Economics of Networks
Diffusion Part 1

Evan Sadler

Massachusetts Institute of Technology

Agenda

- Binary action coordination games
- Contagion in networks
- Mean-field diffusion models
- Connection to Bayesian games

Suggested Reading: EK Chapter 19; Morris (2000), "Contagion;" Jackson and Yariv (2007), "Diffusion of Behavior and Equilibrium Properties in Network Games"

A Binary Action Coordination Game

Consider the payoff matrix

	0	1
0	(q, q)	$(0, 0)$
1	$(0, 0)$	$(1-q, 1-q)$

for some $q \in (0, 1)$

Two actions, players earn the same payoff, positive iff they match

- Action 1 is optimal iff other player chooses 1 w.p $\geq q$

Coordinating on what?

- Meeting place, technology standard, behavioral norm

Local Interaction Systems

We're going to think about player the coordination game with many opponents, not just one

- Simultaneously play the game with a set of neighbors
- Earn a payoff from the interaction with each neighbor
- Must play same action with all neighbors

Action 1 is optimal iff at least a fraction q of my neighbors choose action 1

Definition

A local interaction system is an infinite population $\mathcal{X}$ in which each agent interacts with a finite subset of others. Write $x \sim y$ if x and y are neighbors. Assume $x \sim y \implies y \sim x$, there exists $M < \infty$ such that $|\{y : x \sim y\}| \leq M$ for all x, and there exists a path between any pair of players.

Local Interaction Systems

Write G_x for the set of neighbors of x

Write X for some set of players and $\overline{X}$ for its complement

Define $\pi(X, x)$ as the proportion of x's neighbors in X:

$$\pi(X, x) = \frac{|X \cap G_x|}{|G_x|}$$

Define $\Pi^p(X)$ as the set of players x for whom $\pi(X, x) \geq p$

If set Y chooses action 1, the best response is for players in $X = \Pi^q(Y)$ to choose action 1

Best Response Dynamics

We will study best response dynamics in local interaction systems

Imagine action 1 spreading through the population

Question: Is there a finite group of players that, if they start out playing action 1, can spread action 1 to the entire population?

The contagion threshold ξ is the largest value of q for which this is possible

- Higher threshold $\implies$ easier to get contagion

Example: Interaction on a Line

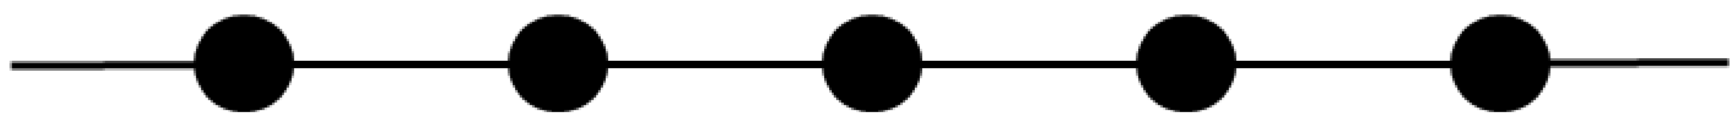

If $q < \frac{1}{2}$, a player switches to action 1 after one neighbor does

Contagion threshold $\frac{1}{2}$

Example: Interaction on an m-Dimensional Lattice

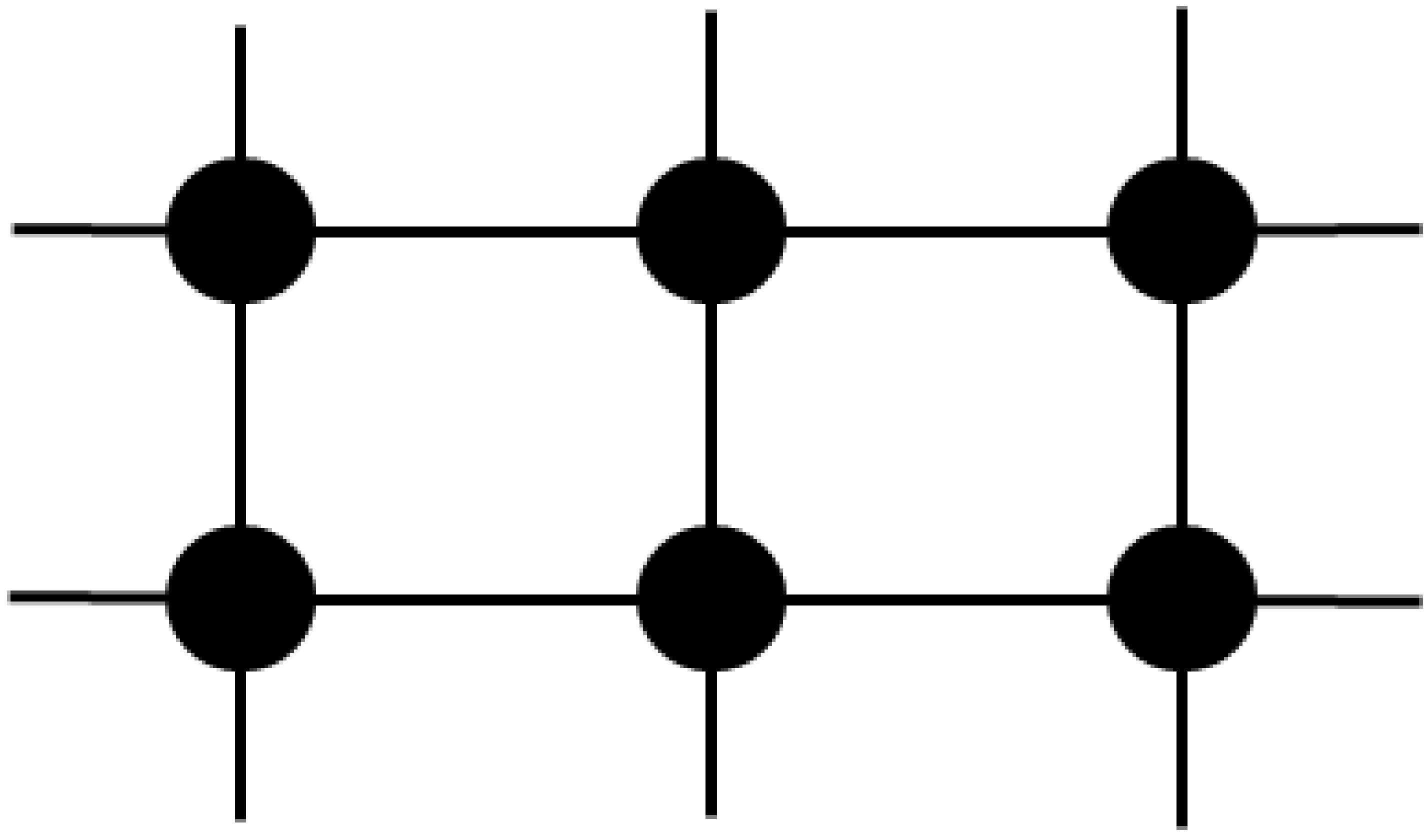

Contagion threshold $\frac{1}{4}$

Example: Trees

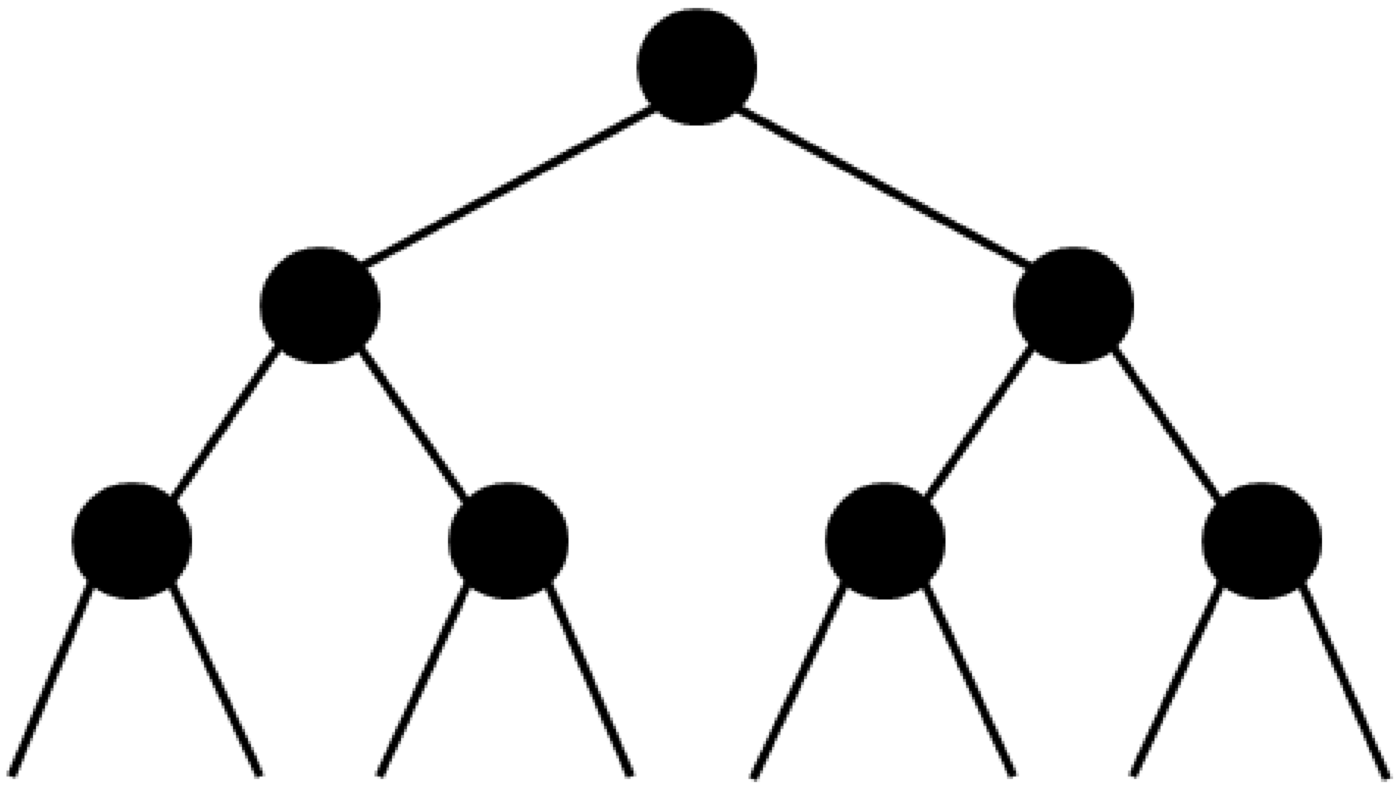

Contagion threshold $\frac{1}{3}$

Group Cohesion

Use several properties in characterization. First: group cohesion

A natural measure: relative frequency of in-group links versus out-group links

Define the cohesion of X as

$$c(X) = \min_{x \in X} \pi(X, x) = \max\{p : X \subseteq \Pi^p(X)\}$$

We say X is p-cohesive if $c(X) \geq p$

Cohesion in our Examples

Line network: non-trivial group is at most $\frac{1}{2}$-cohesive

2-dimensional lattice: consider set of players above some horizonal line

- Group is $\frac{3}{4}$-cohesive

Tree: group consisting of an entire branch is $\frac{2}{3}$ cohesive

Neighbor Growth

Distance between x and X: length of shortest path between x and *some* player in X

Define $G^n(X)$ as set of players within distance n of X

A local interaction system satisfies low neighbor growth if for all $\gamma > 1$, we have

$$\frac{G^n(X)}{\gamma^n} \to 0$$

i.e. $G^n(X)$ grows subexponentially

Intuitively, low neighbor growth correlates with higher cohesion

δ-Uniformity

A labeling l on the set of players is a bijection with the positive integers

Write $\alpha_l(i)$ for the fraction of neighbors of i with a lower label under l

A labeling is an Erdös labeling if there is a finite group X such that for all n we have

$$i \in G^n(X) \text{ and } j \notin G^n(X) \implies l(i) < l(j)$$

A local interaction system satisfies δ-uniformity if there exists an Erdös labeling such that for all sufficiently large K:

$$\max_{i,j \geq K} |\alpha_l(i) - \alpha_l(j)| \leq \delta$$

δ-uniformity

Intuitive idea: for players sufficiently far from the reference set, fraction of neighbors with lower labels (i.e. closer to reference set) tends roughly towards something (within δ)

Example: 2-dimensional lattice

- δ-uniformity fails for $\delta < \frac{1}{4}$

For any n, there are $4(n+1)$ players in $G^{n+1}(\{x\})$ but not in $G^n(\{x\})$

- Four sides of a square

Players in the middle of a side have $\alpha_l(k) = \frac{1}{2}$, but corners have $\alpha_l(k) = \frac{1}{4}$.

Characterizing the Contagion Threshold

Recall the contagion threshold

$$\xi = \max\left\{q : \bigcup_{k\geq 1}[\Pi^q]^k(X) = \mathcal{X} \text{ for some finite } X\right\}$$

A group X is co-finite if its complement $\overline{X}$ is finite

Theorem

The contagion threshold ξ is the smallest p such that every co-finite group contains an infinite $(1-p)$-cohesive subgroup.

Cohesive groups can act as barriers to contagion

Characterizing the Contagion Threshold

Theorem

The contagion threshold is always at most $\frac{1}{2}$.

If the system satisfies low neighbor growth and δ-uniformity, then the contagion threshold satisfies

$$\xi \geq \frac{1}{2} - \delta$$

Intuition: contagion always spreads slowly; if not, eventually too few neighbors are choosing action 1

We need new adopters to have enough interaction with one another

Coexistence of Conventions

An equilibrium of the local interaction game is a best response to itself

- A group X taking action 1 is an equilibrium if everyone in X wants to keep choosing 1 and everyone in $\overline{X}$ wants to keep choosing 0

A group X is an equilibrium iff X is q-cohesive and $\overline{X}$ is $(1-q)$-cohesive

Theorem

Suppose the system satisfies low neighbor growth and has contagion threshold ξ. For all $q \in [\xi, 1-\xi]$, the game has a co-existent equilibrium.

Discussion

Results highlight qualitative features of networks that facilitate contagion

Could derive analogous results for large but finite networks

Conditions not always easy to check, doesn't answer all questions we might ask

Next up: a different approach based on distributional information about the network

A Mean-Field Approximation

We have a large population (a continuum), characterize network via degree distribution D

- Write p_d for probability of degree d
- Write $\overline{d} = \sum_d p_d d$ for the average degree

A random neighbor has degree d with probability

$$\tilde{p}_d = \frac{p_d d}{\overline{d}}$$

Friendship paradox

Agents choose between two actions, 0 or 1

- Refer to 0 as the "default"
- Heterogeneous costs c of choosing 1, continuous distribution F

A Mean-Field Approximation

Added value of choosing action 1 is $v(d, x)$

- Player's degree d
- Expects each neighbor to choose 1 w.p. x
- Assume increasing in x, complementarities

Agent i prefers action 1 if $v(d_i, x) \geq c_i$

- Happens with probability $F(v(d_i, x))$

Define $F_{d,x} = F(v(d, x))$, probability that degree d player wants to adopt, given neighbor adoption probability x

Examples

Suppose $v(d,x) = u(dx)$ for some increasing concave u

- Payoffs depend on expected number of adopting neighbors
- Might be reasonable for adoption of a communication technology

Suppose $v(d,x) = u(x)$ for some increasing concave u

- Care about average play of neighbors, network has no role

Suppose $v(d,x)$ is a step function:

$$v(d,x) = \begin{cases} a & \text{if } x \leq \tau \\ b & \text{if } x > \tau \end{cases}$$

for some $\tau \in (0,1)$

Bayesian Equilibrium

Consider a static game first

Look at symmetric Bayes-Nash Equilibria:

- Agent i observes own degree d_i and cost c_i
- Neighbors are random draw from the population, degree d with probability $\tilde{p}_d$

Existence follows from standard arguments

Equilibrium condition:

$$x = \phi(x) \equiv \sum_d \tilde{p}_d F_{d,x}$$

Bayesian Equilibrium

Can think of $\phi(x)$ as a best response map

- If agents play a best response to neighbor adoption probability x, then the new neighbor adoption probability is $\phi(x)$

In equilibrium, agent i adopts iff $c_i \leq v(d_i, x)$

Neighbor adoption probability x fully characterizes equilibrium behavior

- Refer to x as an equilibrium

A Diffusion Process

Now look at a diffusion process in discrete time

At $t = 0$, some fraction of the population exogenously adopts action 1

- Write x^0 for corresponding neighbor adoption probability

Write x_d^t for fraction of degree d agents adopting at time t

- Neighbor adoption probability $x^t = \sum_d \tilde{p}_d x_d^t$

Next period, assume

$$x_d^{t+1} = F_{d,x^t}$$

A Diffusion Process

Best response dynamics

Two key assumptions:

- Agents are myopic
- Next period, neighobrs are a new independent random draw from the population

Second assumption is what makes this “mean-field”

- Unlike our earlier model, identity of adopting individual doesn’t matter
- Can just keep track of population averages

Equilibrium Structure

Complementarities ensure monotone convergence of the dynamics

Note, if v is weakly increasing in degree d, then in equilibrium, higher degree agents are more likely to adopt

A point x is a fixed point of the dynamics iff

$$x = \sum_d \tilde{p}_d F_{d,x} = \phi(x)$$

Fixed points are exactly equilibria of the static game

Stability and Tipping

An equilibrium x is stable if there exists $\epsilon' > 0$ such that $\phi(x - \epsilon) > x - \epsilon$ and $\phi(x + \epsilon) < x + \epsilon$ for all $\epsilon' > \epsilon > 0$

Best response dynamics will return to equilibrium after a small disturbance

An equilibrium is a tipping point if the inequalities above are reversed

- A small disturbance causes the system to move away from the equilibrium

Let ϕ and $\hat{\phi}$ denote two best response mappings. We say $\hat{\phi}$ generates greater diffusion if:

- For any stable equilibrium in ϕ, there exists a higher one in $\hat{\phi}$
- For any tipping point in ϕ, there exists a lower one in $\hat{\phi}$

Stability and Tipping

If $\hat{\phi}$ is pointwise larger than ϕ, then $\hat{\phi}$ generates greater diffusion

Example:

- Suppose F is uniform on $[0, 1]$
- Assume $v(d, x) = \frac{1}{2}\sqrt{x}$ for all d
- Implies $F_{d,x} = \phi(x) = \frac{1}{2}\sqrt{x}$

Equilibrum at $x = 0$ is unstable, equilibrium at $x = \frac{1}{4}$ is stable

Comparative Statics

Changes in the cost distribution

Theorem

If $\hat{F}$ FOSD F, then we have $\hat{\phi}(x) \leq \phi(x)$ for all x, implying F generates greater diffusion than $\hat{F}$.

If adoption is more costly, you get less of it

Comparative Statics

What about changes in the network structure?

Theorem

Consider two different neighbor degree distributions p and $\hat{p}$, and suppose $F_{d,x}$ is non-decreasing in d. If p FOSD $\hat{p}$, then $\phi(x) \geq \hat{\phi}(x)$, and p generates greater diffusion than $\hat{p}$.

If high-degree agents are (weakly) more inclined to adopt, then increasing density increases diffusion

$$\phi(x) = \sum_d \frac{p_d d}{\sum_k p_k k} F_{d,x} \geq \sum_d \frac{\hat{p}_d d}{\sum_k p_k k} F_{d,x} = \hat{\phi}(x)$$

by definition of FOSD

Comparative Statics

Something less obvious...

Theorem

Suppose p is a mean-preserving spread of $\hat{p}$, and $dF_{d,x}$ is non-decreasing and weakly convex in d. Then $\phi(x) \geq \hat{\phi}(x)$, and p generates greater diffusion than $\hat{p}$.

Can think of MPS as increasing "centralization"

- Intuition: MPS exacerbates the friendship paradox

Proof left as an exercise

Looking Ahead

Mean-field approach has advantages and disadvantages

- Tractability
- Some problematic assumptions

Next time, an approach based on random graphs

- Distributional assumptions on the network
- No reshuffling links