Social Cogitation, Contagious Risk and Network formation

Éva Tardos
Cornell University

Plan for today

- Social Cogitation: a simple model
- Network formation: a simple model

Traditional application: speed of innovation or information

Cascading failure or risk?

- Network formation with cascading failure
- More sophisticated cascade models

Social contagion

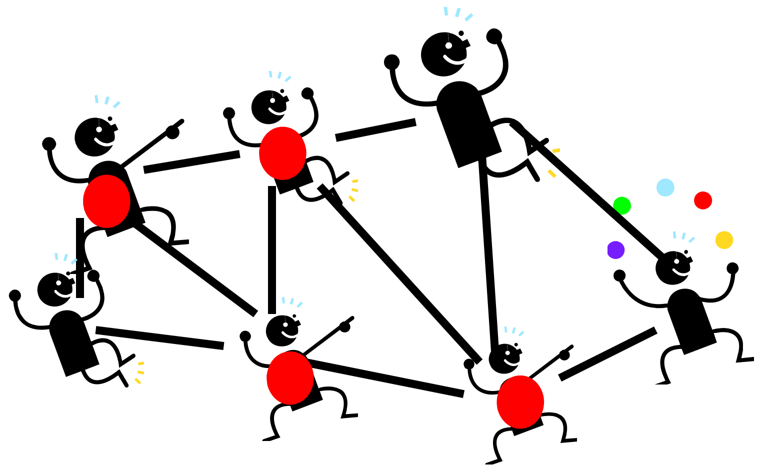

- Node v adopts the behavior,
- It may spread to neighbors

Models:

- spreads with probability p or p_e
- threshold model: adopt if θ_v neighbors adopted
- Adopt if influence $f_v(S) \geq \theta_v$ of adopted neighbors S.

Assumption $f_v(S)$ diminishing return

Domingos - Richardson KDD'01-02

Kempe – J. Kleinberg –Tardos KDD'03

Mossel-Roch STOC'07

Social contagion

Questions asked:

- How to start the largest possible cascade:
  - Starting set S
  - causes cascade $\sigma(S)$
  - Select A to maximize $\sigma(A)$

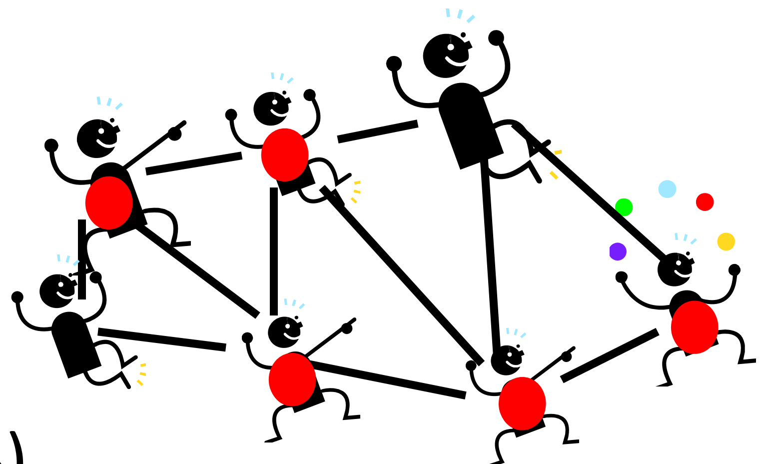

Mossel-Roch STOC'07: if $f_v(S)$ has diminishing return than so has $\sigma(S)$

Diminishing return: greedy algorithm optimizes cascade up to (1-1/*e*) factor

Cascade and Diminishing return

Example:

spreads with probability p_e

Proof idea: life edge subgraph

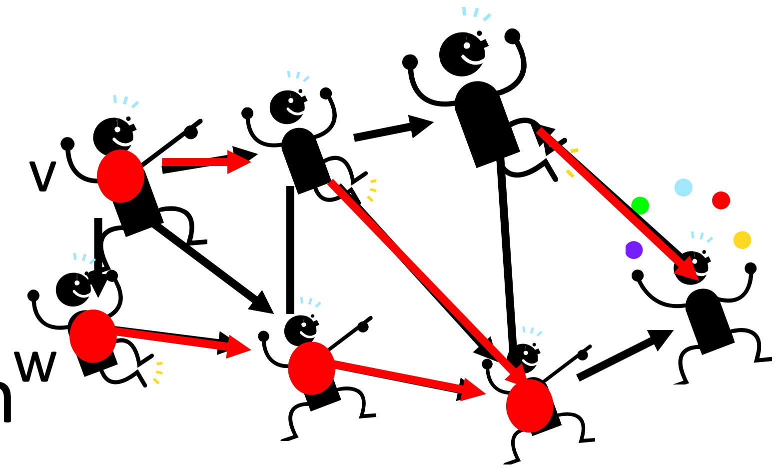

Need to prove:

- $\sigma(v) \leq \sigma(v+w) - \sigma(w)$
- Once w had its effect, some of $\sigma(v)$ already flipped.

Models of Network Formation

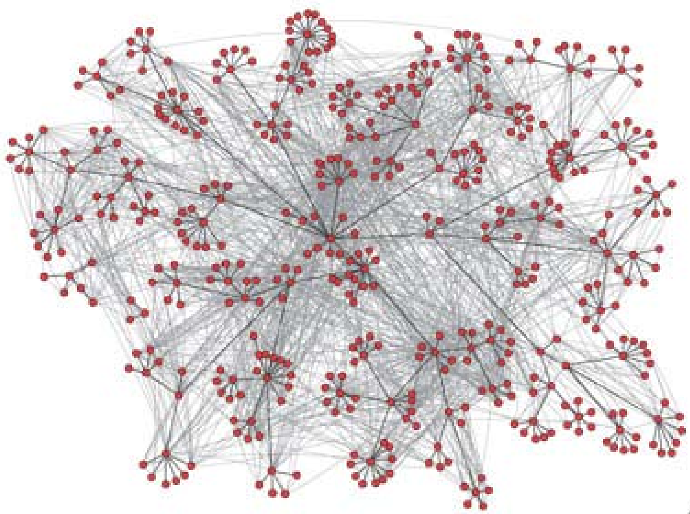

Corporate e-mail communication
(Adamic and Adar, 2005)

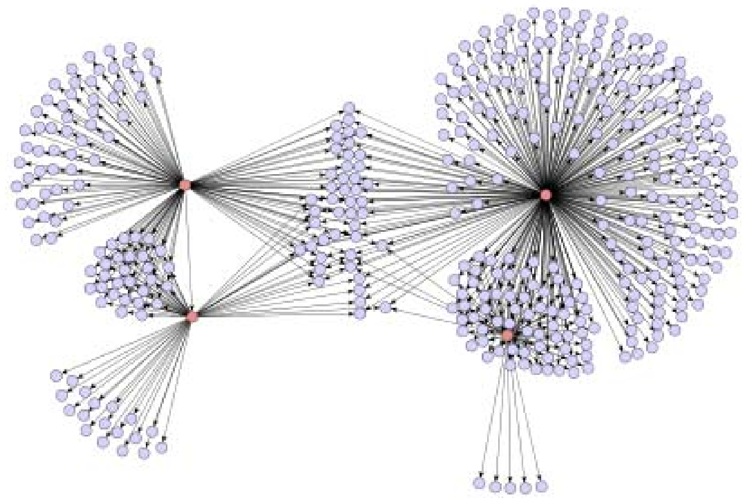

Book recommendations
(Leskovec-Adamic-Huberman 2006)

Why do complex networks look the way they do?

Two classes of models:

- Probabilistic
- Strategic

Models of Network Formation

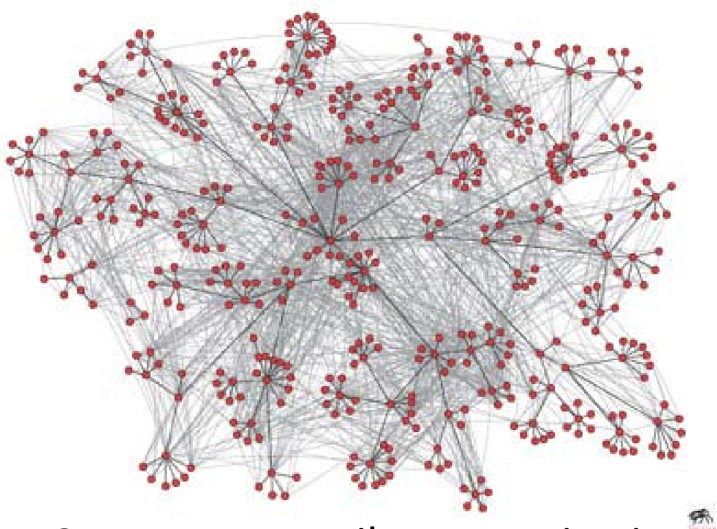

Corporate e-mail communication
(Adamic and Adar, 2005)

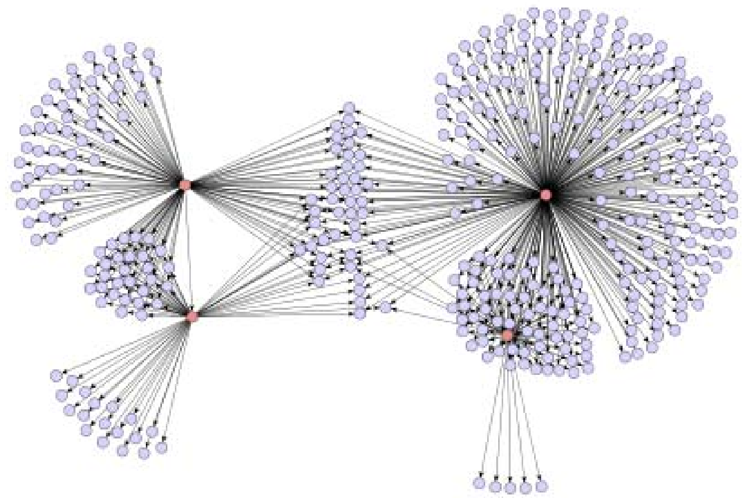

Book recommendations
(Leskovec-Adamic-Huberman 2006)

Probabilistic: Links form by simple probabilistic rules.

- Preferential attachment: Link to other nodes with prob. proportional to their degrees [Barabasi-Albert 1999].
- Small-world models: Link to other nodes with probability decaying in distance [Watts-Strogatz 1998, Kleinberg 2000].

Models of Network Formation

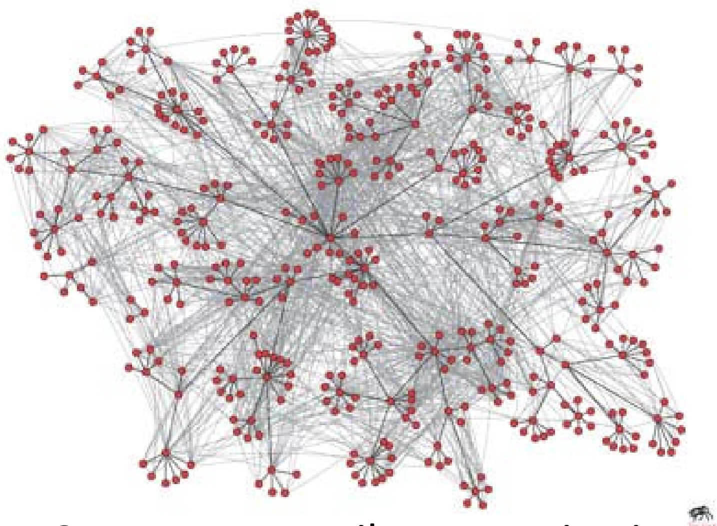

Corporate e-mail communication
(Adamic and Adar, 2005)

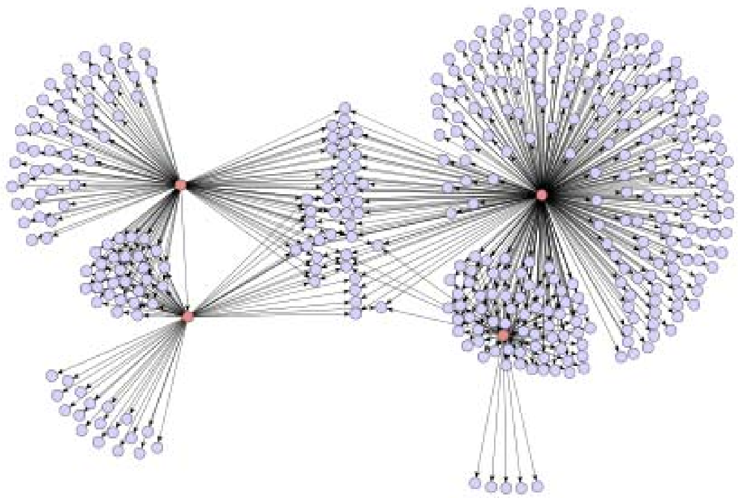

Book recommendations
(Leskovec-Adamic-Huberman 2006)

Strategic (e.g. Jackson-Wolinsky'96, Bala-Goyal'00, Fabrikant et al'03):

- Nodes can construct a subset of the possible links incident to them.
- They receive payoffs based on the structure of the overall network that forms.

Strategic Network Formation

Parameters $\alpha > 0$ and $0 < \delta < 1$.

Jackson-Wolinsky: node i 's payoff is

$$\Pi_i = -\alpha\, d_i + \Sigma_{j \neq i}\, \delta^{\,dist(I,j)}$$

where d_i is degree

Fabrikant at al:

$$\Pi_i = -\alpha\, d^+_i - \Sigma_{j \neq i}\, dist(i,j)$$

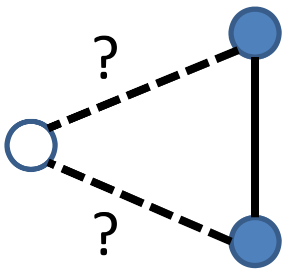

? $0, -\alpha+\delta+\delta^2, -2\alpha+2\delta$

A network is **stable** if no node can change its behavior to increase its payoff, given the choices of other nodes, e.g., Jackson-Wolinsky:

- Node i can't strictly increase its payoff by deleting all its links.
- If (i , j) isn't an edge, then adding (i , j) can't raise the payoffs of both i and j (at least one strictly).

What are stable networks and how good are they?

Fabrikant at al:

$$\Pi_i = -\alpha\, d^+_i - \Sigma_{j \neq i}\, \mathrm{dist}(i,j)$$

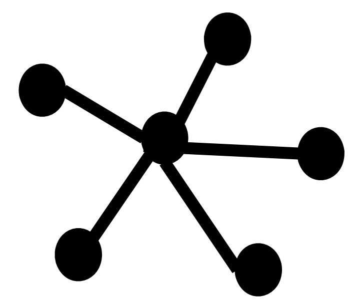

Stable network is connected

Social optimal: star $\sim \alpha n + 2n^2$

Equilibrium cost:

Distances at most $\sim 2\sqrt{\alpha}$

Why?

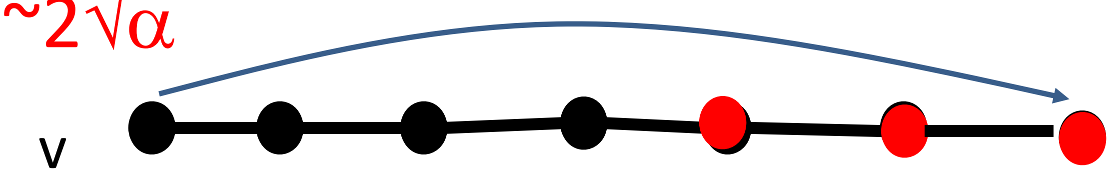

Adding edge decreases v's cost by to many node, with a total of more than α

Total distance cost at most $2n^2 \sqrt{\alpha}$

What are stable networks and how good are they?

Fabrikant at al: $\Pi_i = -\alpha\, d^+_i - \Sigma_{j \neq i}\, \mathrm{dist}(i,j)$

Has at most $\sim n^2/\sqrt{\alpha}$ edges

Few edges shortcut short paths

of less than $\sqrt{\alpha}$

Why?

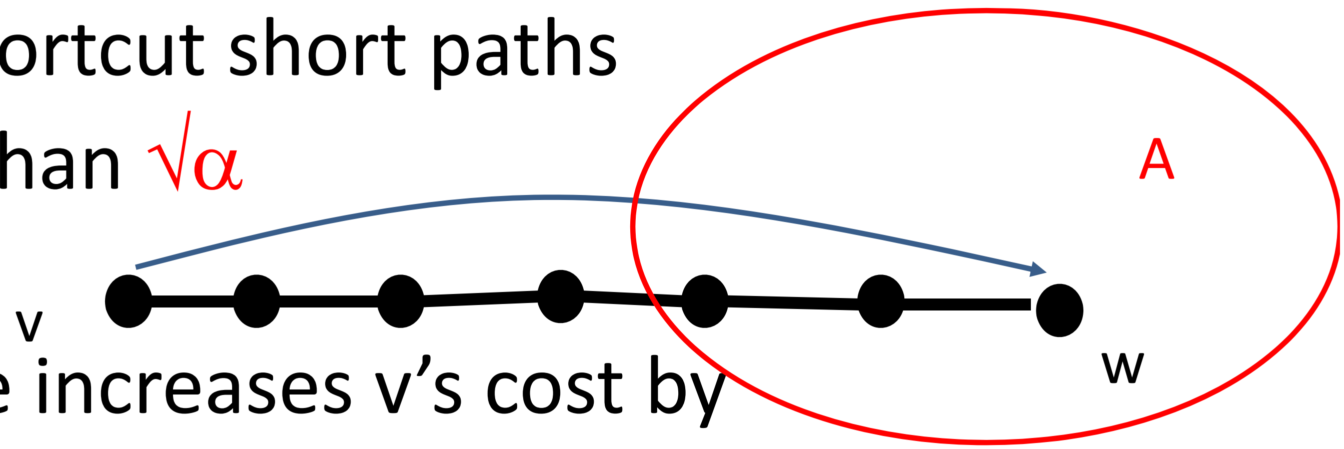

Deleting edge increases v's cost by

at most $|A|\sqrt{\alpha}$.

Beneficial if $|A| < \sqrt{\alpha}$. Hence at most $n/\sqrt{\alpha}$ such edges out of v.

What are stable networks and how good are they?

Fabrikant at al: $\Pi_i = -\alpha\, d^+_i - \Sigma_{j \neq i}$ dist(i,j)

Number of edges

- shortcut $\surd\alpha$ paths: at most n* (n/$\surd\alpha$)
- Don't shortcut path $n^2/\surd\alpha$

Has at most $\sim n^2/\surd\alpha$ edges

Total cost α*#edges+ Σ distances $\leq n^2\surd\alpha + n^2\surd\alpha$

Plan for today

- Social Cogitation: a simple model
- Network formation: a simple model

Traditional application: speed of innovation or information

Cascading failure or risk?

- Network formation with cascading failure
- More sophisticated cascade models

Contagion game: different payoff

- Standard payoff:
  - cost for link $\alpha\, d_i$
  - benefit from reaching others $\Sigma_{j \neq i}\, \delta^{\text{dist}(i,j)}$
  - Or cost of reaching others $\Sigma_{j \neq i}\, \text{dist}(i,j)$
- Contagion:
  - benefit from link
  - Danger from being connected to others

Blume, Easley, Kleinberg, Kleinberg, T. EC'11

Financial Networks

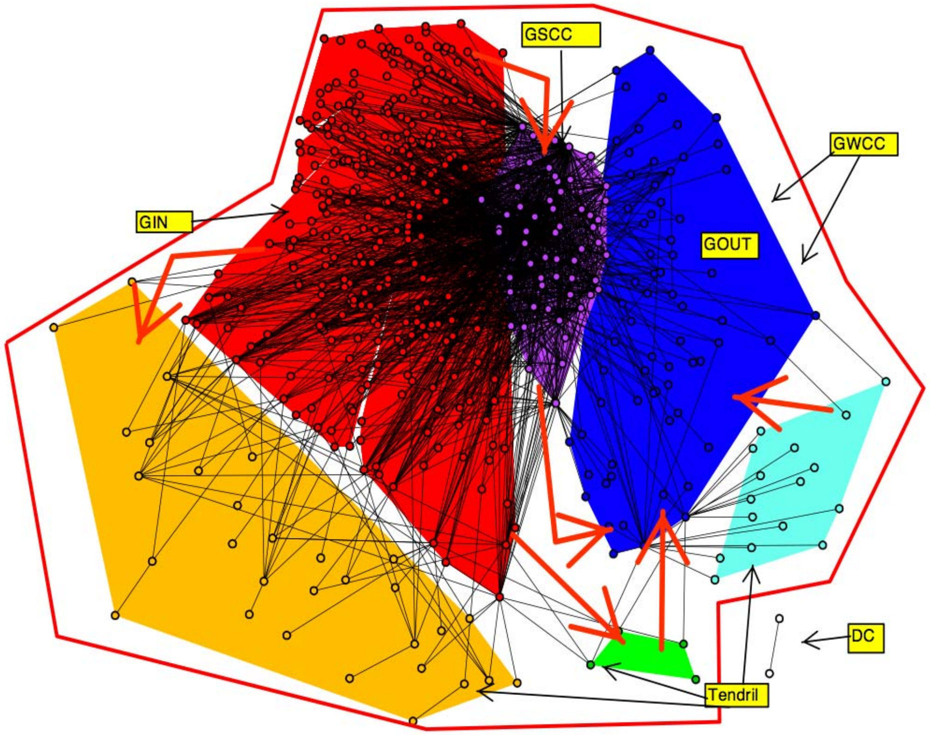

Network of overnight loans among banks (Bech and Atalay, 2008)

Financial networks [Allen-Gale 2000, Haldane-May 2011]

- Nodes benefit from transactions with others, but ...
- Counterparty risk: If X defaults on Y, this can hurt Y (and spread contagiously to other counterparties of Y).

Disease Epidemics

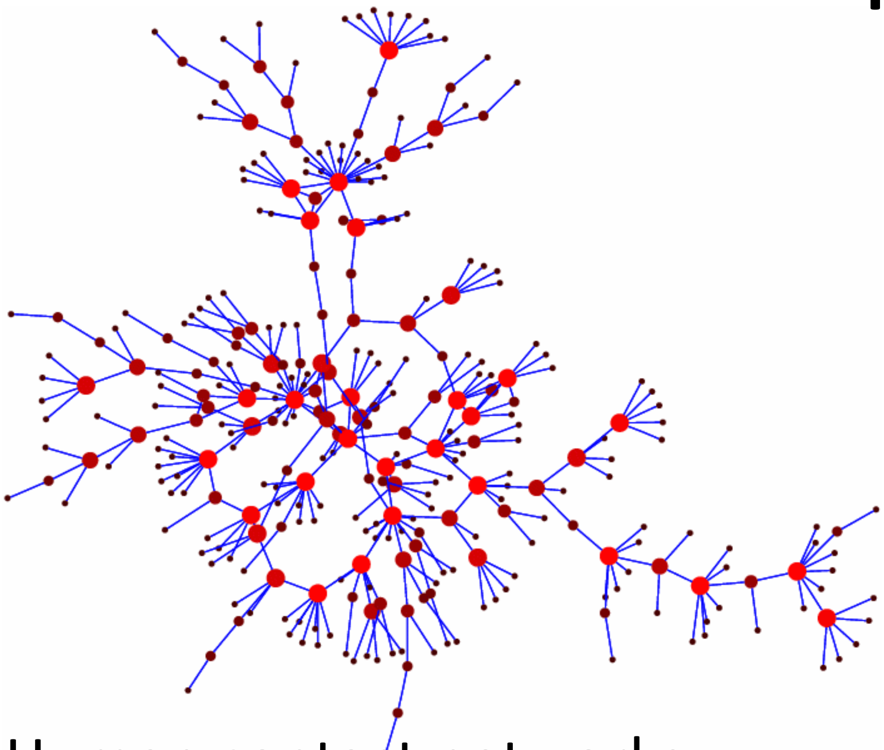

Network of sexual contacts (Potterat et al, 2002)

Human contact networks

- In the presence of an epidemic disease, people will alter their contact patterns to restrict more to in-group members.
- Evidence in case of HIV/AIDS for sexual contacts and needle-sharing [Jacquez et al 1988, Barnard 1993].

Covert Organizations

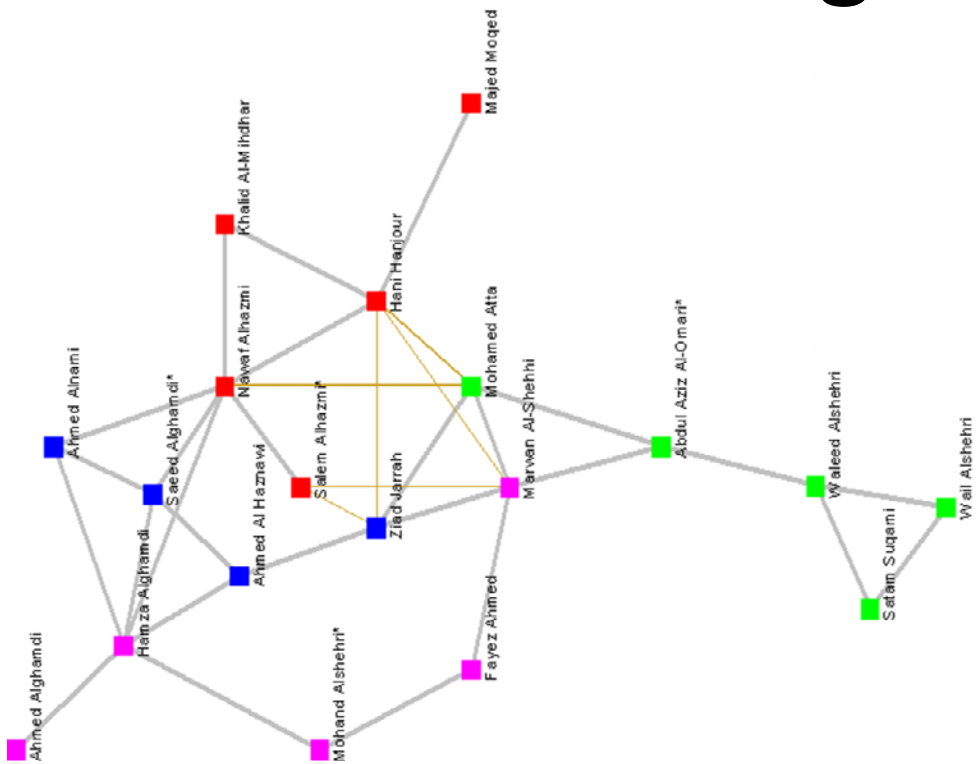

Terrorist network
(Krebs 2001)

Social networks in covert organizations.

- Links are useful for coordinations, but ...
- If one member is compromised, discovery of others can spread across links [Gutfraind 2010].

The Model

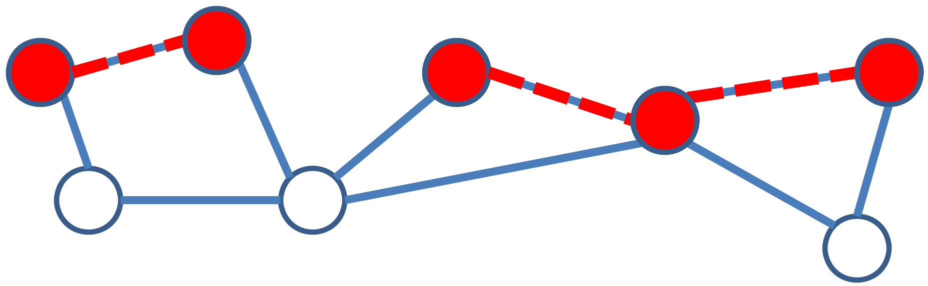

Link formation:

- Players V can form up to Δ bilateral relationships with others, obtaining a benefit of a > 0 from each.
- Resulting in undirected graph G = (V; E).

A failure process then spreads through the graph:

- Each node fails spontaneously with probability q > 0.
- Failure spreads along each edge with probability q > 0.
- Failed nodes lose link benefits and also incur cost of b < 0.

Payoffs

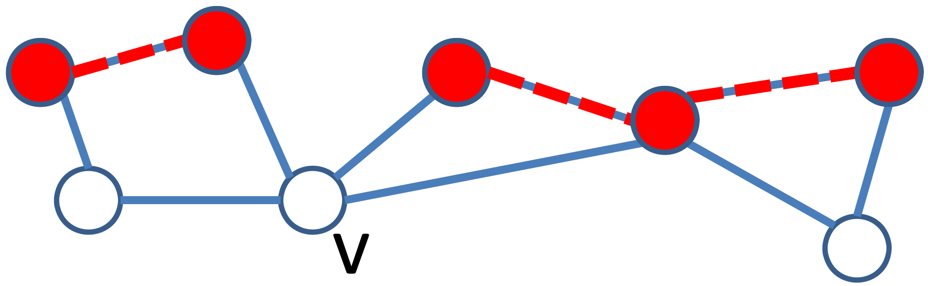

Let Φ_A denoting the probability node v fails, v 's payoff is

$$\Pi_v = a\, d_v\, (1 - \Phi_v) - b\, \Phi_v :$$

We compare graphs G according to min-welfare: $\min_v \Pi_v$

- Socially optimal G: maximizes min-welfare over all graphs.
- Also look at best/worst min-welfare of stable graphs.

  Stable = No node wants to delete all incident edges; no non-edge pair wants to link.

Assumptions

$a > bqp$.

$a < bp$.

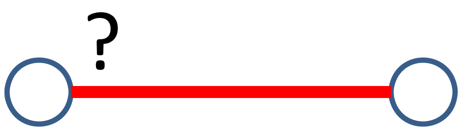

$a < bq$.

Assume that all of these bounds hold by a wide margin.

Condition $P(\delta)$:

$$\delta^{-1}\, bqp < a < \delta \min(bp; bq):$$

Main Results

The optimal min-welfare is $(1 + f(\delta))a/p$ for a non-negative fn $f(\delta) \rightarrow 0$.

- Socially optimal graphs are positioned just past a natural phase transition in the behavior of the payoffs.
- Exposes a difference between clustered and anonymous market structures.

For a sufficiently large number of nodes, the largest min-welfare of a stable graph is $g(\delta)\, a/p$ for a fn $g(\delta) \rightarrow 0$.

- Stable graphs involve slightly too much linking, which quickly burns away almost all the available payoffs.
- A type of tragedy of the commons, with survival probability serving as the shared resource that gets overconsumed.

Techniques: life-edge subgraph

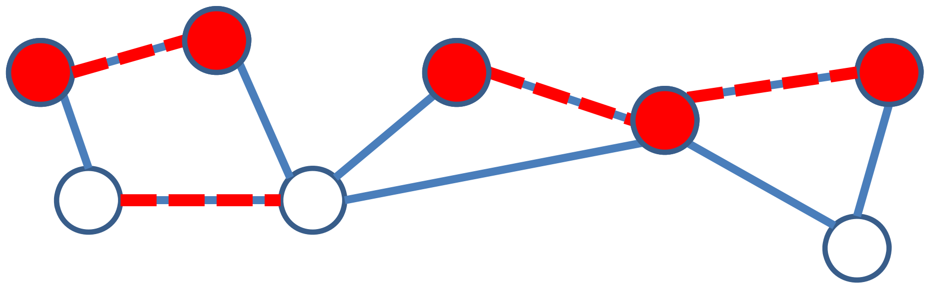

Nodes fail independently with probability $q > 0$.

Edges fail independently with probability $p > 0$.

- Picture the choice of failed edges being made first, before nodes fail.
- Results in a live-edge subgraph.
- A node fails, if any node in his life-edge component fails

Special case: click

G is a complete graph.

Live-edge subgraph is a sample from the random graph $G_{n,p}$.

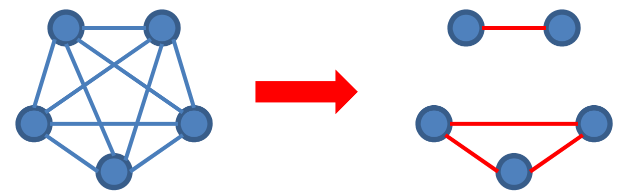

Theorem: If $p = (1-x)/n$ for fixed x, the probability a node v 's component exceeds size c is $\leq \exp(-c)$.

Theorem: If $p = (1+x)/n$ for fixed x, with high probability $\exists$ a component of size $\Theta(x)n$.

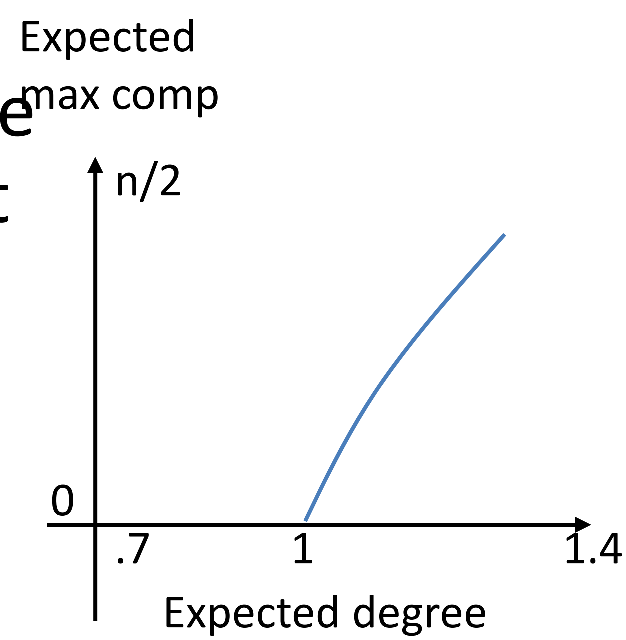

Life-edge components in general

Let G be a graph with minimum degree d, and declare each edge to be life with probability p, where pd = 1 + x for some fixed x > 0.

- Then for any node v , there is a constant prob. that v's live-edge component contains $>\varepsilon d$ nodes, for a constant ε.

A consequence (for fixed x > 0, with p; q; small):

- For all nodes to get payoff (1 + x)a/p, need degree (1 + x)/p

Assume p=q: component of size $\varepsilon(1 + x)/p$ means failure probability $\varepsilon(1 + x)$ constant, and cost $-\varepsilon(1 + x)$ b

- constant probability of being in a large live-edge component $\Rightarrow$high a probability of failure and hence a negative payoff.

Corollary: $\exists$ fn f (δ) $\rightarrow$ 0 such that no network can have min welfare >(1 + f (δ))a/p (critical payoff: a/p)

Proof idea

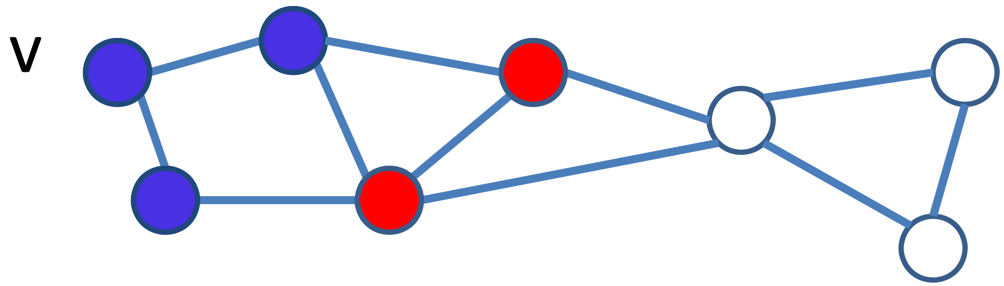

To show (high prob.) existence of a giant component [Karp 1990]:

Imagine performing breadth-first search to discover v's component.

- At the start, we remove one node from the BFS queue and add 1 + x nodes in expectation.
- So initially, the length of the BFS queue behaves like a random walk with positive drift

Super critical payoff possible?

Question: Is it possible to reach payoff $(1+f(\delta))a/p$, where $f(\delta)$ can go to 0 with δ? (assume p=q)

clustered market? union of disjoint cliques of size $(1 + f(\delta))/p$ each.

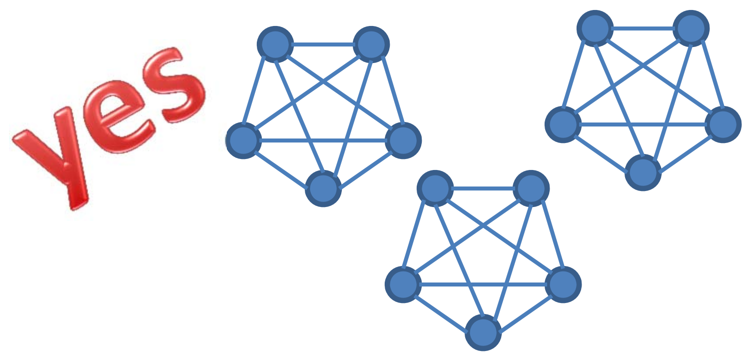

Anonymous market? random d-regular graph on n nodes, with $d = (1 + f(\delta))/p$.

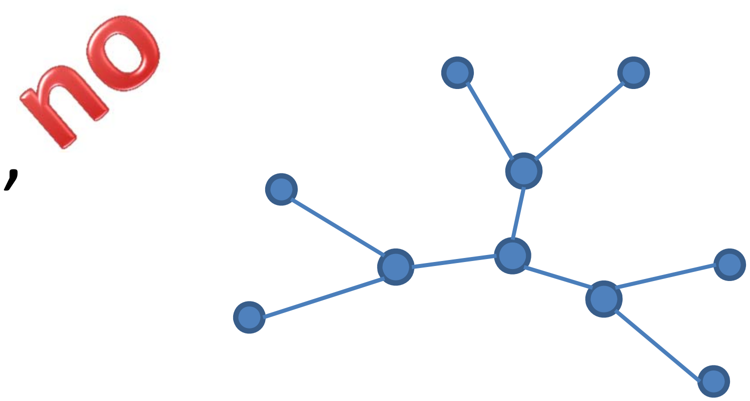

Proof idea: clustered vs anymous

Start with degree 1/p, and ask: is it a good idea to increase the degree to (1 + x)/p for some very small x?

Compare the payoff increase from links with payoff decrease from failure.

- v 's gain in payoff from the links is ax/p.
- v 's loss in payoff is b times the increased probability of failure.
- First-order approximation of payoff loss via differentiation.

Proof idea: anonymous

- v 's gain in payoff from the links is ax/p.
- v 's loss in payoff is b times the increased probability of failure.

Anonymous:

d-regular graph on n nodes, with d = (1 + x)

- Has one giant component with εn nodes, likely contains failed one
- Failure probability ~p*(component size/n) =pε
- Not worth increasing degree

Proof idea: clustered

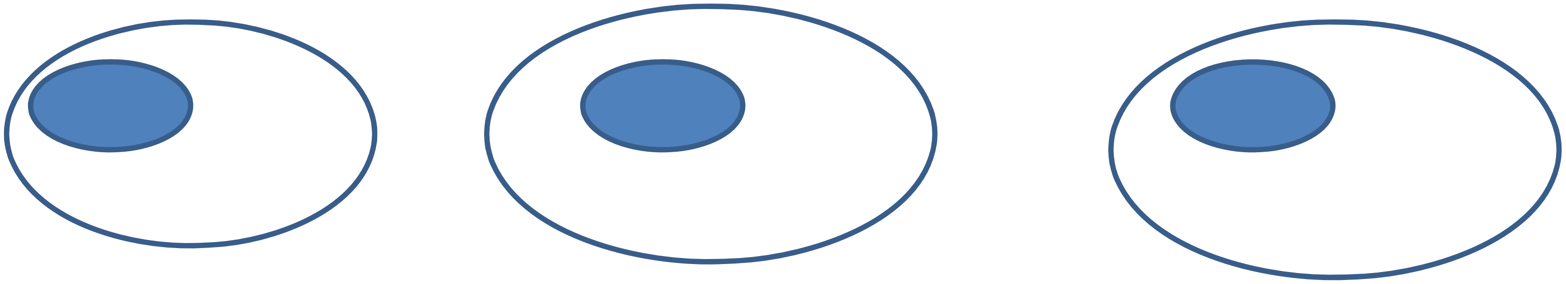

Clustered components of size d = (1 + x) each

- Has one giant component with εd nodes. Probability of containing failed node εdq
- Effect of increasing size: both v and failed node has to be in large component to cause failure: probability $(\varepsilon d)^2 q$
- worth increasing degree a bit.

Stable graphs

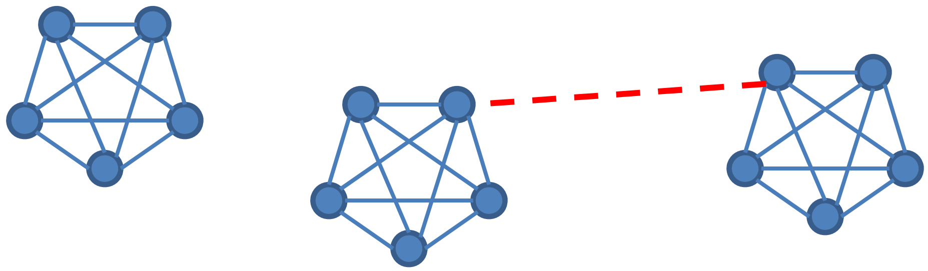

Why is the optimal union of cliques not stable?

- Two nodes v and w in different cliques will want to connect.
- v and w pass on increased risk to other nodes in their cliques,
- reducing payoffs of these other nodes: a negative externality.

Theorem: the union of slightly larger cliques is stable (with near-zero node payoffs).

Theorem: min-welfare of a stable graph of max degree Δ is $g(\delta)$ a/p for a function $g(\delta) \to 0$.

Proof idea: stable graphs

Theorem: min-welfare of a stable graph of max degree Δ is $g(\delta)\ a/p$ for a function $g(\delta) \to 0$.

failure probabilities Φ_v and Φ_w

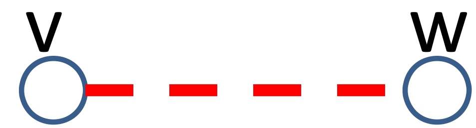

- w doesn't want edge (v,w):
- benefit a < loss $bp\Phi_v$ hence $a/bp < \Phi_v$

Payoff for such high risk nodes

$$\Pi_v = a\ d_v (1 - \Phi_v) - b\ \Phi_v < a\ d_v - a/p$$

hence needs degree $(1+g(\delta))$

Proof idea: stable graphs

Theorem: min-welfare of a stable graph of max degree Δ is $g(\delta)\ a/p$ for a fn $g(\delta) \rightarrow 0$.

Payoff for such high risk nodes

$\Pi_v = a\ d_v\ (1 - \Phi_v) - b\ \Phi_v < a\ d_v - a/p$

hence needs degree $(1+g(\delta))/p$

Low risk nodes

- Either have max degree Δ
- or form a click

Idea: nodes far away from this click have high total risk and hence low payoff

Plan for today

- Social Cogitation: a simple model
- Network formation: a simple model

Traditional application: speed of innovation or information

Cascading failure or risk?

- Network formation with cascading failure
- More sophisticated cascade models

More sophisticated cascade model

Cascade: adopt if influence $f_v(S)$ of adopted neighbors $S \geq \theta_v$.

Assumption $f_v(S)$ diminishing return

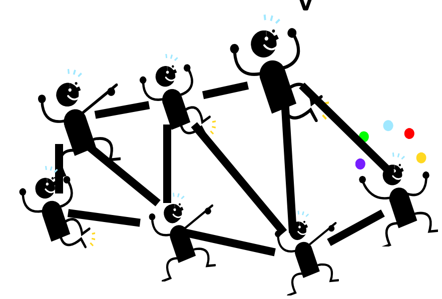

Is diminishing return reasonable (in risk)?

Probability of adoption if $f_v(k)$ as $k=|S|$ only

- $p_1 \geq p_2 \geq p_3 \geq p_4 \geq \ldots.$

~ true in technology adoption (maybe $p_1 < p_2$ and then true)

- In risk: $p_1 \leq p_2 \leq p_3 \leq p_4 \leq \ldots.$

Threshold model

adopt if θ_v neighbors adopted. If θ_v chosen uniformly random, cascade f(S) still diminishing return

But if θ_v deterministic, or not uniformly chosen, then not true

- Blume, Easley, Kleinberg, Kleinberg, T. FOCS'11

Threshold almost surely h: what are best networks

What can be best networks be?

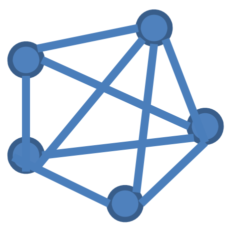

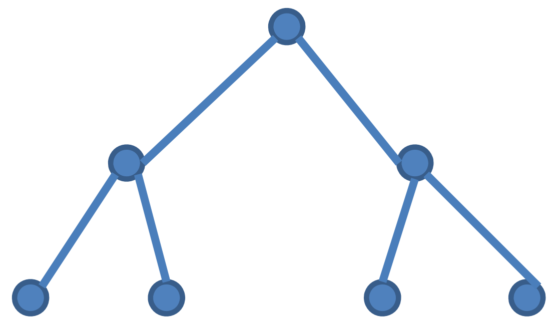

Model: d-regular graph (may be infinite).

- Each nodes draws a failure threshold μ from a common distribution. We will consider μ that high probability h, small probability 0, 1,…,h-1 (threshold distribution)

Intuition: limited size or diversity of neighbors?

What can be best networks be?

- Example 1: (s,1-s):

  Failed node kills the whole component

  click is uniquely optimal:

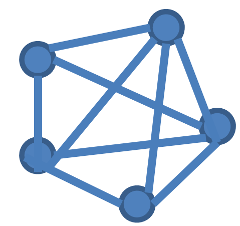

- Example 2: (s,0,1-s)

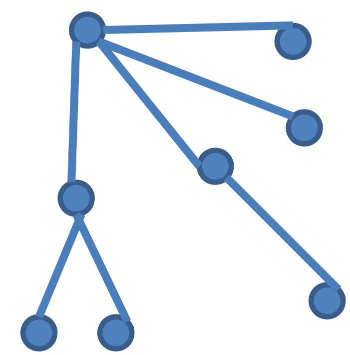

Click is still optimal!

Why? Think of randomness as assign

all nodes threshold 0,2 with prob. s, 0, 1-s.

Click node fails if node two neighbors draw 2

Results

- There is μ that clicks optimal
- Tree can beat click
- But there are others, than can beat click or tree

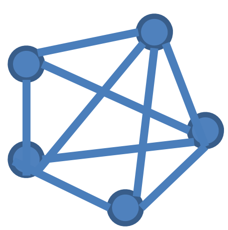

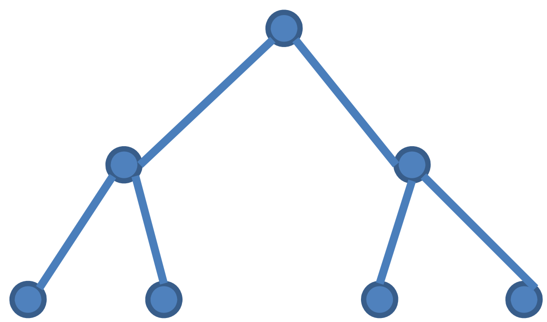

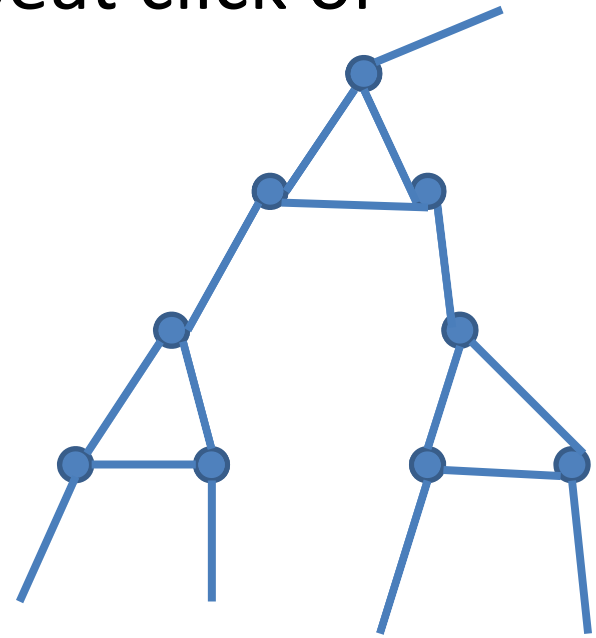

Proof idea

- Proof idea: tree can beat click? Consider $\mu=(s,t1-s-t)$

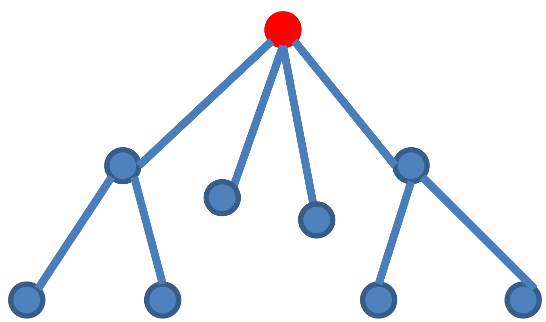

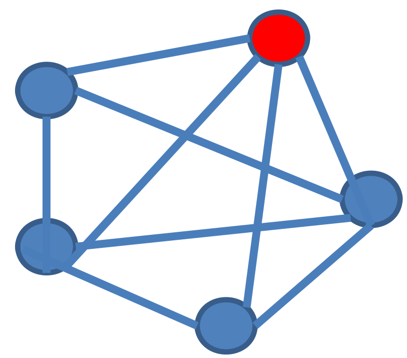

s+dst+d(d-1)/2 s^2+…

s+dst+d(d-1)/2 s^2+d(d-1)/2 s^t …

If s and t $\rightarrow$ 0 then these terms dominate

Summary

Network formation

- With diminishing return and benefit of adoption
- with the prospect of contagious failure.
- Stable graphs are risk-saturated, and this destroys most of the payoff.

Threshold model of adoption

Many open questions