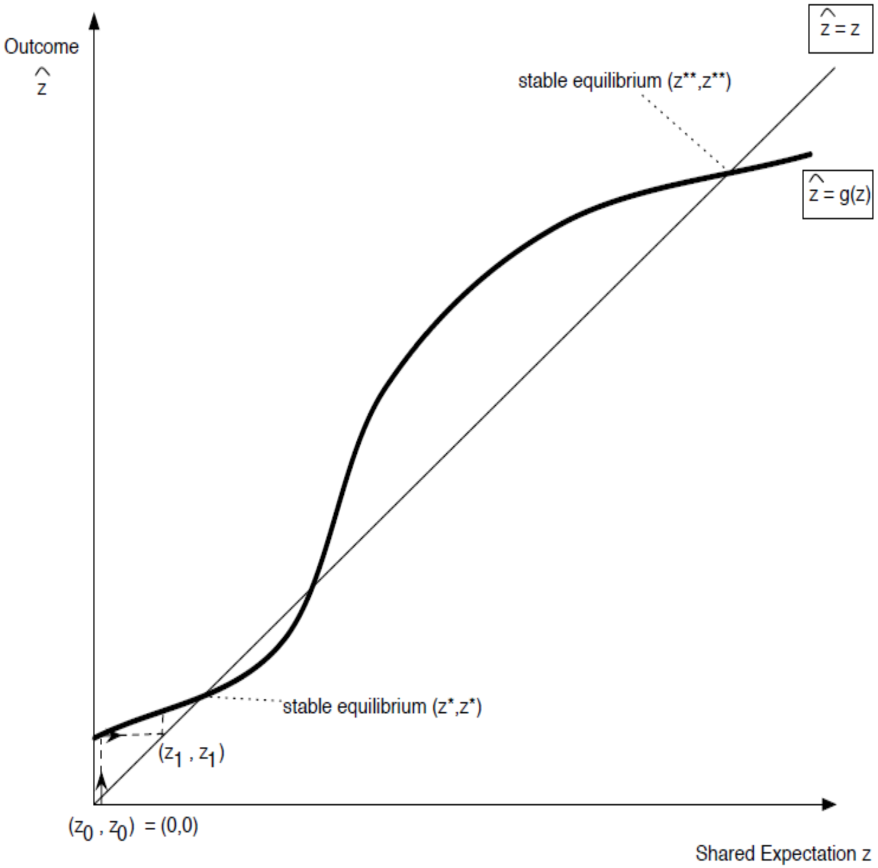

Outcome
$\hat{z}$
$\hat{z} = z$
stable equilibrium (z**,z**)
$\hat{z} = g(z)$
stable equilibrium (z*,z*)
(z_1 , z_1)
$(z_0 , z_0) = (0,0)$
Shared Expectation z

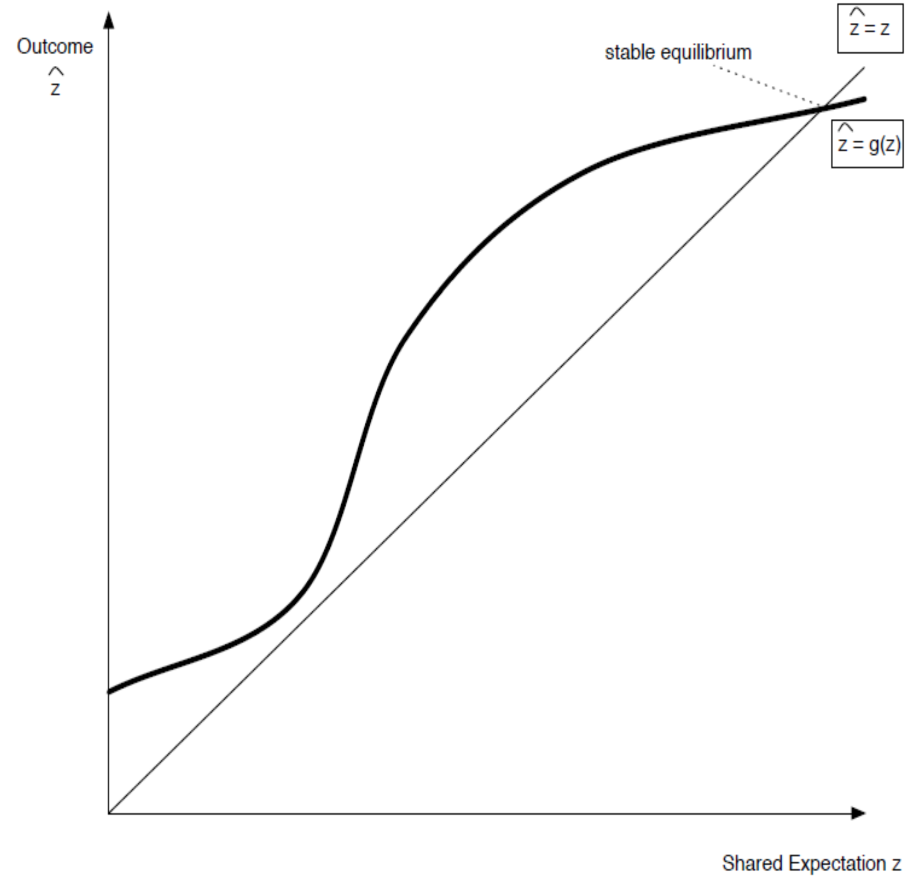

Outcome
$\hat{z}$
$\hat{z} = z$
stable equilibrium
$\hat{z} = g(z)$
Shared Expectation z

Diffusion through social networks: *why* things spread

- Fun: i.e., why do things get popular?
  - Fashion, fads, internet memes, research ideas, …
  - First-order approximation: *preferential attachment* in graphs
- Rational decisions:
  - Decisions made publically with limited information
    - Specifically, decisions where *choice* is **public** but some *evidence* used in the choice is **private**
  - **Decisions made about products** (or behaviors, etc) **that have "network effects"** (aka "externalities")
    - Specifically, the benefits and costs of the behavior are **not completely local** to the decision-maker

Now look at a networked case….

The networked theory

- if v and w both adopt behavior A, they each get a payoff of $a > 0$;
- if they both adopt B, they each get a payoff of $b > 0$; and
- if they adopt opposite behaviors, they each get a payoff of 0.

$$
\begin{array}{cc|c|c|}
 & & \multicolumn{2}{c}{w} \\
 & & A & B \\ \cline{3-4}
\multirow{2}{*}{v} & A & a, a & 0, 0 \\ \cline{3-4}
 & B & 0, 0 & b, b \\ \cline{3-4}
\end{array}
$$

Figure 19.1: A-B Coordination Game

What if *v* is playing the game with many *w*' s ?

If *v* has *d* neighbors and *p*d of them* choose *A*, then *v* should chose *A* iff **pda>-(1-p)db** ie, iff **p>=b/(a+b)**

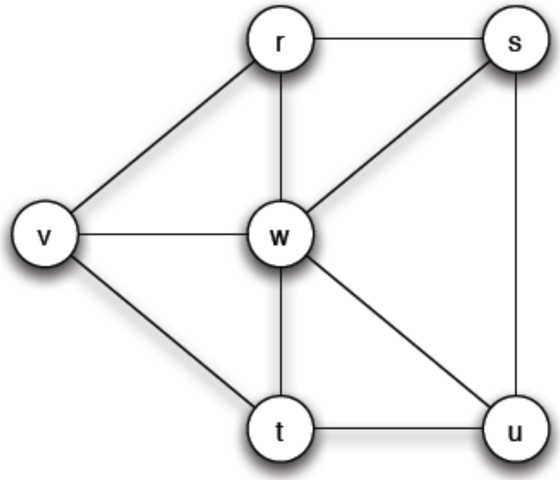

(a) *The underlying network*

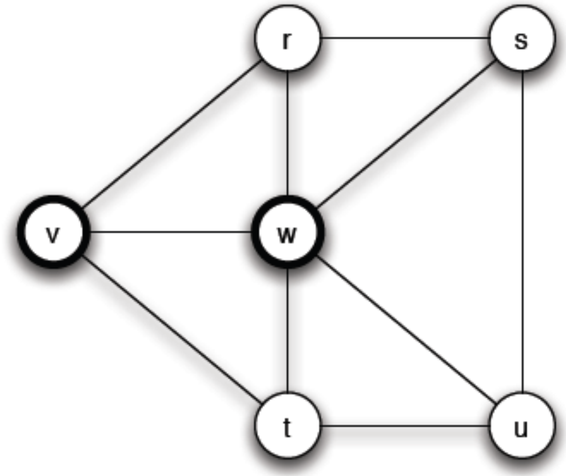

(b) *Two nodes are the initial adopters*

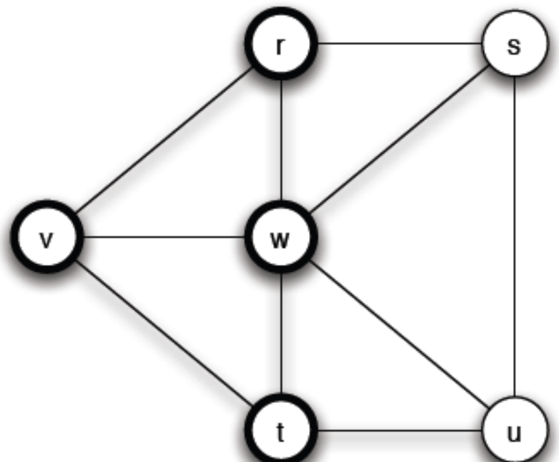

(c) *After one step, two more nodes have adopted*

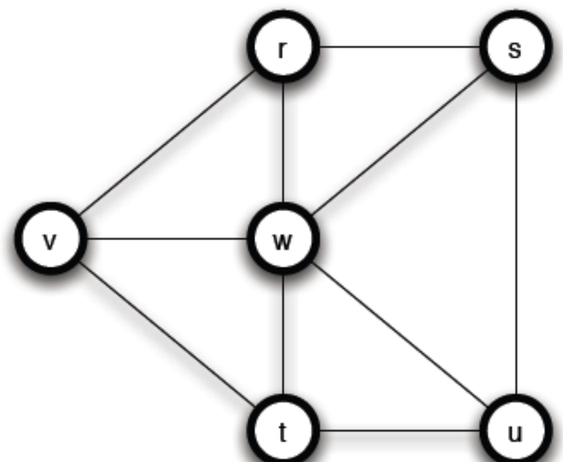

(d) *After a second step, everyone has adopted*

Threshold: switch if 40% of neighbors switched

Figure 19.3: Starting with v and w as the initial adopters, and payoffs $a = 3$ and $b = 2$, the new behavior A spreads to all nodes in two steps. Nodes adopting A in a given step are drawn with dark borders; nodes adopting B are drawn with light borders.

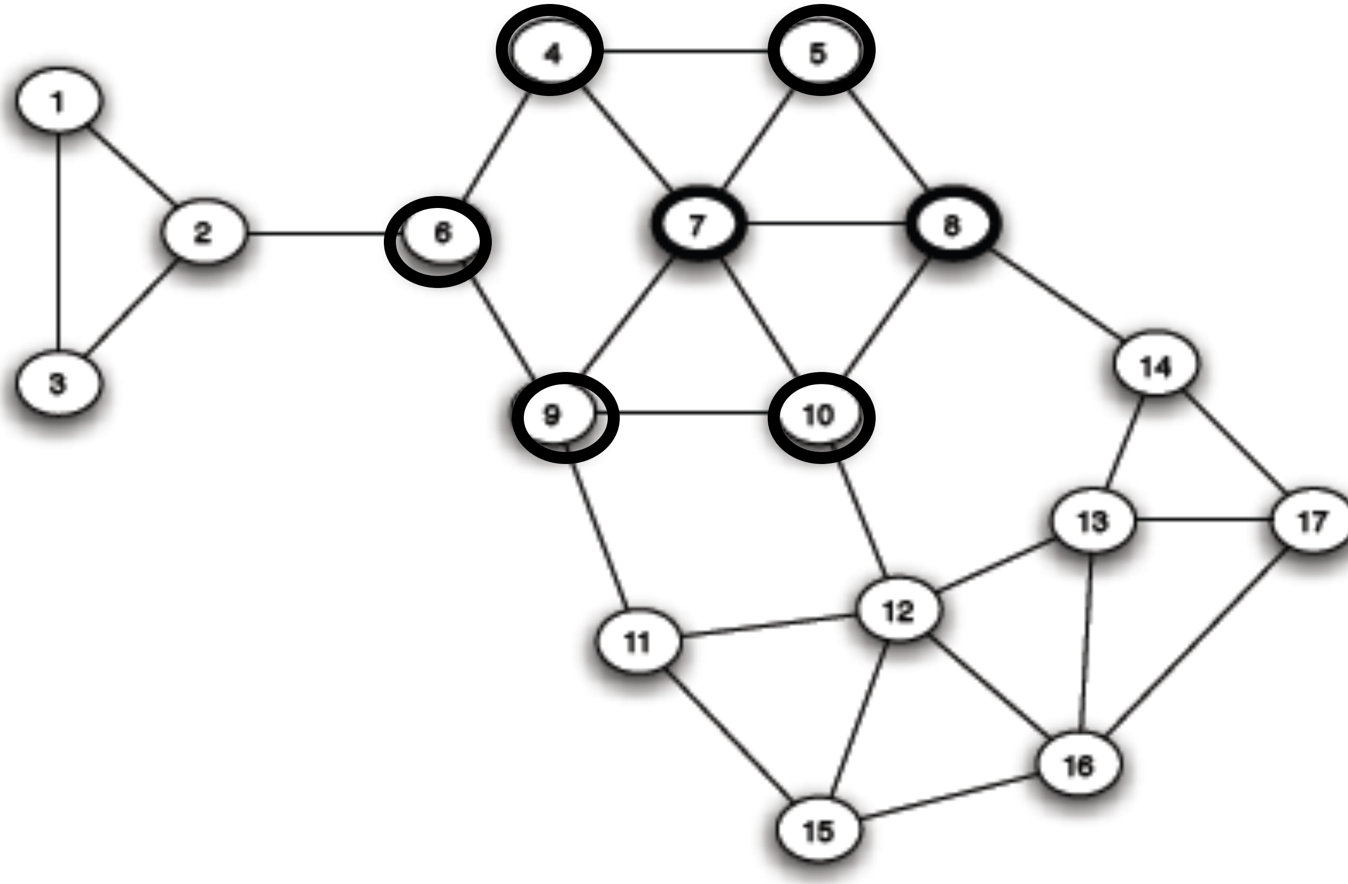

Threshold: switch if 40% of neighbors switched

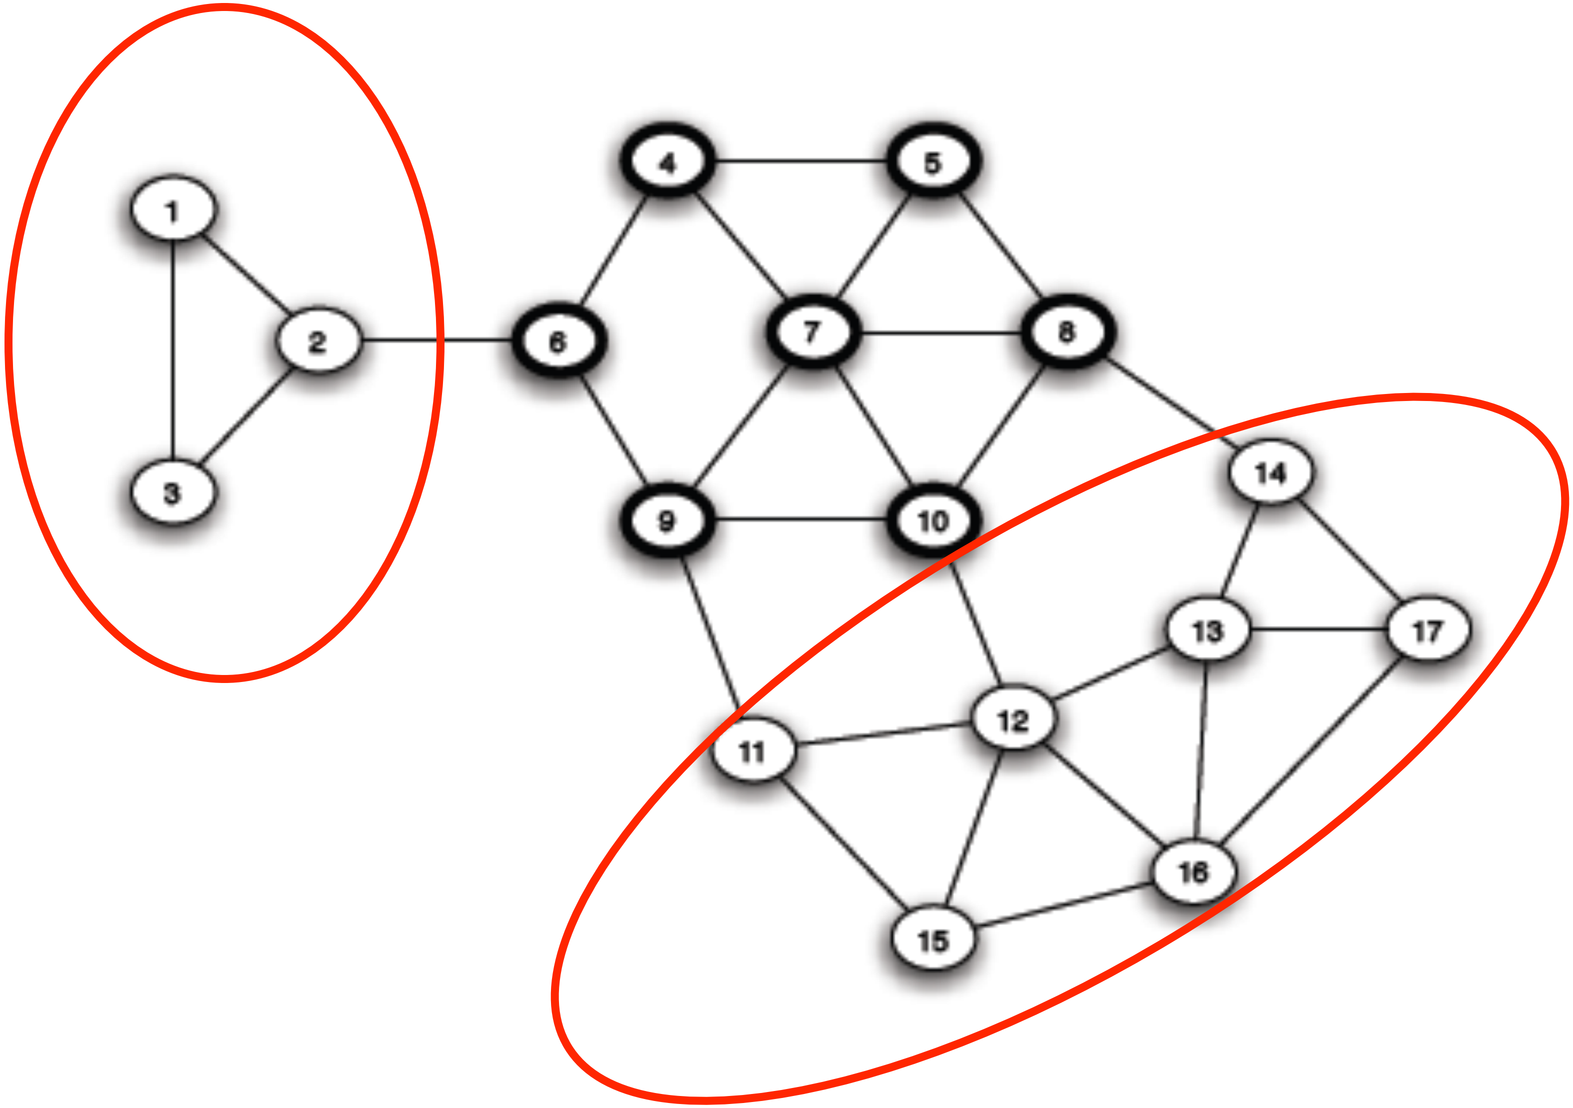

General claim: dense clusters are less susceptible to cascades.

Some simulations and more theory…

Richardson and Domingos

“Mining the Network Value of Customers” – KDD 2001

“Mining Knowledge-Sharing Sites for Viral Marketing” – KDD 2002

Question: who do you *target*?

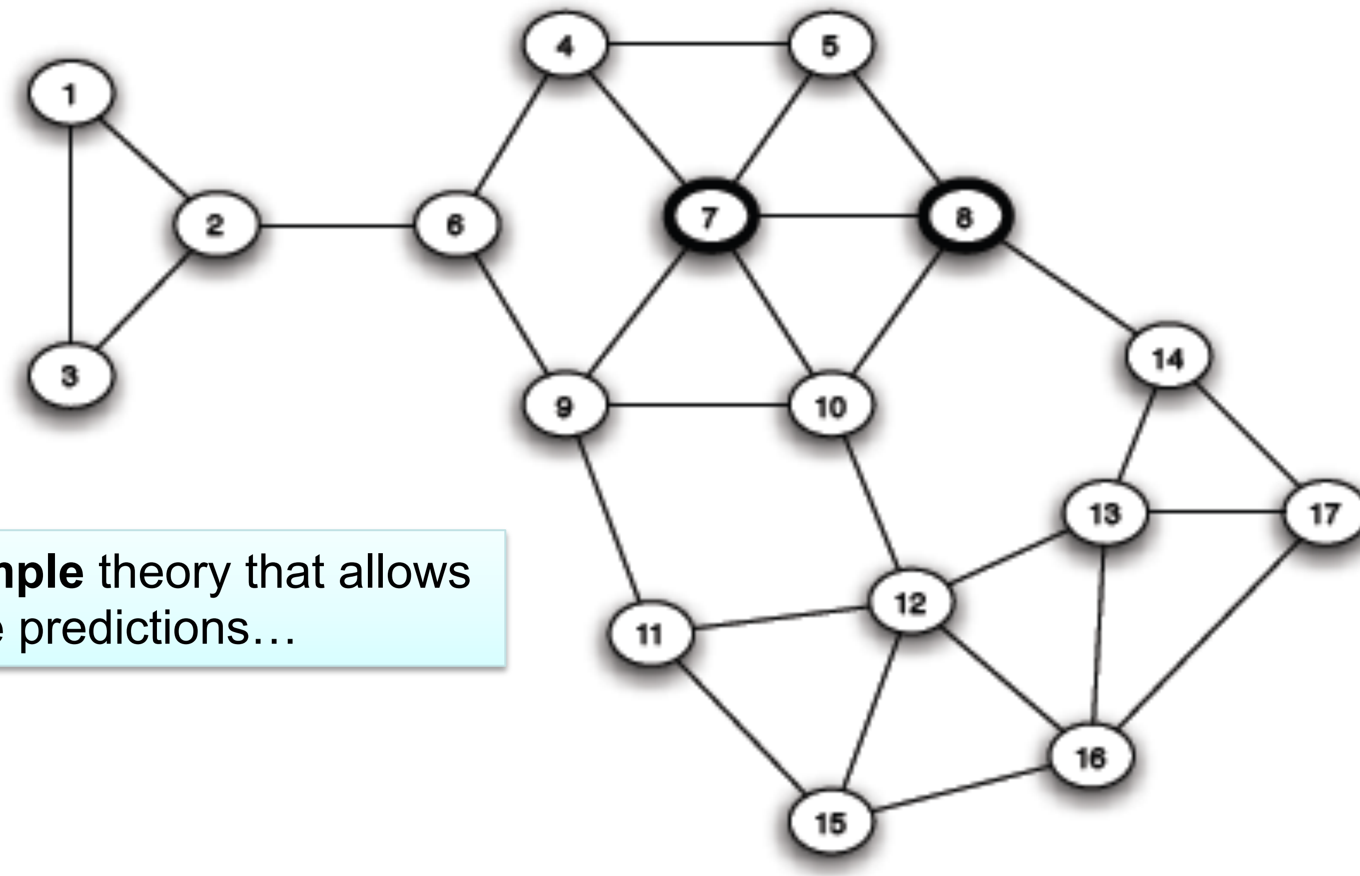

Goal: **simple** theory that allows **tractable** predictions…

Notation

- X_i: did customer i buy it? (yes=1, no=0)
- $\mathbf{N}_i$=neighbors of X_i
- $\mathbf{X}^k$,$\mathbf{X}^u$ = known buyers, unknown buyers
- $\mathbf{Y}$=attributes of product
- M_i=do you market to i? (yes=1, no=0)
- $\mathbf{M}$=all marketing decisions

Model

$$
\begin{aligned}
P(X_i \mid \mathbf{X} - \{X_i\}, \mathbf{Y}, \mathbf{M}) \\
&= P(X_i \mid \mathbf{N_i}, \mathbf{Y}, \mathbf{M}) \qquad (1) \\
&= \beta_i P_0(X_i \mid \mathbf{Y}, M_i) + (1 - \beta_i) P_N(X_i \mid \mathbf{N_i}, \mathbf{Y}, \mathbf{M})
\end{aligned}
$$

internal probability
P0, self-reliance βi

$$
P_N(X_i = 1 \mid \mathbf{N_i}, \mathbf{Y}, \mathbf{M}) = \sum_{X_j \in \mathbf{N_i}} w_{ij} X_j \qquad (2)
$$

$$
\begin{aligned}
P(X_i = 1 \mid \mathbf{Y}, \mathbf{M}) \\
= \beta_i P_0(X_i = 1 \mid \mathbf{Y}, M_i) + (1 - \beta_i) \sum_{X_j \in \mathbf{N_i}} \sum_{\substack{\tilde{\mathbf{N}} \in C(\mathbf{N_i}) \\ \text{with } \tilde{N}_j = 1}} w_{ij} P(\tilde{\mathbf{N}} \mid \mathbf{Y}, \mathbf{M})
\end{aligned}
$$

where $\tilde{N}_j$ is the value of X_j specified by $\tilde{\mathbf{N}}$.

Model

$$P_N(X_i = 1 \mid \mathbf{N_i}, \mathbf{Y}, \mathbf{M}) = \sum_{X_j \in \mathbf{N_i}} w_{ij} X_j \tag{2}$$

$$P(X_i = 1 \mid \mathbf{Y}, \mathbf{M})$$

$$= \beta_i P_0(X_i = 1 \mid \mathbf{Y}, M_i) + (1 - \beta_i) \sum_{X_j \in \mathbf{N_i}} \sum_{\substack{\tilde{\mathbf{N}} \in C(\mathbf{N_i}) \\ \text{with } \tilde{N}_j = 1}} w_{ij} P(\tilde{\mathbf{N}} \mid \mathbf{Y}, \mathbf{M})$$

where $\tilde{N}_j$ is the value of X_j specified by $\tilde{\mathbf{N}}$.

$$P(X_i = 1 \mid \mathbf{Y}, \mathbf{M})$$

$$= \beta_i P_0(X_i = 1 \mid \mathbf{Y}, M_i) + (1 - \beta_i) \sum_{X_j \in \mathbf{N_i}} w_{ij} P(X_j = 1 \mid \mathbf{Y}, \mathbf{M}) \tag{4}$$

PageRank-like recurrence

Model

PageRank-like recurrence

$$P(X_i = 1 \mid \mathbf{Y}, \mathbf{M})$$
$$= \beta_i P_0(X_i = 1 \mid \mathbf{Y}, M_i) + (1 - \beta_i) \sum_{X_j \in \mathbf{N_i}} w_{ij} P(X_j = 1 \mid \mathbf{Y}, \mathbf{M}) \quad (4)$$

Definitions:

- c = cost of marketing to any i
- r0 = revenue without marketing to i
- r1 = revenue with marketing to I
- expected lift in profit from marketing to I is

change M_i to 1, leave rest of **M** unchanged

change M_i to 0…

$$ELP_i^1(\mathbf{Y}, \mathbf{M}) = r_1 P(X_i = 1 \mid \mathbf{Y}, f_i^1(\mathbf{M})) - r_0 P(X_i = 1 \mid \mathbf{Y}, f_i^0(\mathbf{M})) - c$$

Model

Goal: If $\mathbf{M_0}$ is no marketing, maximize:

$$ELP(\mathbf{Y},\mathbf{M}) =$$

$$\sum_{i=1}^{n}\left[r_i P(X_i = 1 \mid \mathbf{Y},\mathbf{M}) - r_0 P(X_i = 1 \mid \mathbf{Y},\mathbf{M}_0) - c_i\right]$$

where $r_i = r_1$ and $c_i = c$ if $M_i = 1$, and $r_i = r_0$ and $c_i = 0$ if $M_i = 0$.

Definitions:

- c = cost of marketing to any i
- r0 = revenue without marketing to i
- r1 = revenue with marketing to I
- expected lift in profit from marketing to I is

$$ELP_i^1(\mathbf{Y},\mathbf{M}) = r_1 P(X_i = 1 \mid \mathbf{Y}, f_i^1(\mathbf{M})) - r_0 P(X_i = 1 \mid \mathbf{Y}, f_i^0(\mathbf{M})) - c$$

Model

Goal: If $\mathbf{M_0}$ is no marketing, maximize:

$$ELP(\mathbf{Y},\mathbf{M}) = \sum_{i=1}^{n}\left[r_i P(X_i = 1 \mid \mathbf{Y},\mathbf{M}) - r_0 P(X_i = 1 \mid \mathbf{Y},\mathbf{M}_0) - c_i\right]$$

where $r_i=r_1$ and $c_i=c$ if $M_i=1$, and $r_i=r_0$ and $c_i=0$ if $M_i=0$.

Extension: assume marketing actions are *continuous* and response is *linear:*

$$ELP_i^z(\mathbf{Y},\mathbf{M}) = r(z)P(X_i = 1 \mid \mathbf{Y}, f_i^z(\mathbf{M})) - r(0)P(X_i = 1 \mid \mathbf{Y}, f_i^0(\mathbf{M})) - c(z) \qquad (5)$$

Model

Goal: If $\mathbf{M_0}$ is no marketing, maximize:

$$ELP(\mathbf{Y},\mathbf{M}) =$$

$$\sum_{i=1}^{n}\left[r_i P(X_i=1\mid\mathbf{Y},\mathbf{M}) - r_0 P(X_i=1\mid\mathbf{Y},\mathbf{M}_0) - c_i\right]$$

where $r_i=r_1$ and $c_i=c$ if $M_i=1$, and $r_i=r_0$ and $c_i=0$ if $M_i=0$.

Key point: the **network effect** of marketing to Xi has a *linear* effect on the rest of the network….so you can prove:

$$P(X_i=1\mid\mathbf{Y},\mathbf{M})$$

$$=\beta_i P_0(X_i=1\mid\mathbf{Y},M_i)+(1-\beta_i)\sum_{X_j\in\mathbf{N_i}} w_{ij}\,P(X_j=1\mid\mathbf{Y},\mathbf{M}) \quad (4)$$

network P

$$\Delta_i(\mathbf{Y}) = \sum_{j=1}^{n}\frac{\partial P(X_j=1\mid\mathbf{Y},\mathbf{M}_0)}{\partial P_0(X_i=1\mid\mathbf{Y},M_i)} = \sum_{j=1}^{n} w_{ji}\Delta_j(\mathbf{Y})$$

non-network i/P_0

(Needs proof)

Model

- With linearity network *effect* doesn't depend on **M**
  - Network *value* depends on **M,** also susceptibility to marketing, cost of marketing to I, …
- With linearity we can estimate network *effect* quickly
- If we assume revenue doesn't depend on **M** (advertising only, no discounts) then we can build on this to compute network *value* and *ELP* from marketing to *I*

- *Without linearity:* this story gets complicated fast (KDD 2001 paper)

Experiments

- Mine Epinions for network (trust ratings)
- Assume uniform w_{ij} weights, constant? self-reliance, and NB model of internal probability (estimated from “purchases” of products, equating review=purchase)
- Vary effectiveness of marketing strategy alpha, revenue, and cost of marketing

Sample result---network values

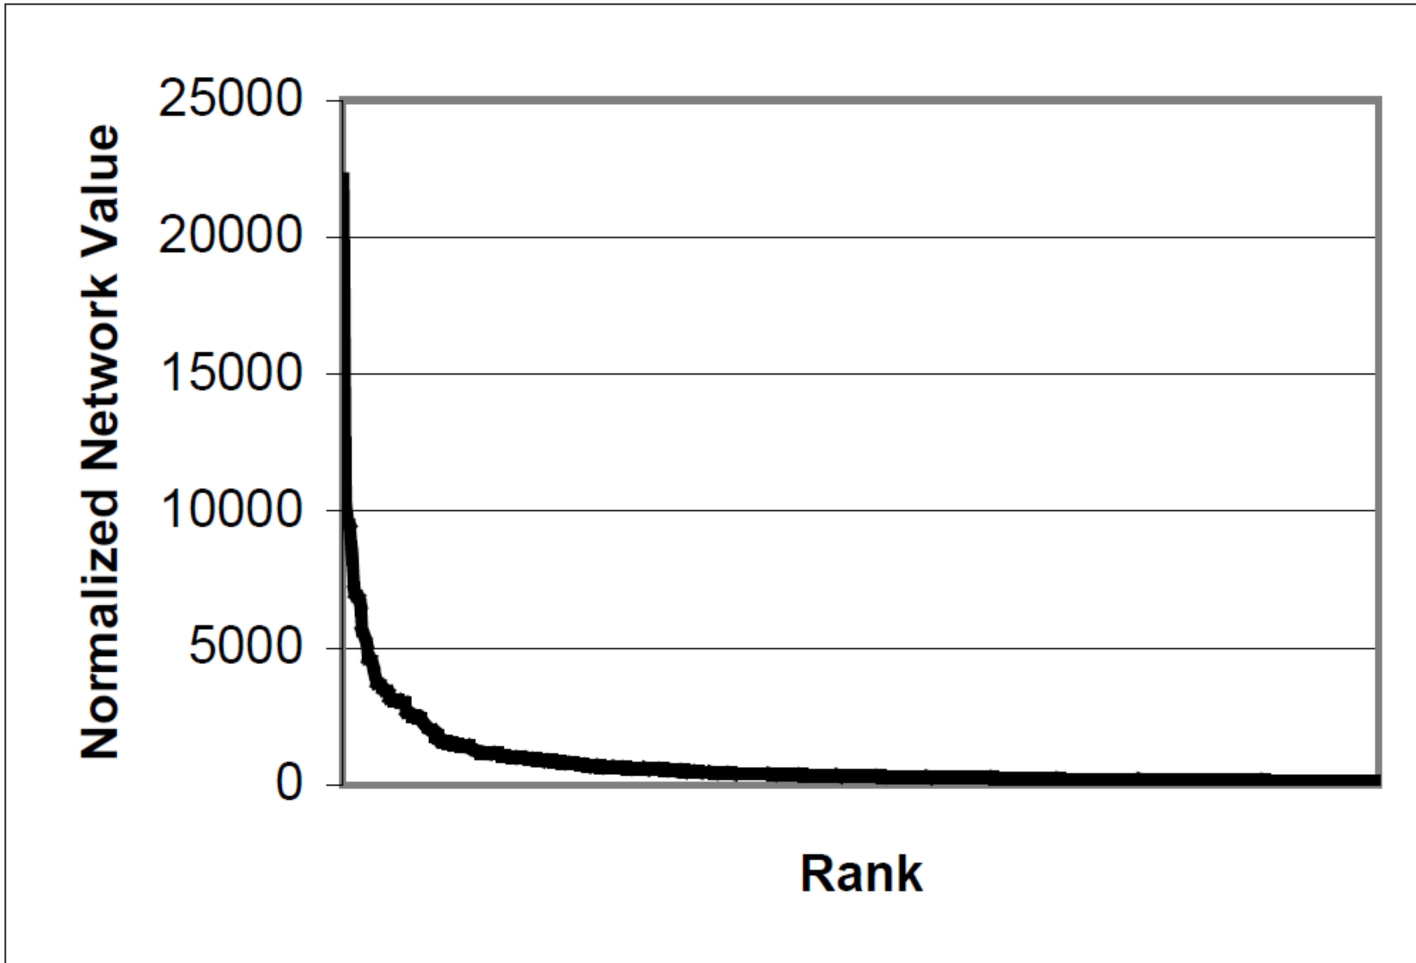

Figure 1: Typical distribution of network value.

Sample result---profits

Table 1: Profit results for Boolean marketing scenario for various costs of marketing.

	$\alpha=2, r_0=1, r_1=1$		
	$c = 0.1$	$c = 0.01$	$c = 0.001$
No Marketing	37.78	37.78	37.78
Direct Marketing	37.78	42.71	66.08
Viral Marketing	47.25	60.54	70.23

Sample result--robustness

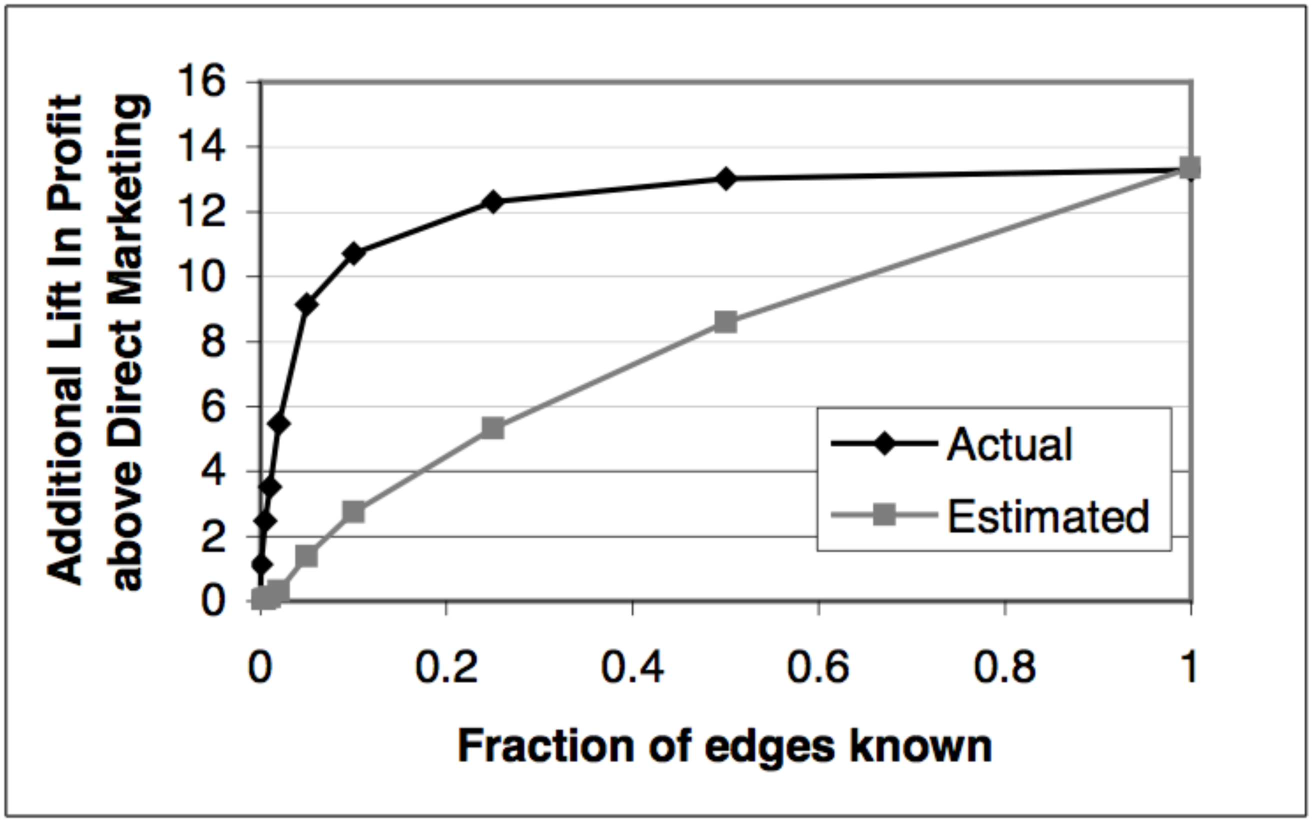

Figure 3: Actual and estimated difference between viral marketing and direct marketing profits with only partial network knowledge.

Some real-world experiments

Statistical Science
2006, Vol. 21, No. 2, 256–276
DOI: 10.1214/088342306000000222

Network-Based Marketing: Identifying Likely Adopters via Consumer Networks

Shawndra Hill, Foster Provost and Chris Volinsky

Opportunity

- Companies like AT&T sell products (e.g., data services, ringtones, ….)
- … and have (partial) network data

- Can you use network data to do better marketing?

Existing marketing approach

This section details our data set, derived primarily from a direct-mail marketing campaign to potential customers of a new communications service (later we augment the primary data with a large set of consumer-specific attributes). The firm's marketing team identified and marketed to a list of prospects using its standard methods.

actions

TABLE 1
Descriptive statistics for the marketing segments (see Section 4.1 for details)

Segment	Loyalty	Intl	Tech1	Tech2	Early Adopt	Offer	% of list
1	3	Y	Hi	1–7	Med–Hi	P1	1.6
2	3	Y	Med	1–7	Med–Hi	P1	2.4
3	2	Y	Hi	1–4	Hi	P1	1.7
4	2	Y	Med	1–4	Hi	P1	1.7
5	1	Y	Hi	1–4	Hi	P1	0.1
6	1	Y	Med	1–4	Hi	P1	0.1
7	3	N	Hi	1–7	Med–Hi	P2	10.9
8	3	N	Med	1–7	Med–Hi	P2	13.1
9	2	N	Hi	1–4	Hi	P2	17.5
10	2	N	Med	1–4	Hi	P2	11.0

Opportunity

- Hypothesis:
  - someone that has communicated with a current subscriber (of the new service) is more likely to adopt it

  - model communication with an existing subscriber as a **binary** flag (network neighbor)

Opportunity?

0.3% are NN

Descriptive statistics for the marketing segments (see Section 4.1 for details)

Segment	Loyalty	Intl	Tech1	Tech2	Early Adopt	Offer	% of list	%NN
1	3	Y	Hi	1–7	Med–Hi	P1	1.6	0.63
2	3	Y	Med	1–7	Med–Hi	P1	2.4	1.26
3	2	Y	Hi	1–4	Hi	P1	1.7	0.08
4	2	Y	Med	1–4	Hi	P1	1.7	0.10
5	1	Y	Hi	1–4	Hi	P1	0.1	0.22
6	1	Y	Med	1–4	Hi	P1	0.1	0.25
7	3	N	Hi	1–7	Med–Hi	P2	10.9	0.50
8	3	N	Med	1–7	Med–Hi	P2	13.1	0.83
9	2	N	Hi	1–4	Hi	P2	17.5	0.04
10	2	N	Med	1–4	Hi	P2	11.0	0.07
11	1	N	Hi	1–4	Hi	P2	5.3	0.14
12	1	N	Med	1–4	Hi	P2	7.7	0.25
13	3	N	Hi	1–7	Med–Hi	P2	2.0	0.63
14	1, 2	N	Hi	1–4	Hi	P2	2.0	0.15
15	1	Y	?	?	?	P3	2.0	1.01
16	1	N	?	?	?	P2	1.6	0.46
17	3	N	Hi	1–7	Med–Hi	P2+	2.0	0.70
18	1, 2	N	Hi	1–4	Hi	P2+	2.0	0.15
19	1, 2, 3	Y	Hi	1–7	Med–Hi	P3	1.8	0.67
20	2	N	Hi, Med	1–4	Hi	L1	6.0	0.05
21	2	N	Hi, Med	1–4	Hi	L2	6.0	0.05

Experiment 1: Use NN flag to predict "takes" for the offer for each segment

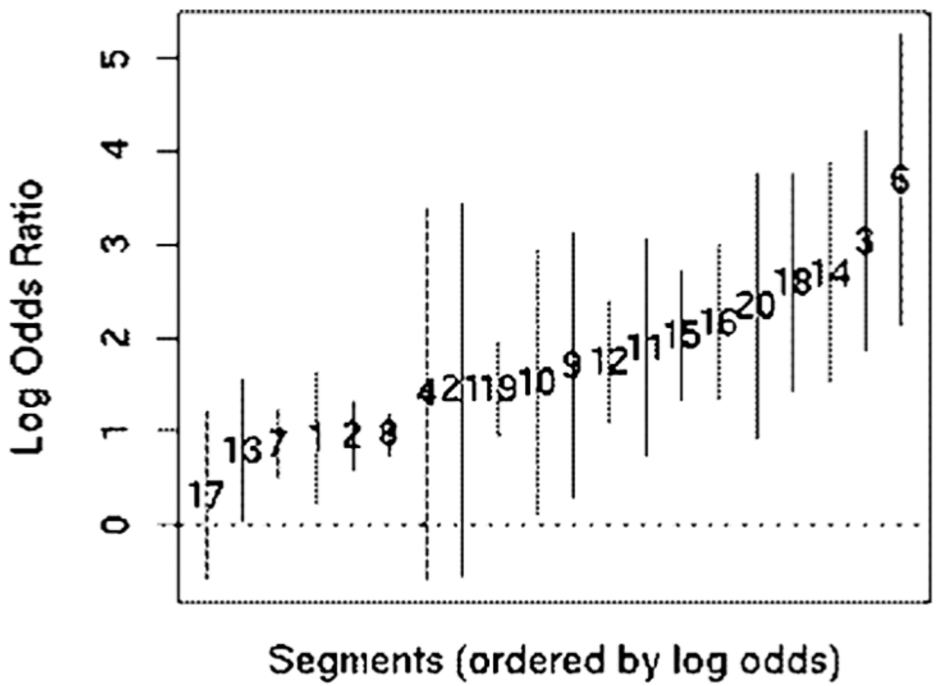

FIG. 2. *Results of logistic regression. Parameter estimates plotted as log-odds ratios with 95% confidence intervals. The number plotted at the value of the parameter estimate refers back to segment numbers from Table* 1.

Experiment 1: Use NN flag to predict "takes" for the offer for each segment

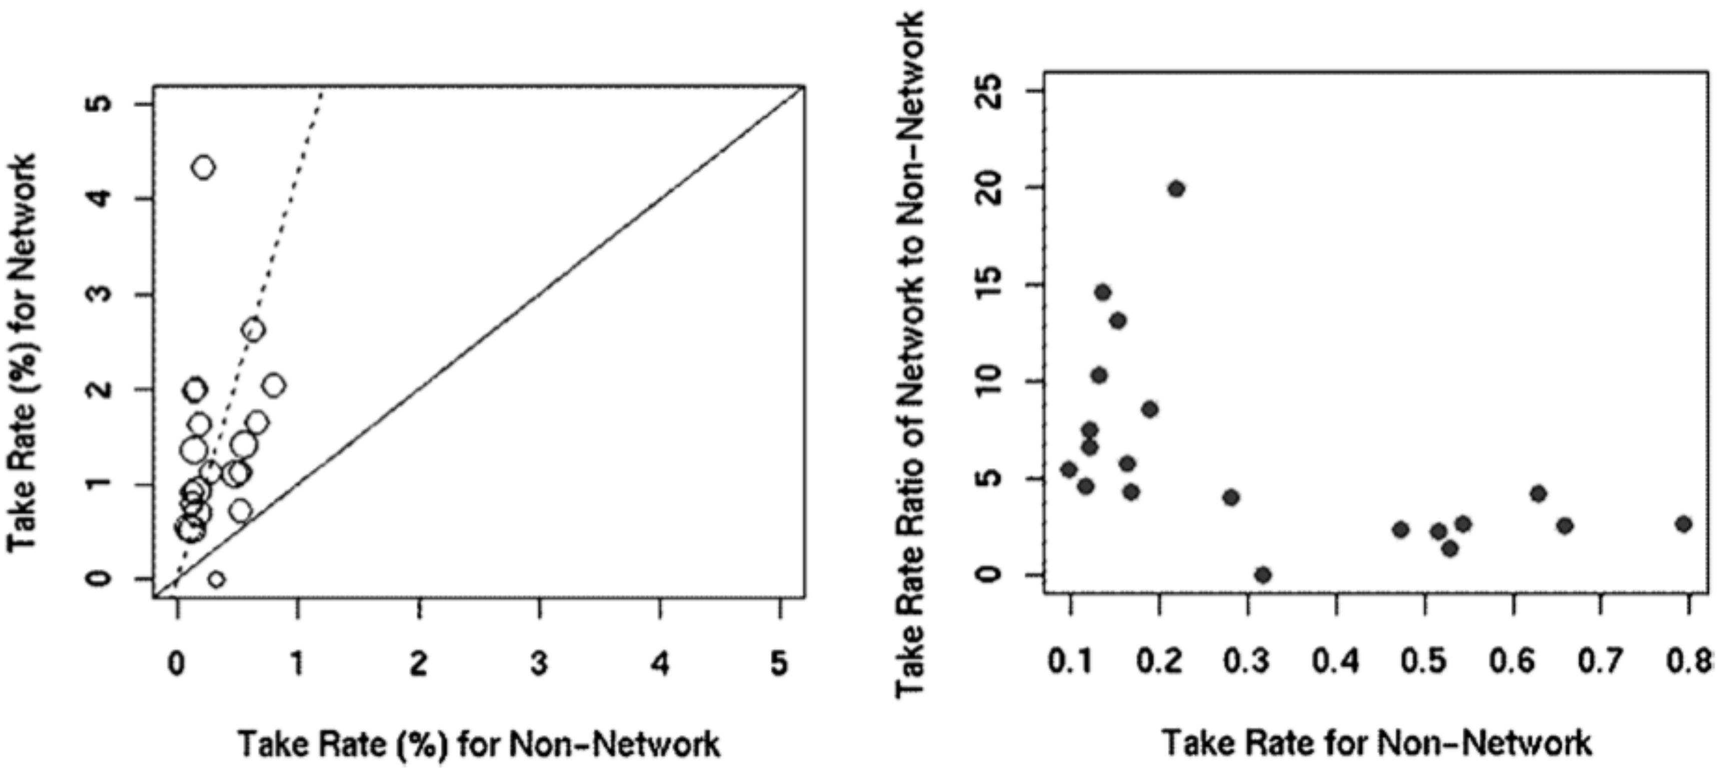

FIG. 3. *Take rates for marketing segments. Left: For each segment, comparison of the take rate of the non-network neighbors with that of the network neighbors. The size of the glyph is proportional to the log size of the segment. There is one outlier not plotted, with a take rate of* 11% *for the network neighbors and* 0.3% *for the non-network neighbors. Reference lines are plotted at* $x = y$ *and at the overall take-rate ratio of* 3.4. *Right: Plot of the take rate for the non-network group versus lift ratio for the network neighbors.*

Experiment 2: Market to "segment 22" (near-misses to segments 1-21 + NN)

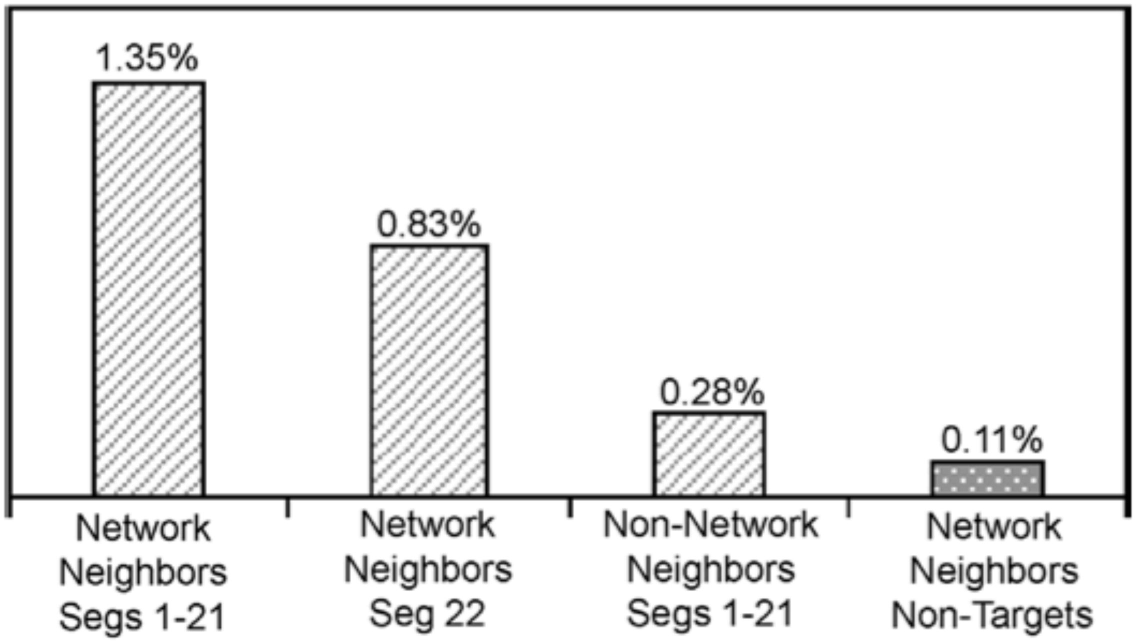

FIG. 4. *Take rates for marketing segments. Take rates for the network neighbors and non-network neighbors in segments* 1–21 *compared with the all-network-neighbor segment* 22 *and with the nontarget network neighbors. All take rates are relative to the non-network-neighbor group* (*segments* 1–21).

Statistical Science
2006, Vol. 21, No. 2, 256–276
DOI: 10.1214/088342306000000222

Network-Based Marketing: Identifying Likely Adopters via Consumer Networks

Shawndra Hill, Foster Provost and Chris Volinsky

It works…but we can't tell you

- what the product was (does it have a network effect?)
- whether this was really worth bothering with (only 0.3% of original market)

And….this isn't really "viral" since there's no iteration

“Big Seed” marketing and network multipliers

Viral Marketing for the Real World

Duncan J. Watts, Jonah Peretti, and Michael Frumin

“Big Seed” marketing

- Suppose you sell a product to K people
  - and each person sells it to R friends
  - …and they sell it to R friends…

- What’s the size of the market?
  - $K^*(1 + R + R^2 + R^3 + \ldots)$

“Big Seed” marketing

- Suppose you sell a product to K people
  - and each person sells it to R friends
  - …and they sell it to R friends….

- What’s the size of the market?
  - $K^*(1 + R + R^2 + R^3 + \ldots) = \underline{K/(1-R)}$

    assuming $R<1$

For R=0.5, your marketing power is doubled
For R=0.9, your marketing power is increased by 10x
For R=0.1 your marketing power is increased by 10%

“Big Seed” marketing

- ForwardTrack
  - designed to encourage viral campaigns
  - participants can tell a friend and watch their “cascades” grow

“Big Seed” marketing

ForwardTrack

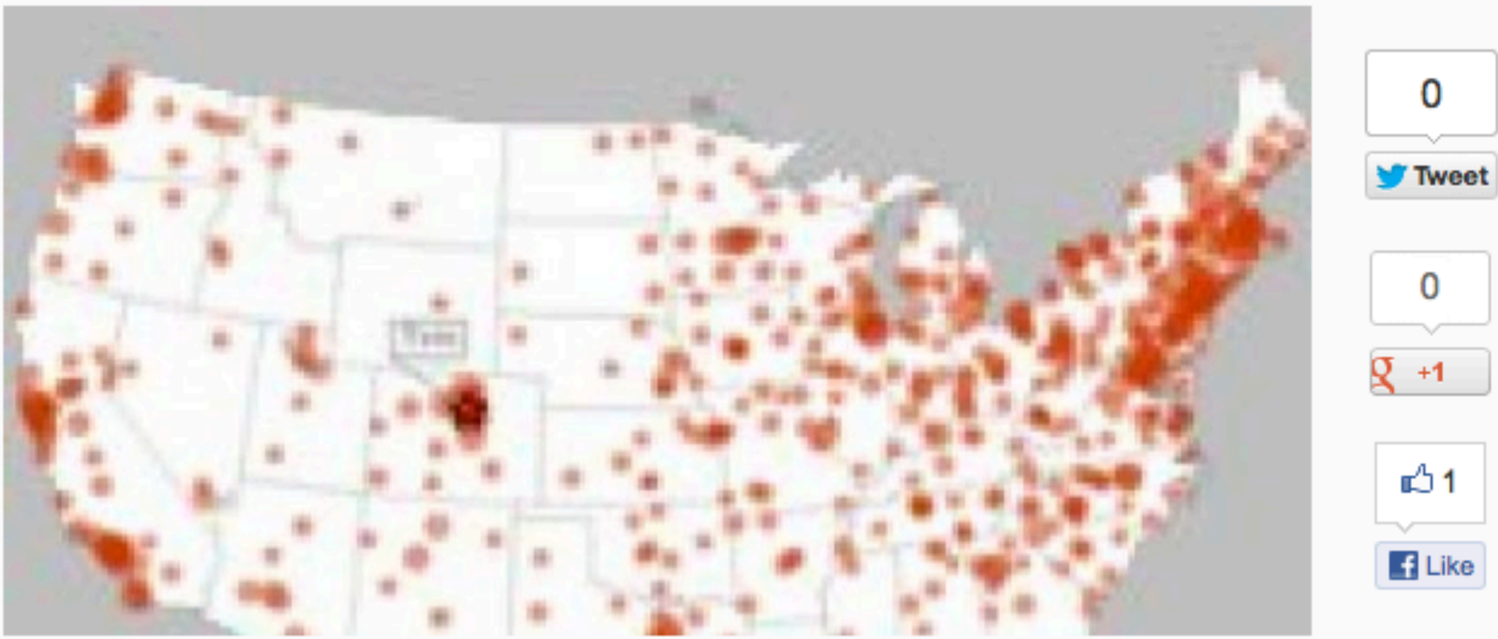

The latest release from Eyebeam R&D is ForwardTrack, a system for tracking and mapping the circulation of email forwards, political calls-to-action and petitions. ForwardTrack is designed to encourage activism by graphically revealing the power of social-networks and demonstrating the impact of the individual's voice in the political process. ForwardTrack is currently being beta tested via Tom's Petition and is available to the public in an open source version at the ForwardTrack site.

“Big Seed” marketing

Drowning NYC – Recombinant Fiction

0
Tweet
0
+1
0
Like

Drowning NYC is a transmedia storytelling; an experimental pilot of a story that is told by actors and narrative devices staged over the Internet and in public spaces of a few selected New York City neighborhoods. The story informs the audiences about rising sea levels due to global warming and how urban populations will cope with it. The genre is theorized by the artist as Recombinant Fiction, a political and pervasive form of transmedia fiction.
This project proposes new pedagogical instruments, innovative activist strategies, elaborate media experiments, cutting-edge forms of theatre and cinema, questions about reality perception/construction.

http://www.drowning-nyc.net
http://futurewaterproofcorp.com

Degree of Separation from Root	StopTheNRA: Tom's Petition	Oxygen Network: Katrina Benefit	P&G: Tide Coldwater Challenge[4]
1 (Seeds)	22,582[1]	7,064	960,954
2	10,698	5,298	34,679
3	6,979	4,087	4,846
4	4,798	3,533	913
5	9,115[2]	2,403	188
6		2,374	38
7		2,039	12
8		1,431	3
9		899	
10		593	
11		481	
12		233	
13		91	
14		55	
15		21	
16		4	
17		2	
Seed	**22,582**	**7,064**	**960,954**
Total Reached	**54,172**	**30,608**	**1,001,633**
Bonus	**31,590**	**23,544**	**40,679**
Gain	**2.399**	**4.333**	**1.042**
R	**0.583** [3]	**0.769**	**0.041** [5]

Analysis of marketing cascades

The Dynamics of Viral Marketing

JURE LESKOVEC
Carnegie Mellon University
LADA A. ADAMIC
University of Michigan
and
BERNARDO A. HUBERMAN
HP Labs

Analysis of marketing cascades

- Dataset:
  - after *s* purchases product *p*, she can send recommendations to her friends *n1,n2,…*
  - first recommendee *ni* to purchase p gets a 10% discount
  - and sender *x* also gets 10% discount
  - everything is tracked and timestamped
  - products have types (DVDs, …) and categories
  - Size: about 500k products, 4M people, 15M recommendations, 100k takes

Analysis of marketing cascades

Lognormal/Power-law for number of recommendations

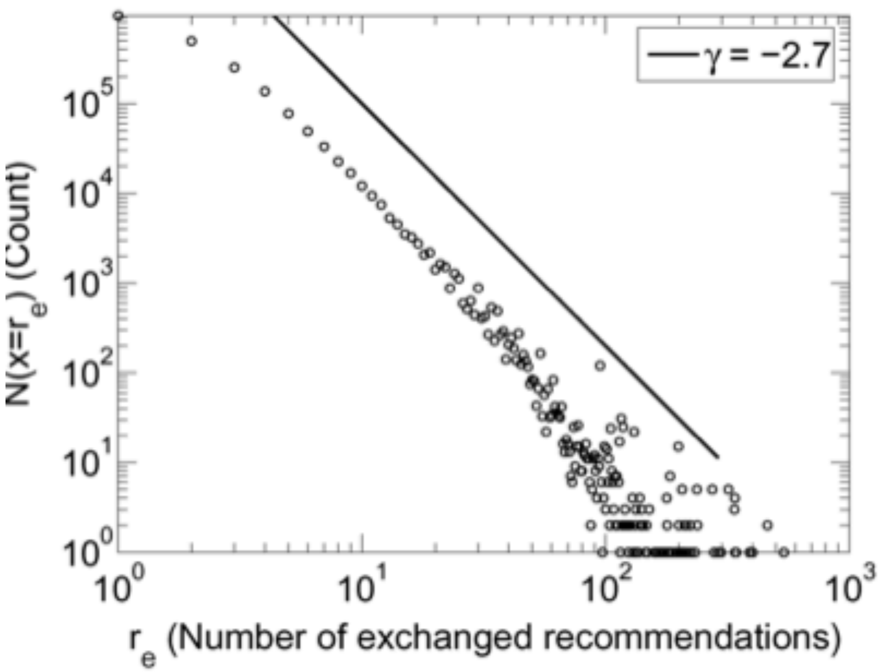

(b) Books

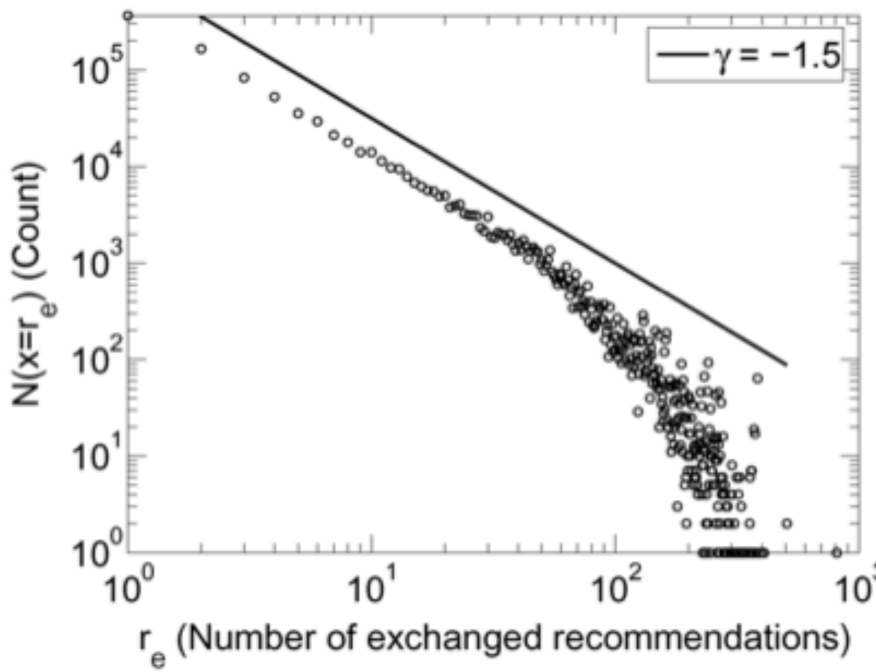

(c) DVD

Analysis of marketing cascades

LCC grows only to about 2.5% of all nodes

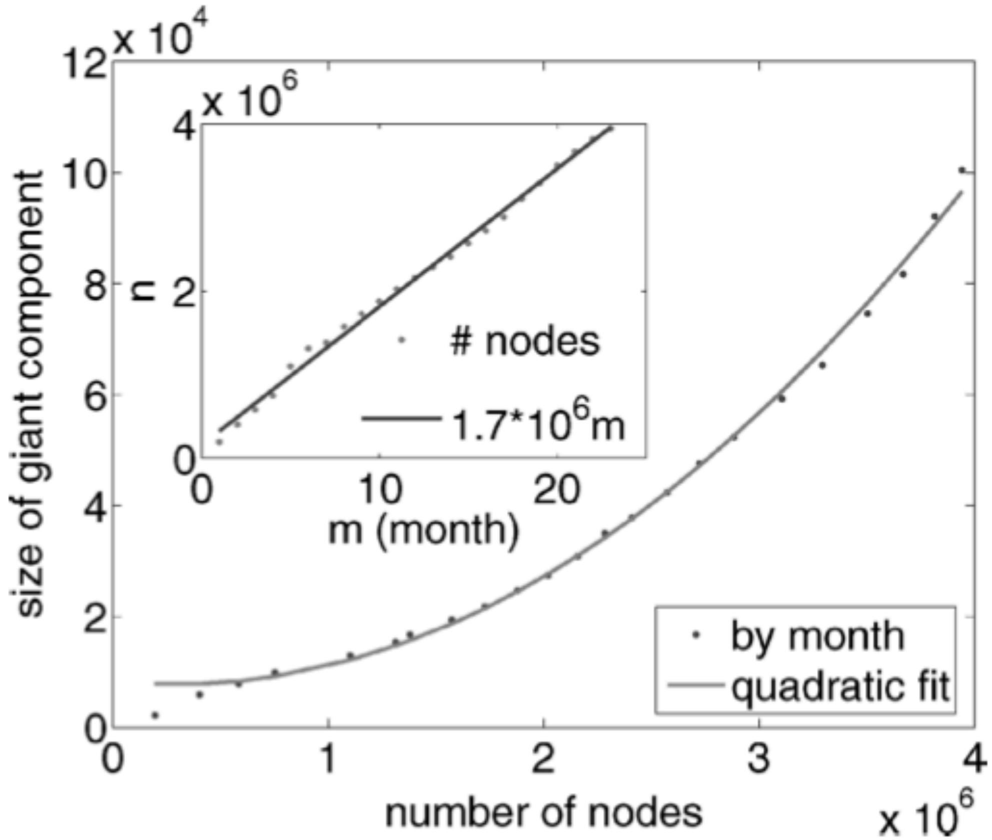

Analysis of marketing cascades

Most recommendation edges are between small clusters

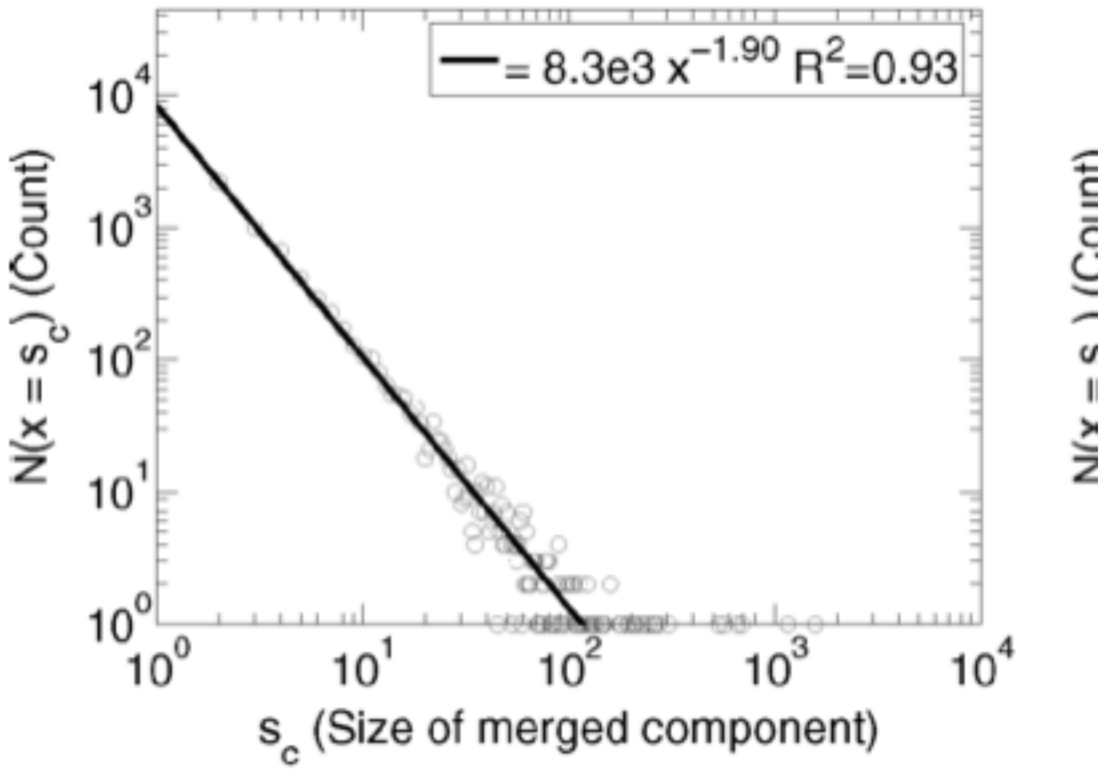

(a) LCC growth

Sample recommendation CC’s

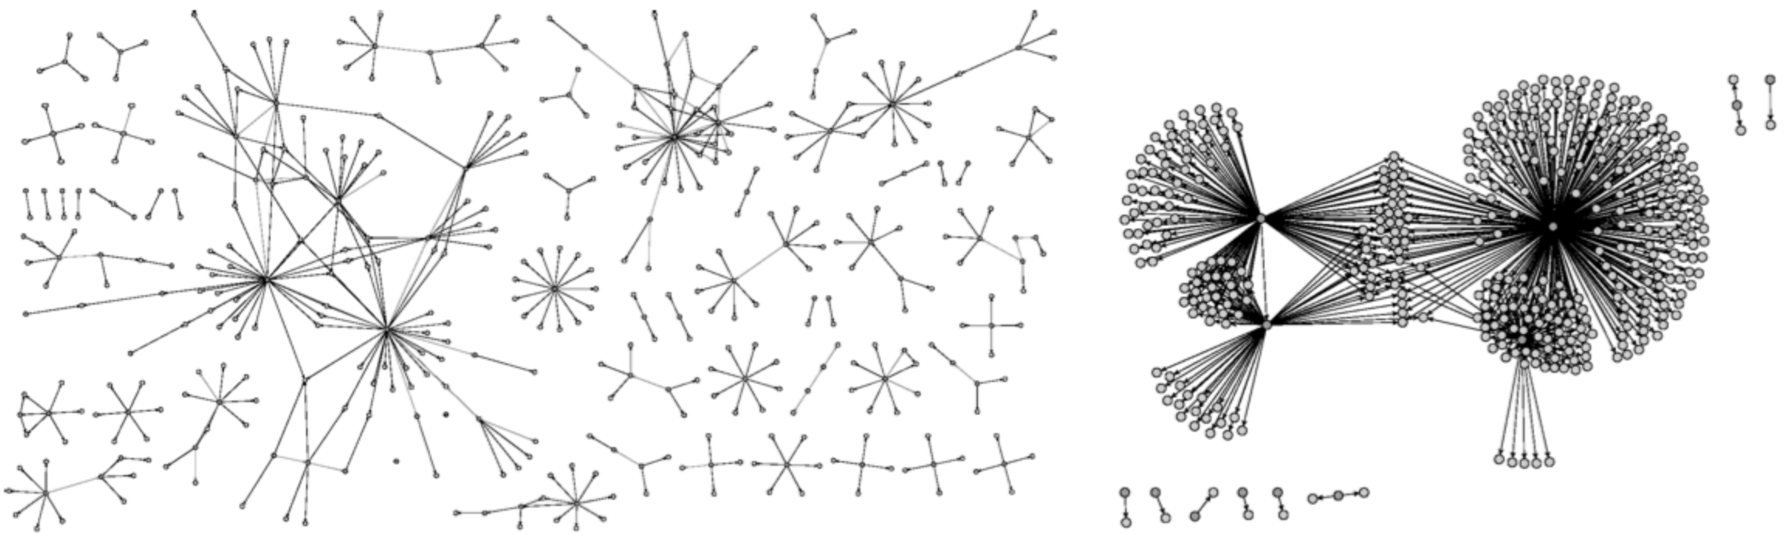

(a) Medical book (b) Japanese graphic novel

Fig. 3. Examples of two product recommendation networks: (a) First-aid study guide *First Aid for the USMLE Step*, (b) Japanese graphic novel (manga) *Oh My Goddess!: Mara Strikes Back*.

Probability of buying saturates quickly

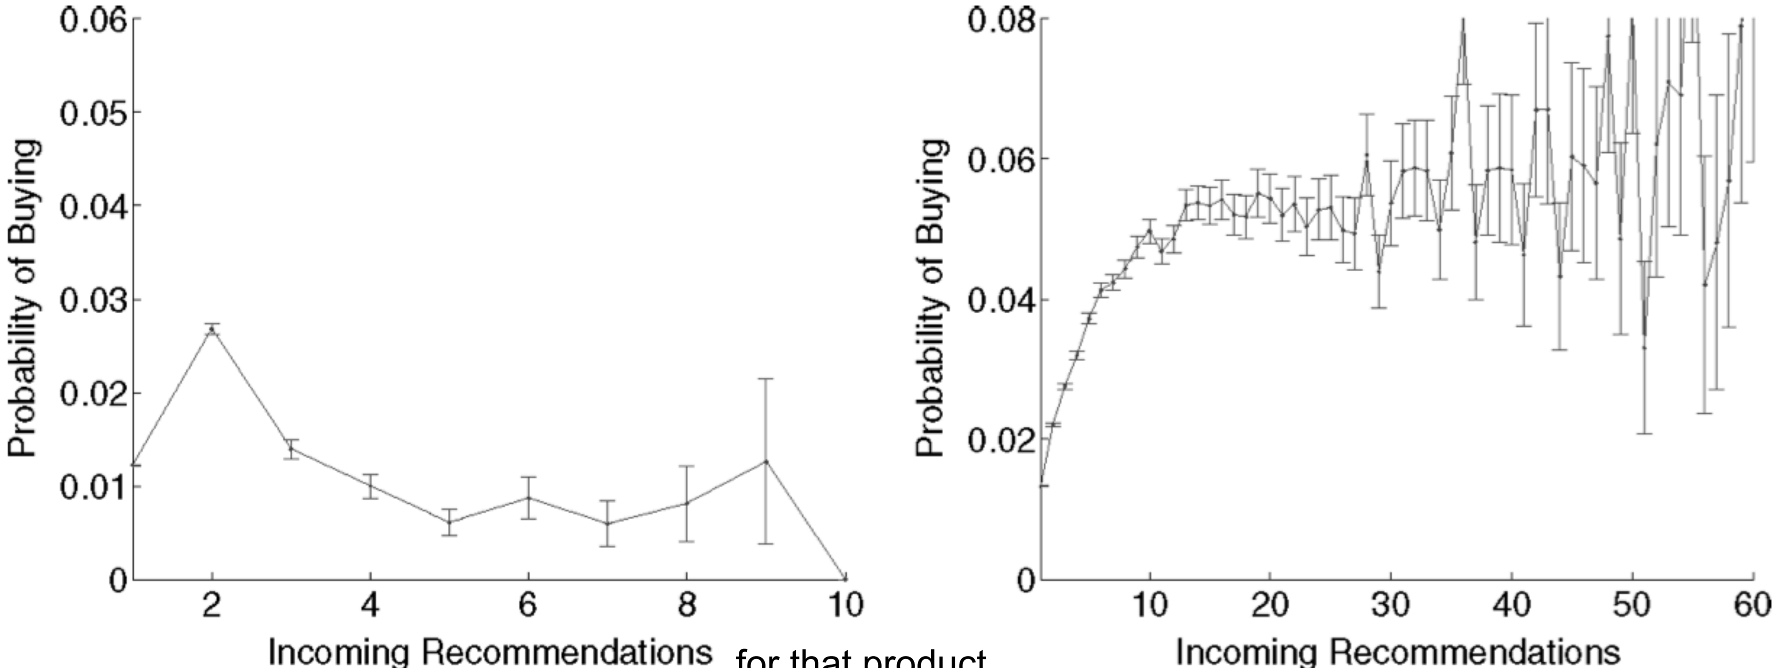

for that product

(a) Books

(b) DVD

Probability of buying saturates quickly

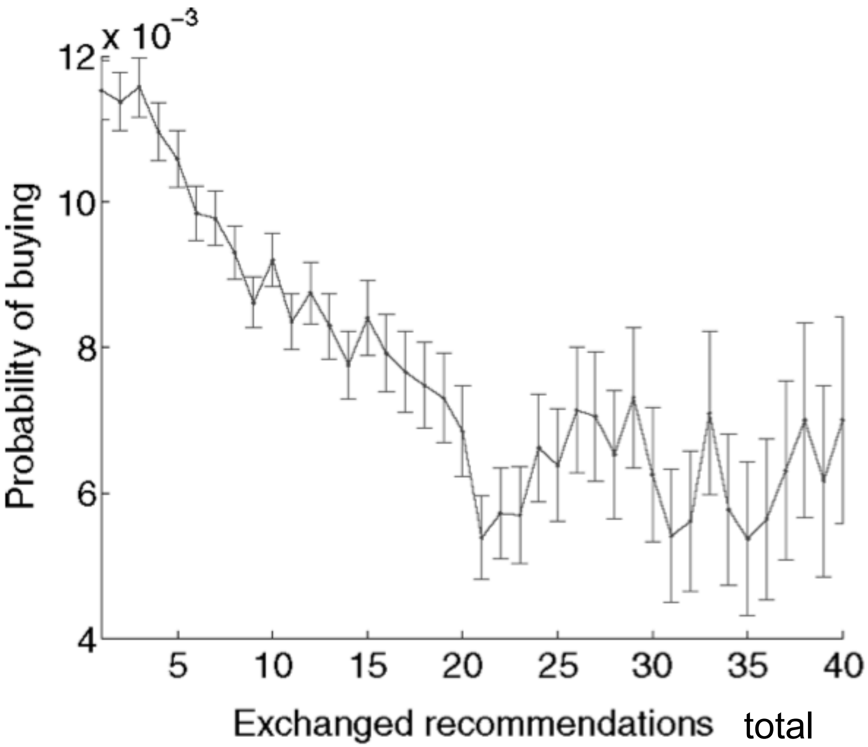

(a) Books

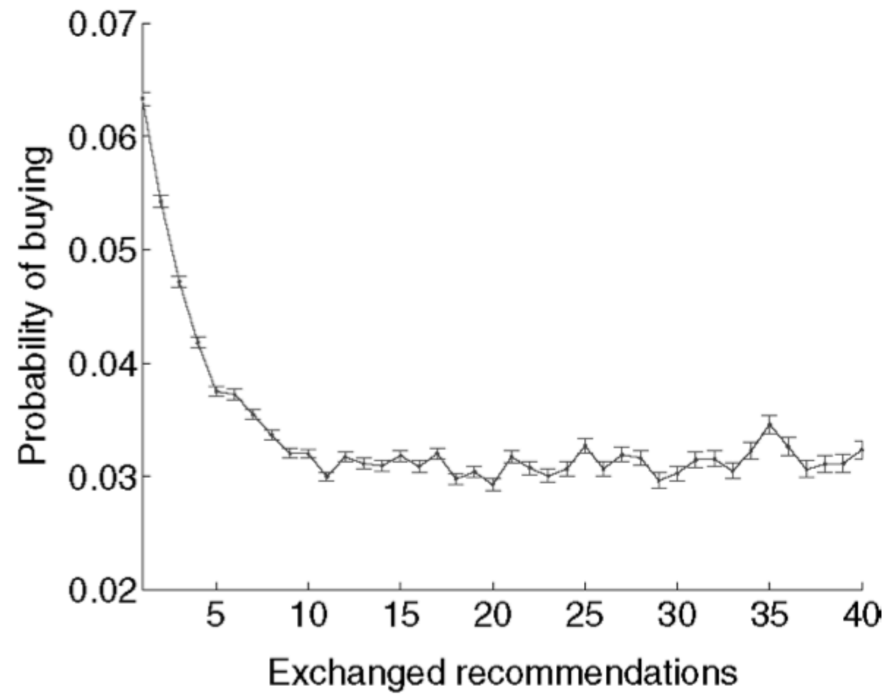

(b) DVD

Probability of buying saturates quickly

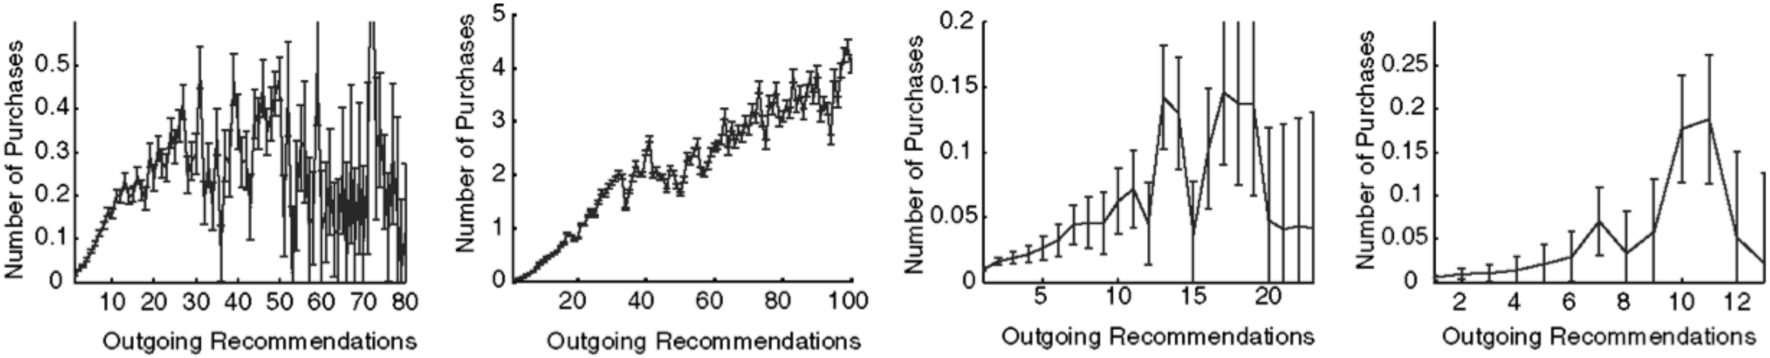

(a) Books (b) DVD (c) Music (d) Video

Some fraction of DVD purchases were from web sites where you *solicit* recommendations from past customers