A simple theory of cascades in networks

Center for Institutional Studies Birthday Workshop
Higher School of Economics

30 May 2015

Yongwhan Lim, Asuman Ozdaglar, Alex Teytelboym

MIT, MIT, Oxford

My microeconomics tutor


Russian Summer
Analysis


Russian Summer

Motivation

- Consumers care what products their friends and relatives use.
- Examples: innovation/technology adoption, social platform use, mobile phone contracts.
- Switching costs are often high: product adoption is **irreversible** (at least temporarily).
- Firms' **initial seeds** in the social network really matter for profit and market share.

Motivation

- Consumers care what products their friends and relatives use.
- Examples: innovation/technology adoption, social platform use, mobile phone contracts.
- Switching costs are often high: product adoption is **irreversible** (at least temporarily).
- Firms' **initial seeds** in the social network really matter for profit and market share.

Motivation

- Consumers care what products their friends and relatives use.
- Examples: innovation/technology adoption, social platform use, mobile phone contracts.
- Switching costs are often high: product adoption is **irreversible** (at least temporarily).
- Firms' **initial seeds** in the social network really matter for profit and market share.

Motivation

- Consumers care what products their friends and relatives use.
- Examples: innovation/technology adoption, social platform use, mobile phone contracts.
- Switching costs are often high: product adoption is **irreversible** (at least temporarily).
- Firms' **initial seeds** in the social network really matter for profit and market share.

Previous work

- This work is mostly closely related to: Goyal and Kearns (2012); Bimpikis, Ozdaglar, and Yildiz (2014) (...and Hotelling, 1929)
- Quality and seeding: Fazeli and Jadbabaie (2012a,b,c); Fazeli, Ajorlou, and Jadbabaie (2014).
- Other papers where consumers can switch products many times: Bharathi, Kempe, and Salek (2007); Alon, Feldman, Procaccia, and Tennenholtz (2010); Apt and Markakis (2011); Simon and Apt (2012); Tzoumas, Amanatidis, and Markakis (2012); Borodin, Braverman, Lucier, and Oren (2013); Apt and Markakis (2014); Mei and Bullo (2014).

Outline of this talk

- We develop a tractable model of cascades in networks.
- We introduce a measure of node influence called **cascade centrality**.
- We study a competitive diffusion game on the network.
- We also characterize the expected number of adopters using cascade centrality in general graphs and find analytical expressions for many graphs.
- In a follow-up paper, we tackle network design questions: maximizing adoption and minimizing failures.

Outline of this talk

- We develop a tractable model of cascades in networks.
- We introduce a measure of node influence called **cascade centrality**.
- We study a competitive diffusion game on the network.
- We also characterize the expected number of adopters using cascade centrality in general graphs and find analytical expressions for many graphs.
- In a follow-up paper, we tackle network design questions: maximizing adoption and minimizing failures.

Outline of this talk

- We develop a tractable model of cascades in networks.
- We introduce a measure of node influence called **cascade centrality**.
- We study a competitive diffusion game on the network.
- We also characterize the expected number of adopters using cascade centrality in general graphs and find analytical expressions for many graphs.
- In a follow-up paper, we tackle network design questions: maximizing adoption and minimizing failures.

Outline of this talk

- We develop a tractable model of cascades in networks.
- We introduce a measure of node influence called **cascade centrality**.
- We study a competitive diffusion game on the network.
- We also characterize the expected number of adopters using cascade centrality in general graphs and find analytical expressions for many graphs.
- In a follow-up paper, we tackle network design questions: maximizing adoption and minimizing failures.

Outline of this talk

- We develop a tractable model of cascades in networks.
- We introduce a measure of node influence called **cascade centrality**.
- We study a competitive diffusion game on the network.
- We also characterize the expected number of adopters using cascade centrality in general graphs and find analytical expressions for many graphs.
- In a follow-up paper, we tackle network design questions: maximizing adoption and minimizing failures.

Outline of this talk

- We develop a tractable model of cascades in networks.
- We introduce a measure of node influence called **cascade centrality**.
- We study a competitive diffusion game on the network.
- We also characterize the expected number of adopters using cascade centrality in general graphs and find analytical expressions for many graphs.
- In a follow-up paper, we tackle network design questions: maximizing adoption and minimizing failures.

Model

Preliminaries

- Simple, undirected graph $G(V, E)$.
- A *adoption threshold* for agent i is a random variable Θ_i drawn from a probability distribution with support $[0, 1]$.
- The associated multivariate probability distribution for all the agents in the graph is $f(\boldsymbol{\theta})$.
- Each agent is $i \in V$ assigned a threshold θ_i. Let's define the threshold profile of agents as $\boldsymbol{\theta} := (\theta_i)_{i \in V}$. A **network** G_θ is a graph endowed with a threshold profile.

Preliminaries

- Simple, undirected graph $G(V, E)$.
- A *adoption threshold* for agent i is a random variable Θ_i drawn from a probability distribution with support $[0, 1]$.
- The associated multivariate probability distribution for all the agents in the graph is $f(\theta)$.
- Each agent is $i \in V$ assigned a threshold θ_i. Let's define the threshold profile of agents as $\boldsymbol{\theta} := (\theta_i)_{i \in V}$. A **network** G_θ is a graph endowed with a threshold profile.

Preliminaries

- Simple, undirected graph $G(V, E)$.
- A *adoption threshold* for agent i is a random variable Θ_i drawn from a probability distribution with support $[0, 1]$.
- The associated multivariate probability distribution for all the agents in the graph is $f(\boldsymbol{\theta})$.
- Each agent is $i \in V$ assigned a threshold θ_i. Let's define the threshold profile of agents as $\boldsymbol{\theta} := (\theta_i)_{i \in V}$. A **network** G_θ is a graph endowed with a threshold profile.

Preliminaries

- Simple, undirected graph $G(V, E)$.
- A *adoption threshold* for agent i is a random variable Θ_i drawn from a probability distribution with support $[0, 1]$.
- The associated multivariate probability distribution for all the agents in the graph is $f(\boldsymbol{\theta})$.
- Each agent is $i \in V$ assigned a threshold θ_i. Let's define the threshold profile of agents as $\boldsymbol{\theta} := (\theta_i)_{i \in V}$. A **network** G_θ is a graph endowed with a threshold profile.

Seeding

- Two firms: A selling product a and B selling product b. Products are perfectly substitutable.
- The state of agent i at time t is denoted $x_i(t) \in \{0, a, b\}$.
- Denote by $S_t^A(G_\theta)$ and $S_t^B(G_\theta)$ the sets of **new** adopters of products A and B in network G_θ at time t resp.
- At time $t = 0$, $x_i(0) = 0$ for all i, and each firm simultaneously chooses **one** agent $S_0^A, S_0^B \in V$ as a seed for their product. Overlap in seed sets resolved randomly.

Seeding

- Two firms: A selling product a and B selling product b. Products are perfectly substitutable.
- The state of agent i at time t is denoted $x_i(t) \in \{0, a, b\}$.
- Denote by $S_t^A(G_\theta)$ and $S_t^B(G_\theta)$ the sets of **new** adopters of products A and B in network G_θ at time t resp.
- At time $t = 0$, $x_i(0) = 0$ for all i, and each firm simultaneously chooses **one** agent $S_0^A, S_0^B \in V$ as a seed for their product. Overlap in seed sets resolved randomly.

Seeding

- Two firms: A selling product a and B selling product b. Products are perfectly substitutable.
- The state of agent i at time t is denoted $x_i(t) \in \{0, a, b\}$.
- Denote by $S_t^A(G_\theta)$ and $S_t^B(G_\theta)$ the sets of **new** adopters of products A and B in network G_θ at time t resp.
- At time $t = 0$, $x_i(0) = 0$ for all i, and each firm simultaneously chooses **one** agent $S_0^A, S_0^B \in V$ as a seed for their product. Overlap in seed sets resolved randomly.

Seeding

- Two firms: A selling product a and B selling product b. Products are perfectly substitutable.
- The state of agent i at time t is denoted $x_i(t) \in \{0, a, b\}$.
- Denote by $S_t^A(G_\theta)$ and $S_t^B(G_\theta)$ the sets of **new** adopters of products A and B in network G_θ at time t resp.
- At time $t = 0$, $x_i(0) = 0$ for all i, and each firm simultaneously chooses **one** agent $S_0^A, S_0^B \in V$ as a seed for their product. Overlap in seed sets resolved randomly.

Linear threshold process dynamics

- Any agent who has not adopted any product by some period t, **decides to adopt one of the products** in time period $t+1$ iff

$$\frac{\text{total friends who adopted } a + \text{ total friends who adopted } b}{\text{total friends}} \geq \theta_i$$

  i.e. Granovetter's linear threshold model.

- If the threshold is reached, the **probability of adopting product** a is

$$\frac{\#\text{ friends who adopted } a \text{at } t}{\#\text{ friends who adopted } a \text{at } t + \#\text{ friends who adopted } b \text{at } t}$$

- Once an agent adopted product a, he remains in state a in all subsequent periods.
- This process converges to a random set: eventual adopters S^A of product a and S^B of product b.

Linear threshold process dynamics

- Any agent who has not adopted any product by some period t, **decides to adopt one of the products** in time period $t+1$ iff

$$\frac{\text{total friends who adopted } a + \text{ total friends who adopted } b}{\text{total friends}} \geq \theta_i$$

i.e. Granovetter's linear threshold model.

- If the threshold is reached, the **probability of adopting product** a is

$$\frac{\#\text{ friends who adopted } a \text{at } t}{\#\text{ friends who adopted } a \text{at } t + \#\text{ friends who adopted } b \text{at } t}$$

- Once an agent adopted product a, he remains in state a in all subsequent periods.
- This process converges to a random set: eventual adopters S^A of product a and S^B of product b.

Linear threshold process dynamics

- Any agent who has not adopted any product by some period t, **decides to adopt one of the products** in time period $t+1$ iff

$$\frac{\text{total friends who adopted } a + \text{ total friends who adopted } b}{\text{total friends}} \geq \theta_i$$

i.e. Granovetter's linear threshold model.

- If the threshold is reached, the **probability of adopting product** a is

$$\frac{\# \text{ friends who adopted } a \text{at } t}{\# \text{ friends who adopted } a \text{at } t + \# \text{ friends who adopted } b \text{at } t}$$

- Once an agent adopted product a, he remains in state a in all subsequent periods.
- This process converges to a random set: eventual adopters S^A of product a and S^B of product b.

Linear threshold process dynamics

- Any agent who has not adopted any product by some period t, **decides to adopt one of the products** in time period $t+1$ iff

$$\frac{\text{total friends who adopted } a + \text{ total friends who adopted } b}{\text{total friends}} \geq \theta_i$$

  i.e. Granovetter's linear threshold model.

- If the threshold is reached, the **probability of adopting product** a is

$$\frac{\#\text{ friends who adopted } a \text{at } t}{\#\text{ friends who adopted } a \text{at } t + \#\text{ friends who adopted } b \text{at } t}$$

- Once an agent adopted product a, he remains in state a in all subsequent periods.
- This process converges to a random set: eventual adopters S^A of product a and S^B of product b.

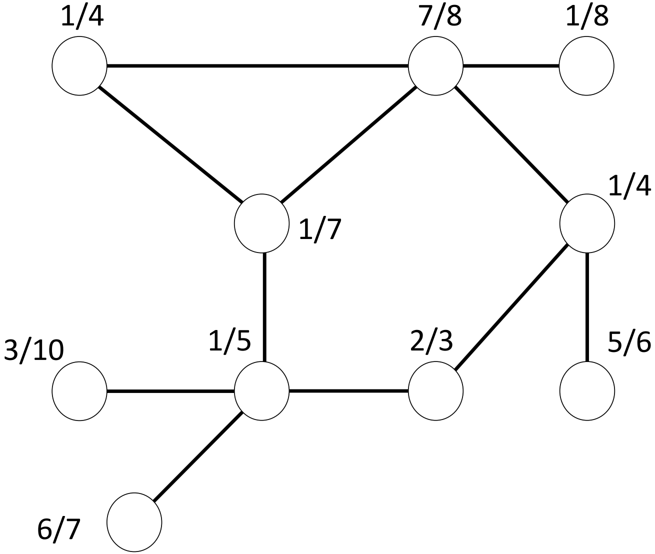
1/4
7/8
1/8
1/4
1/7
3/10
1/5
2/3
5/6
6/7

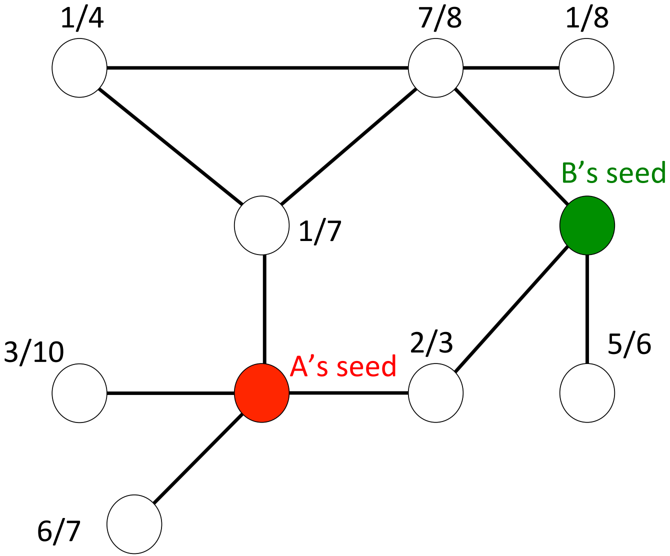
1/4
7/8
1/8
B’s seed
1/7
3/10
A’s seed
2/3
5/6
6/7

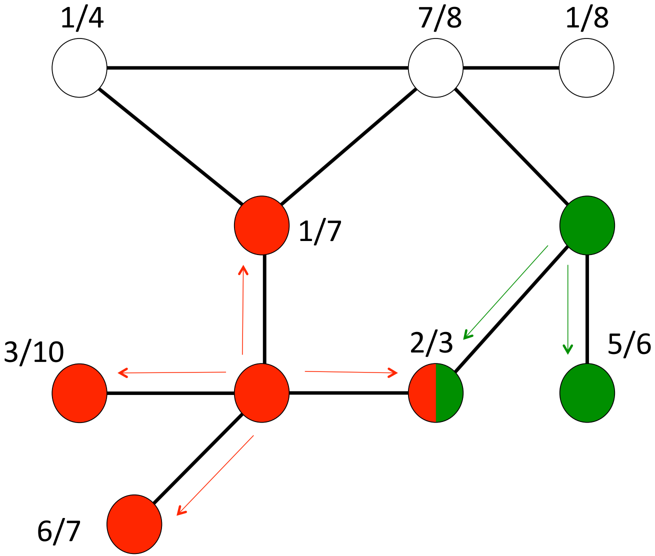
1/4
7/8
1/8
1/7
3/10
2/3
5/6
6/7

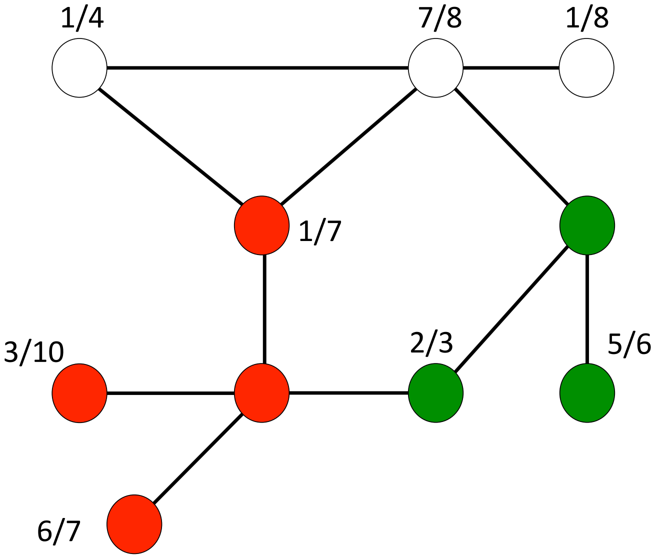
1/4
7/8
1/8
1/7
3/10
2/3
5/6
6/7

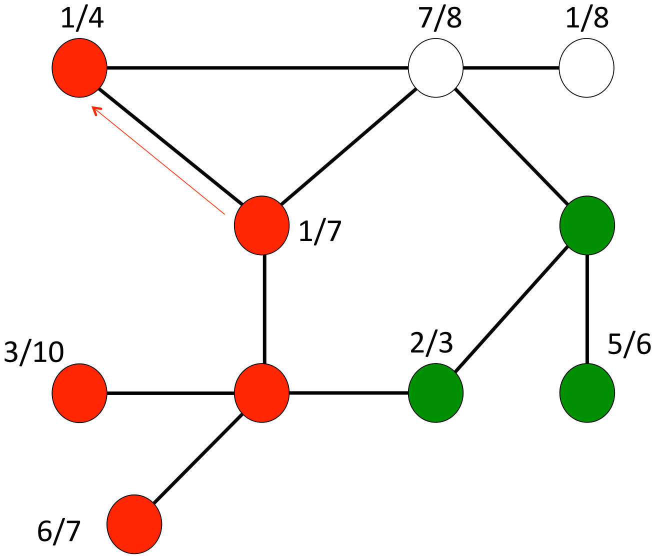
1/4
7/8
1/8
1/7
3/10
2/3
5/6
6/7

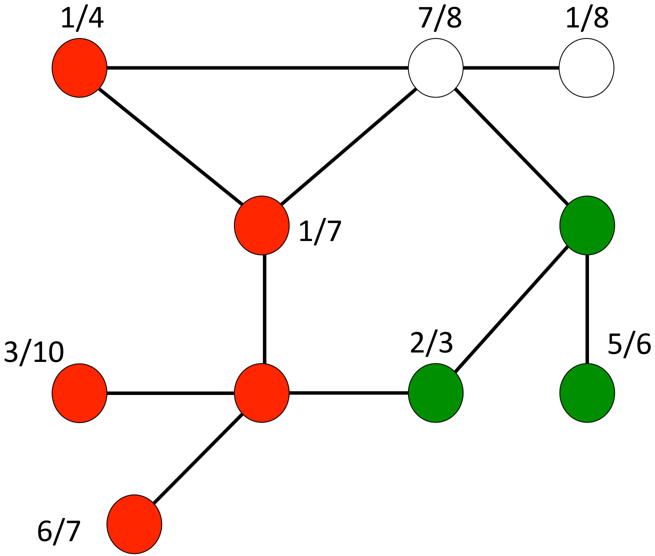
1/4
7/8
1/8
1/7
3/10
2/3
5/6
6/7

Expected number of adopters

- Fixing seeds S_0^A and S_0^B and a graph G, and re-run the process by drawing the agents' thresholds from $f(\theta)$ each time.
- Denoting the probability of any agent adopting product a is

$$\mathbb{P}_i^A(G, S_0^A, S_0^B) = \int_{\mathbb{R}^n} |S^A(G_\theta, S_0^A, S_0^B) \cap \{i\}| f(\theta) d\theta$$

- Expected number of adopters of product a is

$$\mathbb{E}[S^A(G, S_0^A, S_0^B)] = \int_{\mathbb{R}^n} |S^A(G_\theta, S_0^A, S_0^B)| f(\theta) d\theta = \sum_{i=1}^{n} \mathbb{P}_i^A(G, S_0^A, S_0^B)$$

Expected number of adopters

- Fixing seeds S_0^A and S_0^B and a graph G, and re-run the process by drawing the agents' thresholds from $f(\boldsymbol{\theta})$ each time.
- Denoting the probability of any agent adopting product a is

$$\mathbb{P}_i^A(G, S_0^A, S_0^B) = \int_{\mathbb{R}^n} |S^A(G_\theta, S_0^A, S_0^B) \cap \{i\}| f(\boldsymbol{\theta}) d\boldsymbol{\theta}$$

- Expected number of adopters of product a is

$$\mathbb{E}[S^A(G, S_0^A, S_0^B)] = \int_{\mathbb{R}^n} |S^A(G_\theta, S_0^A, S_0^B)| f(\boldsymbol{\theta}) d\boldsymbol{\theta}$$
$$= \sum_{i=1}^{n} \mathbb{P}_i^A(G, S_0^A, S_0^B)$$

Expected number of adopters

- Fixing seeds S_0^A and S_0^B and a graph G, and re-run the process by drawing the agents' thresholds from $f(\boldsymbol{\theta})$ each time.
- Denoting the probability of any agent adopting product a is

$$\mathbb{P}_i^A(G, S_0^A, S_0^B) = \int_{\mathbb{R}^n} |S^A(G_\theta, S_0^A, S_0^B) \cap \{i\}| f(\boldsymbol{\theta}) d\boldsymbol{\theta}$$

- Expected number of adopters of product a is

$$\begin{aligned} \mathbb{E}[S^A(G, S_0^A, S_0^B)] &= \int_{\mathbb{R}^n} |S^A(G_\theta, S_0^A, S_0^B)| f(\boldsymbol{\theta}) d\boldsymbol{\theta} \\ &= \sum_{i=1}^{n} \mathbb{P}_i^A(G, S_0^A, S_0^B) \end{aligned}$$

Consider what happens when firm A is a monopolist

Uniform distribution

Assumption

For any G_θ and every $i \in V$, $\Theta_i \sim \mathcal{U}(0,1)$ and independent.

It's the Laplacian prior for the firms. Moreover, we prove that

$$\mathbb{P}_i^A(G) = \sum_{j \in N_i(G)} \frac{\mathbb{P}_j^A(G | i \notin S^A)}{d_i}$$

if and only if $\Theta_i \sim \mathcal{U}(0,1)$.

Paths

Definition

A sequence of nodes $P = (i_0, \cdots, i_k)$ on a graph G is a path if $i_j \in N_{i_{j-1}}(G)$ for all $1 \leq j \leq k$ and each $i_j \in P$ is distinct.

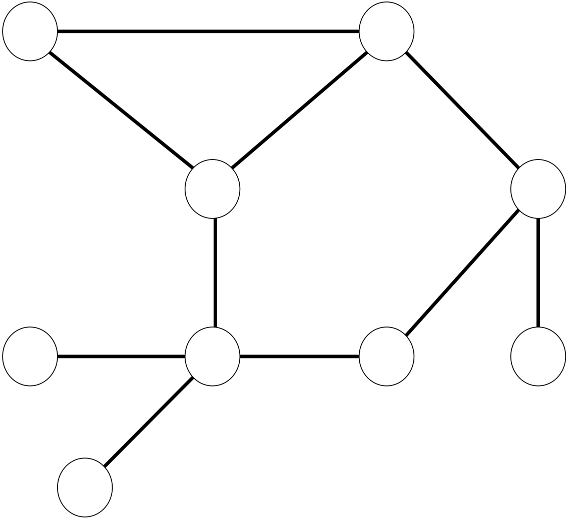

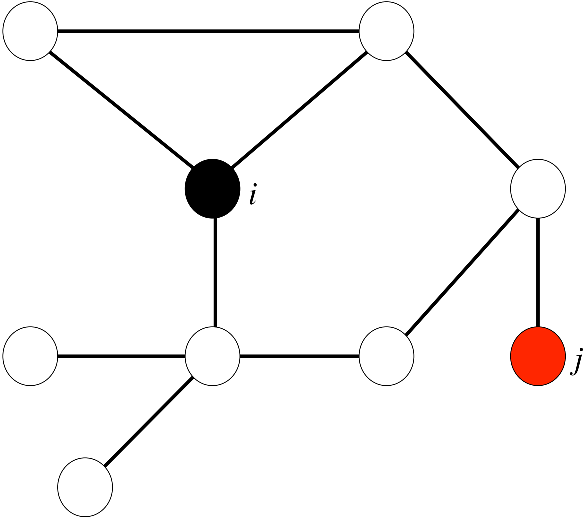

i
j

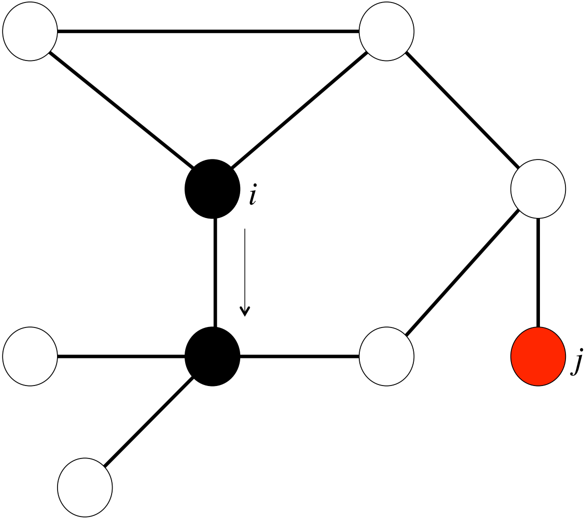
i
j

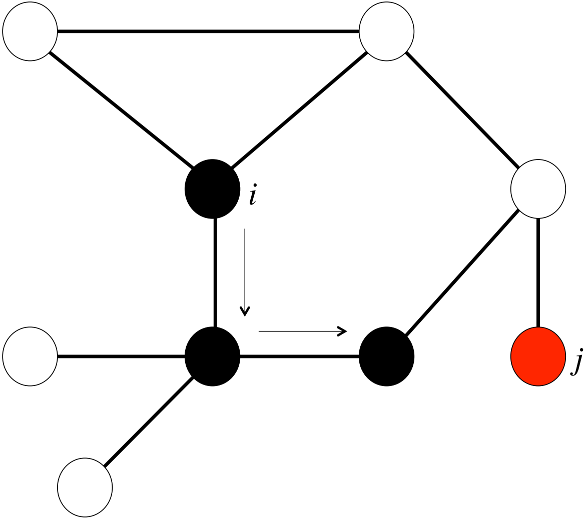

i
j

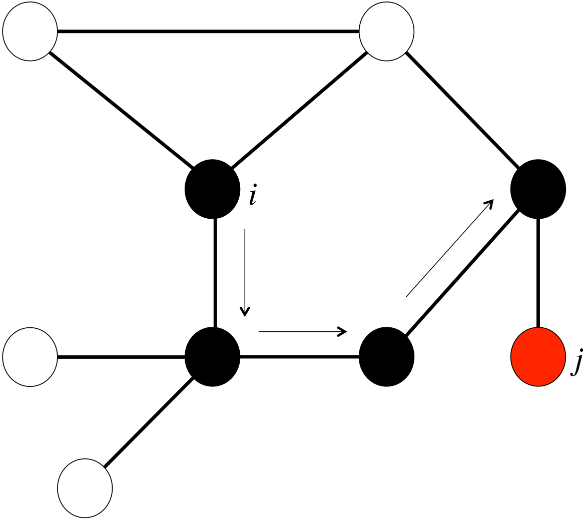

i
j

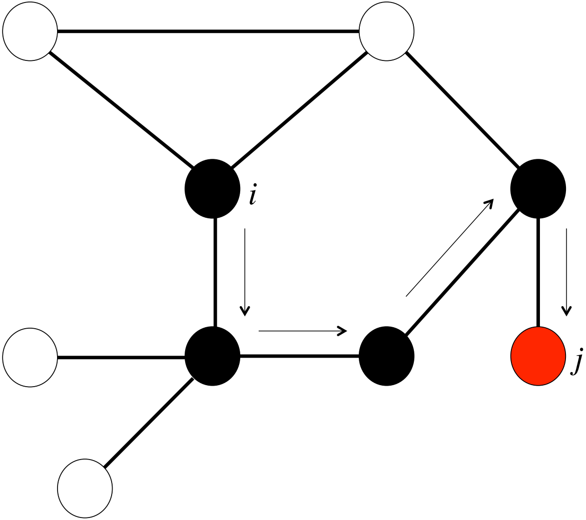
i
j

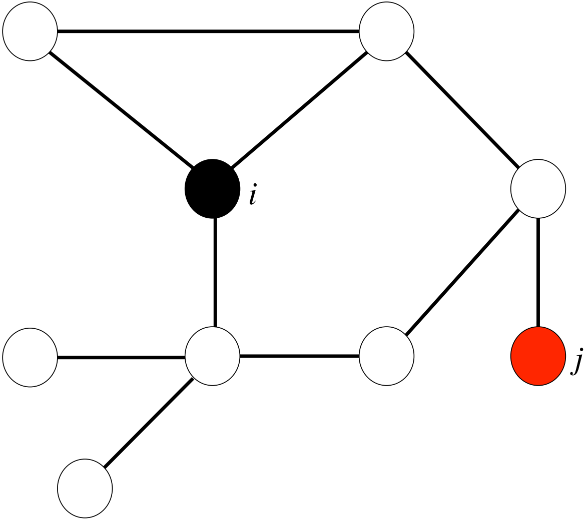
i
j

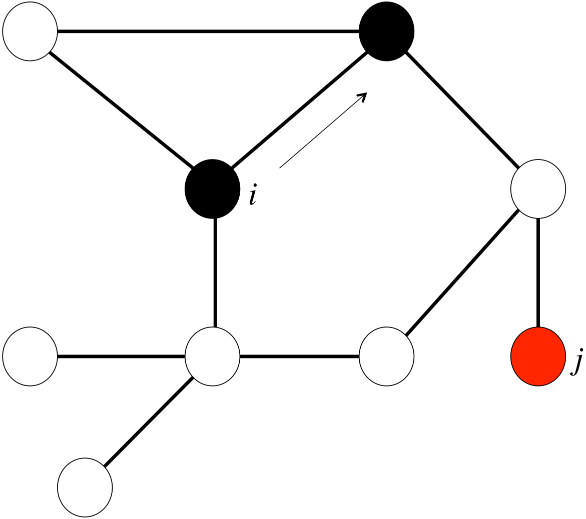
i
j

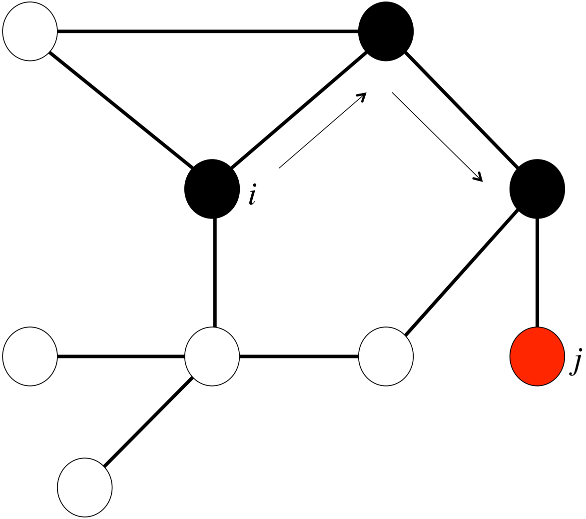
i
j

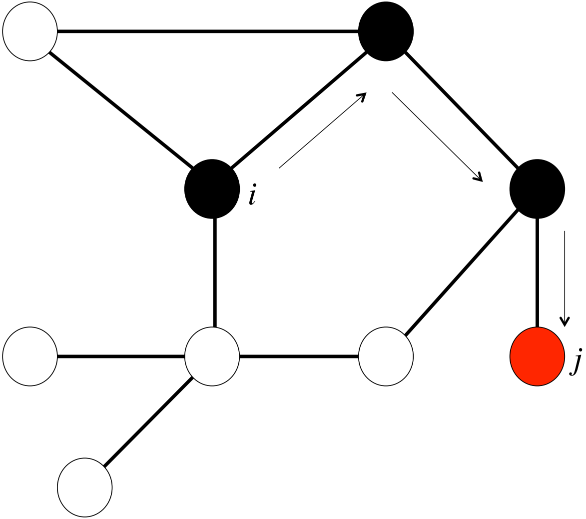

i
j

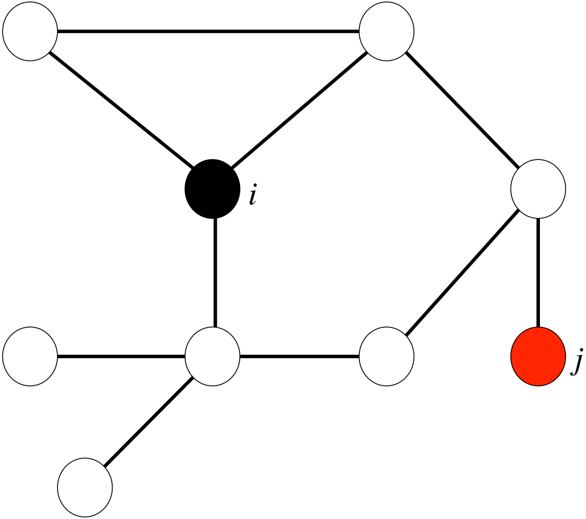
i
j

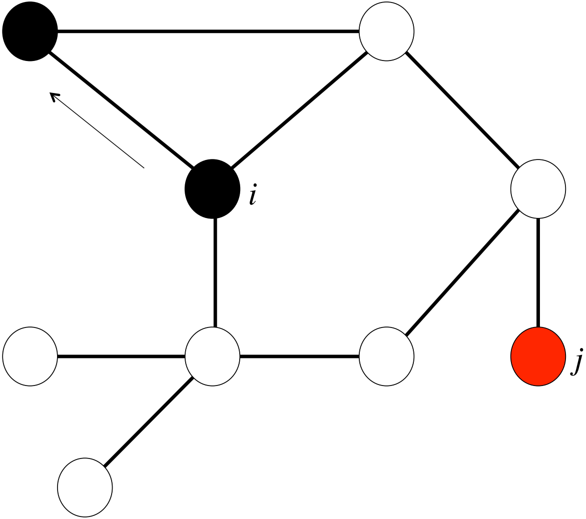
i
j

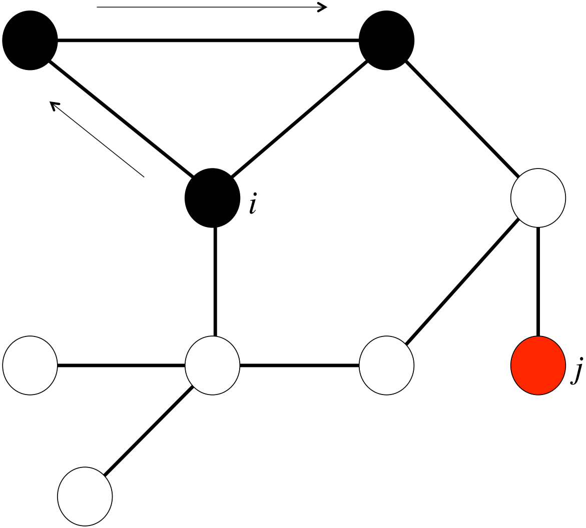

i
j

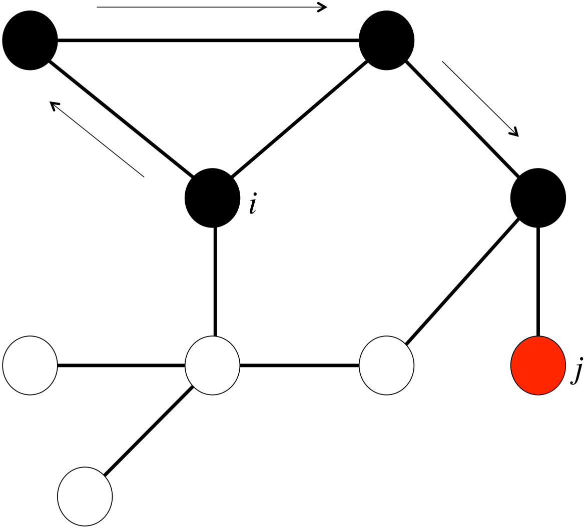
i
j

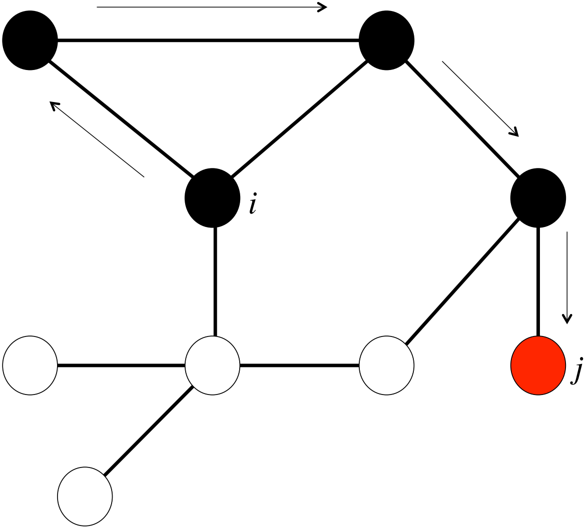
i
j

Degree sequence product

Definition

For a path P, a *degree sequence* along any path P is $(d_i(G))_{i \in P \setminus \{i_0\}}$.

Definition

A *degree sequence product* along P is:

$$\chi_P := \prod_{i \in P \setminus \{i_0\}} d_i(G)$$

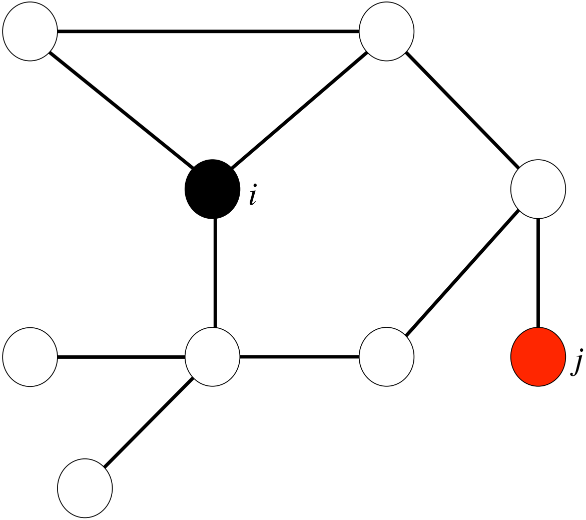
i
j

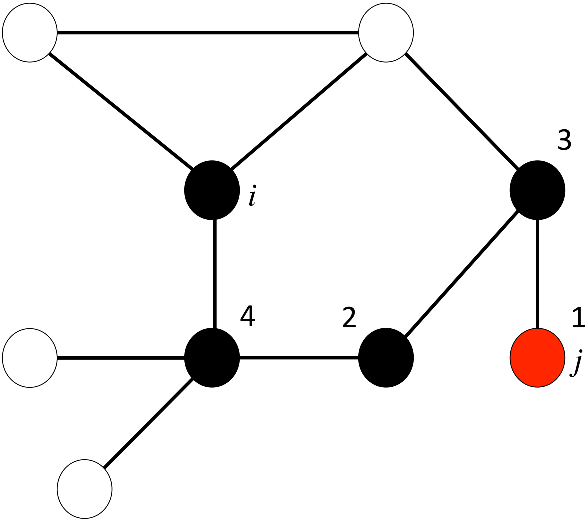
3
i
4
2
1
j

Key proposition

For any G and S_0, let $\mathcal{P}_{ji}$ be the set of all paths beginning at $j \in S_0$ and ending at $i \in V \setminus S_0$ and $\mathcal{P}^*_{ji} \subseteq \mathcal{P}_{ji}$ denote the subset of those paths that exclude any other node in S_0.

Proposition

Suppose firm A is a monopolist. Given a graph G and seed S_0, the probability that node $i \in V \setminus S_0$ adopts product a is:

$$\mathbb{P}_i^A(G, S_0^A) = \sum_{j \in S_0} \sum_{P \in \mathcal{P}^*_{ji}} \frac{1}{\chi_P}$$

See Kempe et al. (2003); Chen et al. (2010).

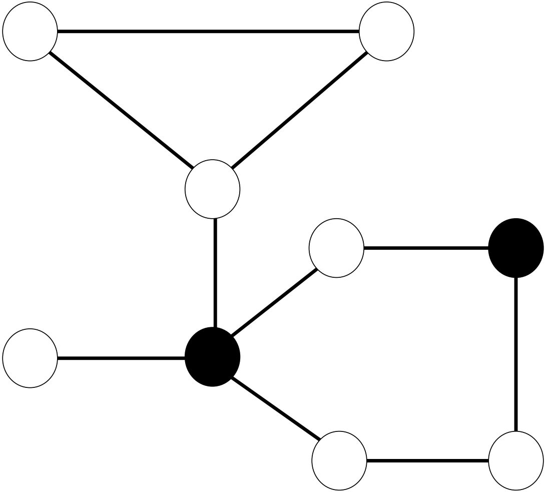

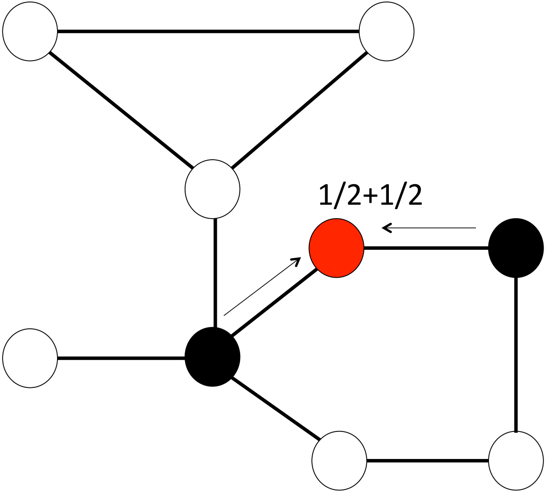
1/2+1/2

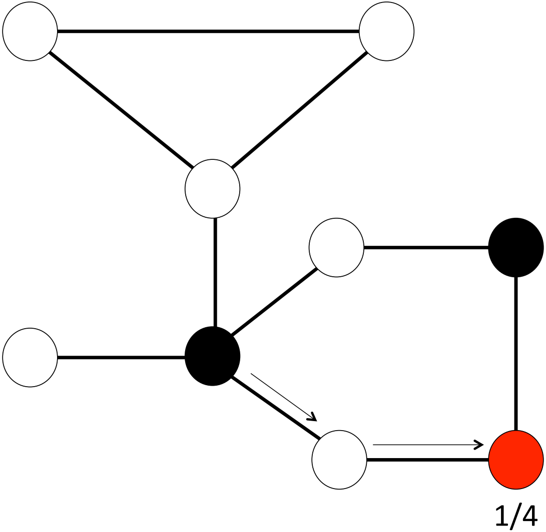

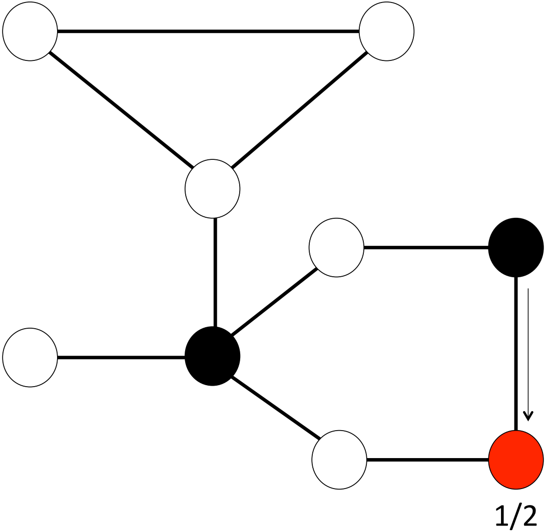
1/2

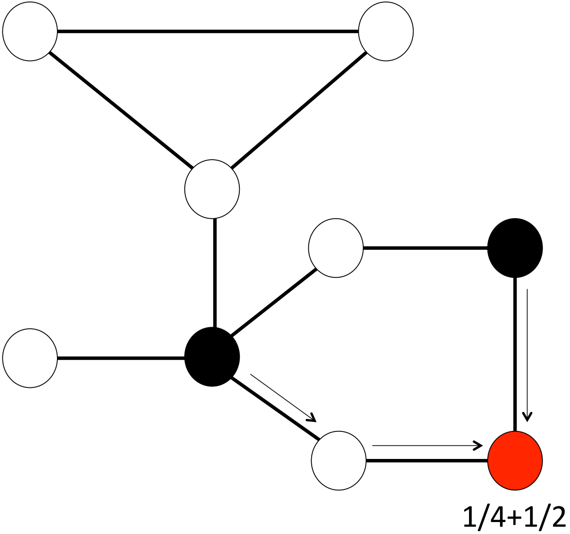
1/4+1/2

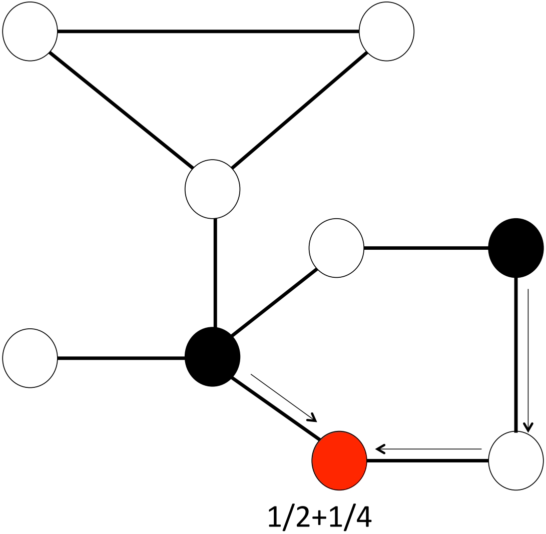
1/2+1/4

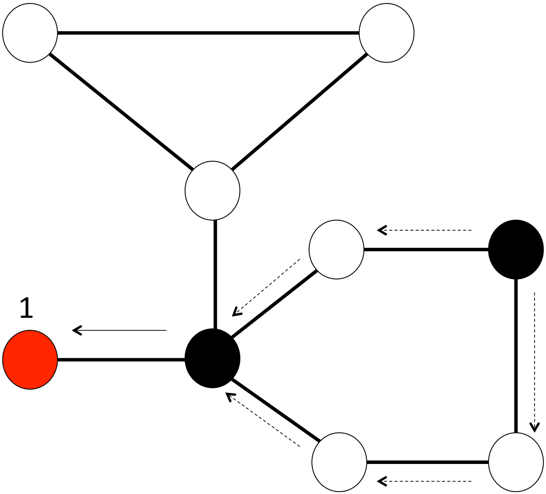
1

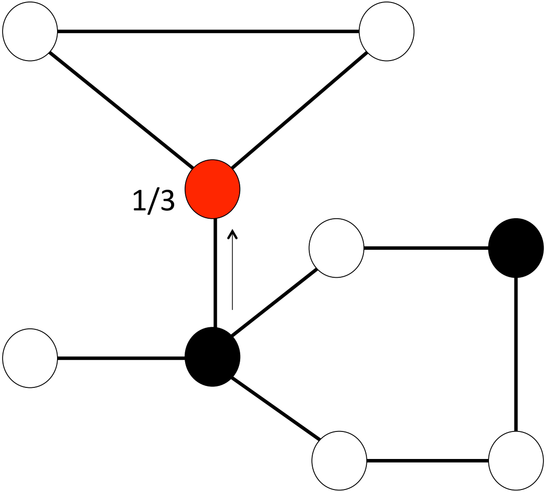
1/3

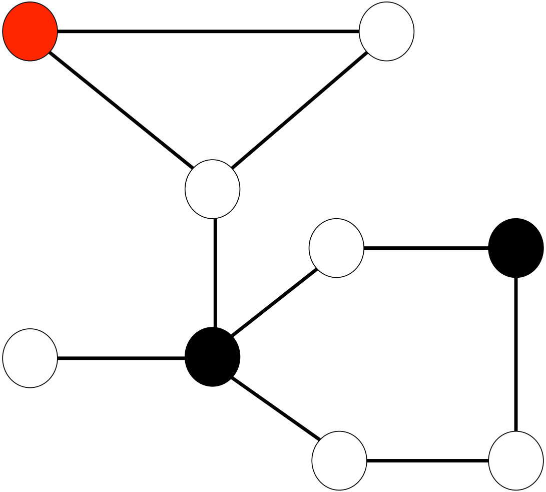

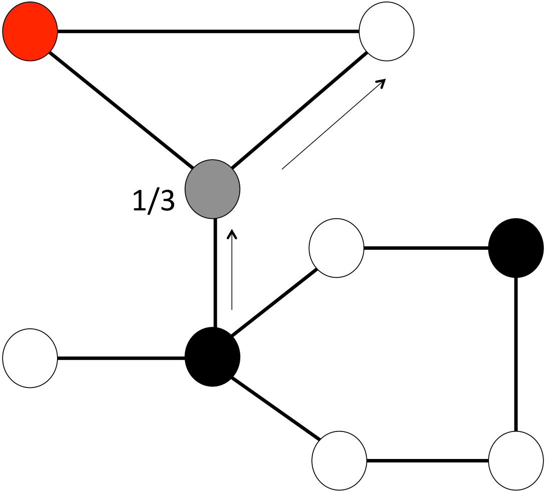
1/3

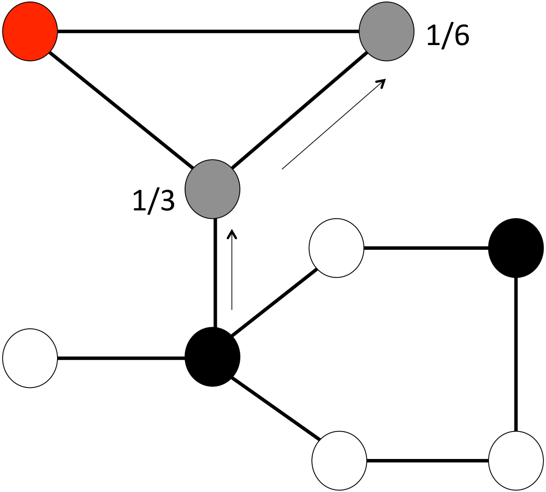
1/6
1/3

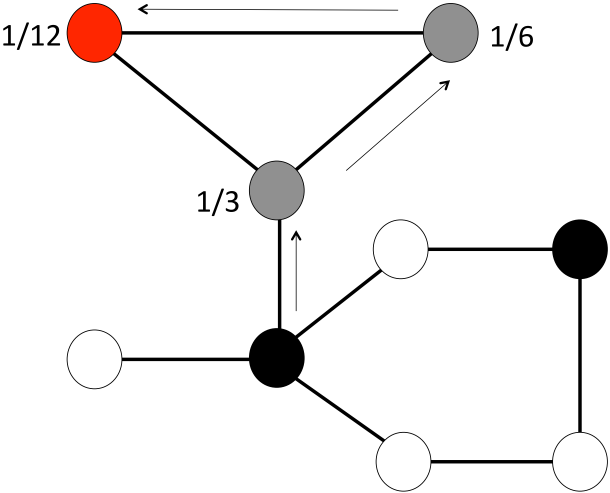
1/12
1/6
1/3

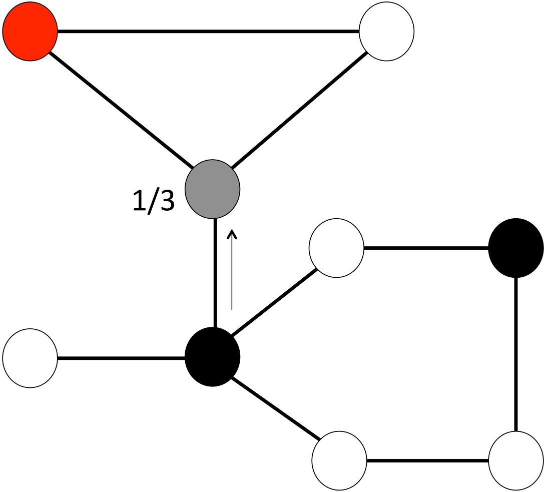
1/3

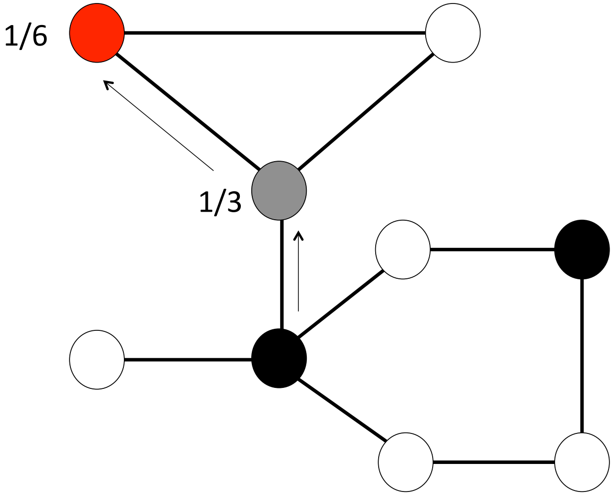
1/6
1/3

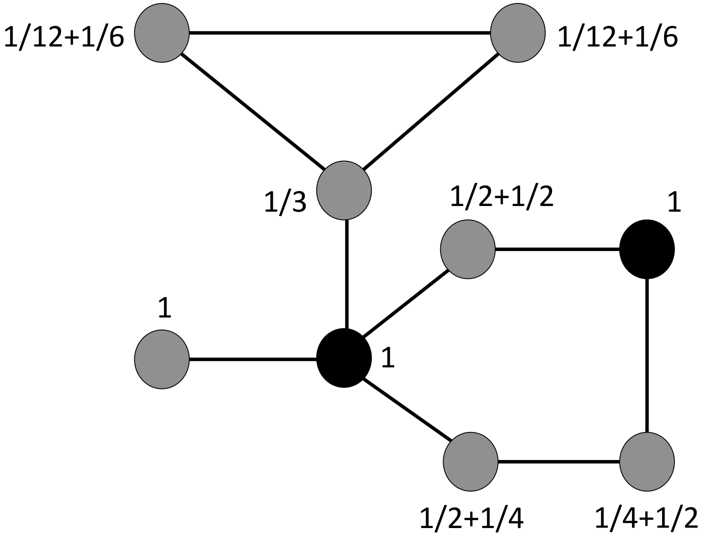
1/12+1/6
1/12+1/6
1/3
1/2+1/2
1
1
1
1/2+1/4
1/4+1/2

Cascade centrality

Definition

Cascade centrality of node i in graph G is the expected number of adopters of product a in that graph given i is the seed and firm A is a monopolist, namely

$$\mathcal{C}_i(G) := \mathbb{E}[S^A(G,\{i\})] = 1 + \sum_{j \in V \setminus \{i\}} \mathbb{P}_j^A(G,\{i\}) = 1 + \sum_{j \in V \setminus \{i\}} \sum_{P \in \mathcal{P}_{ij}} \frac{1}{\chi_P}$$

Back to the duopoly...

Game: uniform thresholds

- Action space of firms A and B: $\Sigma := \Sigma_A \times \Sigma_B := V \times V$
- Action profile $\sigma := (\sigma_A, \sigma_B)$ is simply a pair of nodes.
- Payoff profile: $\pi := (\pi_A(\sigma), \pi_B(\sigma))$ is the expected number of adopter of products a and b.

Game: uniform thresholds

- Action space of firms A and B: $\Sigma := \Sigma_A \times \Sigma_B := V \times V$
- Action profile $\sigma := (\sigma_A, \sigma_B)$ is simply a pair of nodes.
- Payoff profile: $\pi := (\pi_A(\sigma), \pi_B(\sigma))$ is the expected number of adopter of products a and b.

Game: uniform thresholds

- Action space of firms A and B: $\Sigma := \Sigma_A \times \Sigma_B := V \times V$
- Action profile $\sigma := (\sigma_A, \sigma_B)$ is simply a pair of nodes.
- Payoff profile: $\pi := (\pi_A(\sigma), \pi_B(\sigma))$ is the expected number of adopter of products a and b.

Game: uniform thresholds

- For $i \neq j$, let us denote $\Xi(i,j)$ as the set of all paths that begin at i and include (but do not necessarily end) at j.

Proposition

The expected number of adopters of product a (i.e. firm A's payoff) is

$$\pi_A(\sigma_A, \sigma_B) = \begin{cases} \frac{\mathcal{C}_{\sigma_A}}{2} & \text{if } \sigma_A = \sigma_B \\ \mathcal{C}_{\sigma_A} - \epsilon(\sigma_A, \sigma_B) & \text{if } \sigma_A \neq \sigma_B \end{cases}$$

where

$$\epsilon(i,j) = \sum_{P \in \Xi(i,j)} \frac{1}{\chi_P}$$

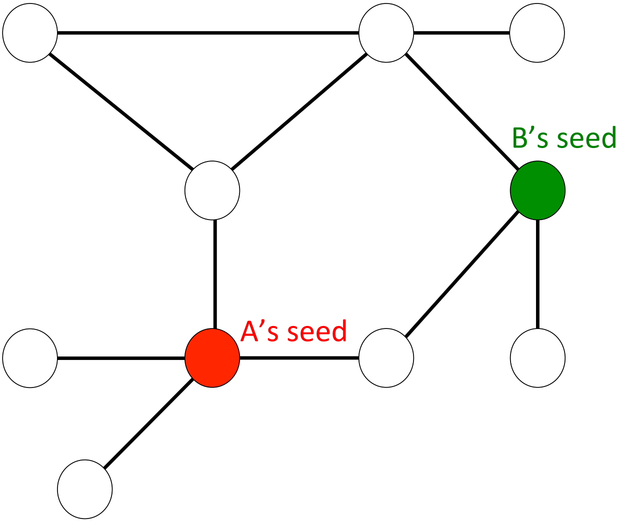
B's seed
A's seed

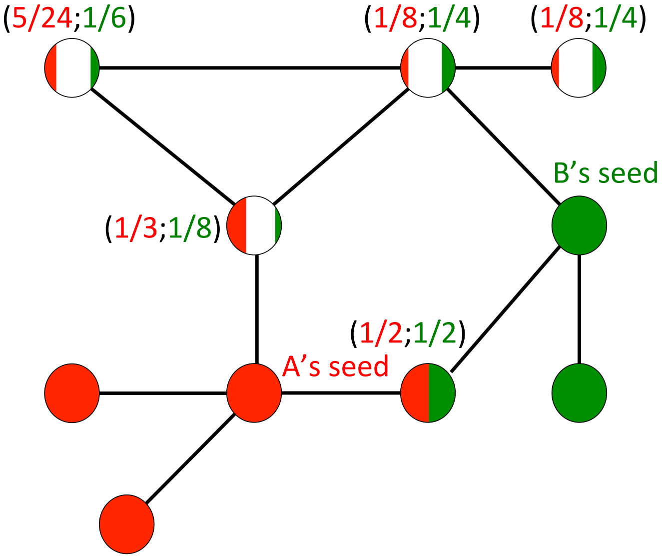

(5/24;1/6)
(1/8;1/4)
(1/8;1/4)
B's seed
(1/3;1/8)
(1/2;1/2)
A's seed

PSNE

- The game is defined as $\Gamma := (\Sigma, \pi)$.

Definition

A profile of actions $\sigma^* := (\sigma_A^*, \sigma_B^*) \in \Sigma$ is a pure-strategy Nash equilibrium if:

- $\pi_A(\sigma_A^*, \sigma_B^*) \geq \pi_A(\sigma_A, \sigma_B^*)$ for all actions $\sigma_A \in \Sigma_A$
- $\pi_B(\sigma_A^*, \sigma_B^*) \geq \pi_B(\sigma_A^*, \sigma_B)$ for all actions $\sigma_B \in \Sigma_B$

- Define Σ^* as the set of all pure-strategy Nash equilibria.

PSNE

- The game is defined as $\Gamma := (\Sigma, \pi)$.

Definition

A profile of actions $\sigma^* := (\sigma_A^*, \sigma_B^*) \in \Sigma$ is a pure-strategy Nash equilibrium if:

- $\pi_A(\sigma_A^*, \sigma_B^*) \geq \pi_A(\sigma_A, \sigma_B^*)$ for all actions $\sigma_A \in \Sigma_A$
- $\pi_B(\sigma_A^*, \sigma_B^*) \geq \pi_B(\sigma_A^*, \sigma_B)$ for all actions $\sigma_B \in \Sigma_B$

- Define Σ^* as the set of all pure-strategy Nash equilibria.

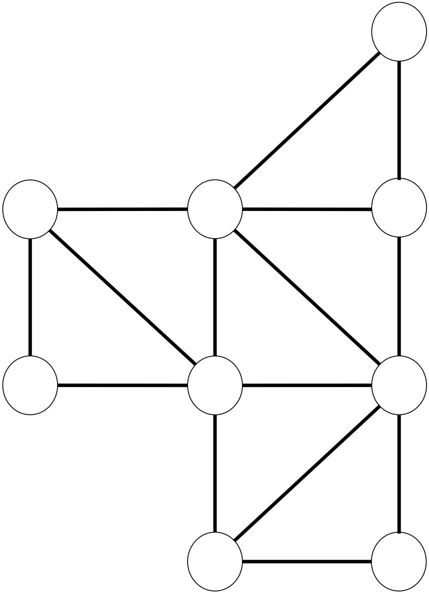

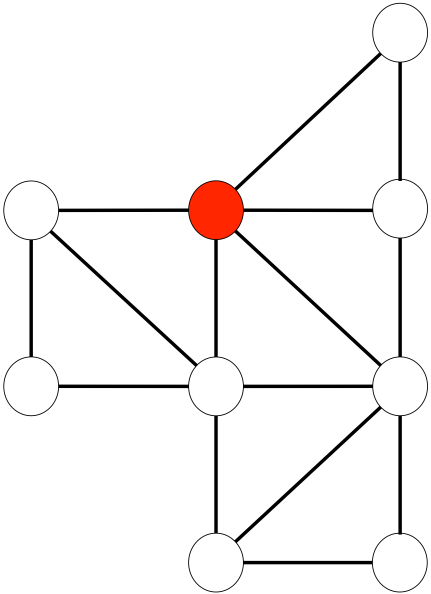

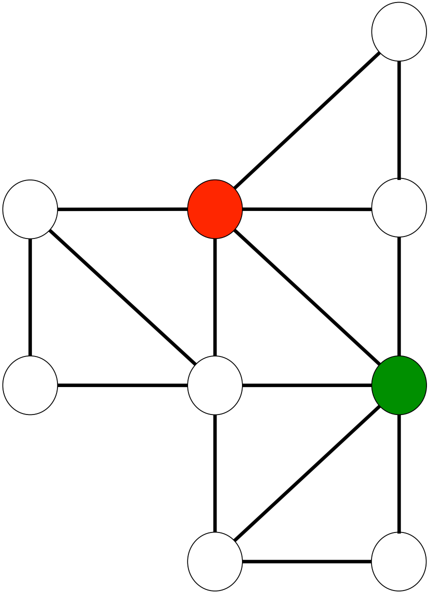

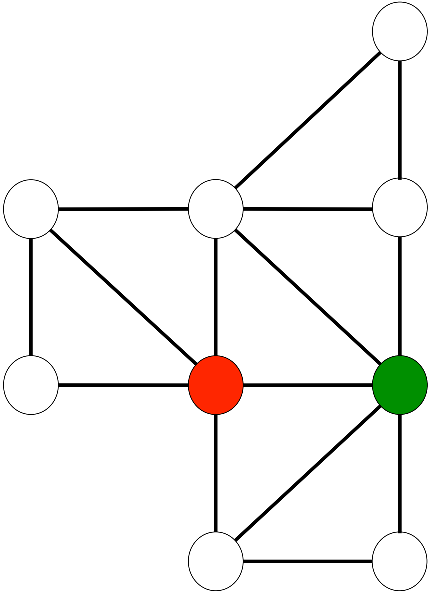

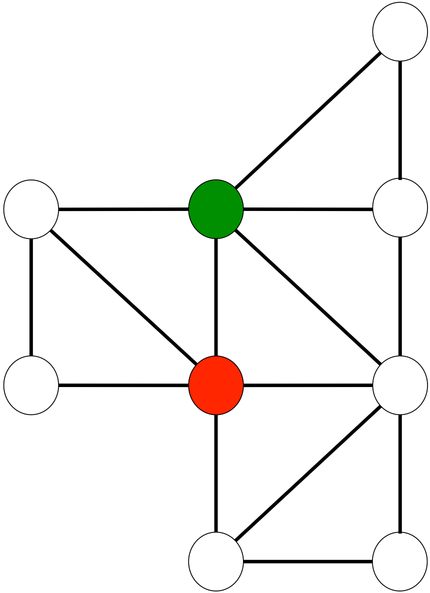

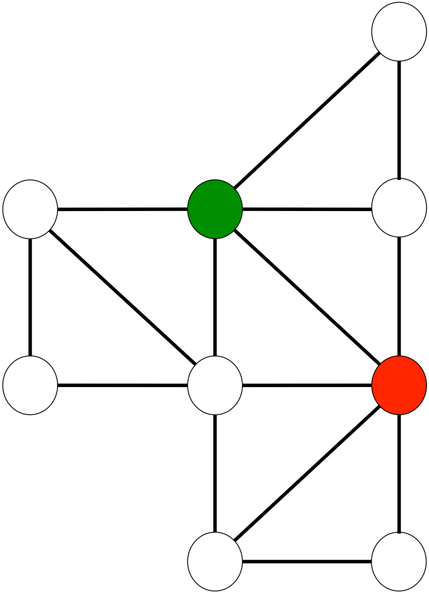

PSNE: existence characterization

Theorem

Consider a duopoly with unit budgets Γ. *Then* Γ *admits at least one PSNE if and only if either:*

1. *There exists* $i \in V$ *such that, for any* $j \in V \setminus \{i\}$*:*

- $$\frac{\mathcal{C}_i}{\mathcal{C}_j} \geq 2 - 2 \cdot \left(\frac{\epsilon(j,i)}{\mathcal{C}_j} \right)$$

then there exists a $\sigma^* = (i, i)$ *PSNE, or...*

PSNE: existence characterization

Theorem

*Consider a duopoly with unit budgets Γ. Then Γ admits at least one PSNE if and only if either Condition **1** is satisfied or*

2. *There exist $i, j \in V$ such that, $\mathcal{C}_i \geq \mathcal{C}_j$ and for any $k \in V \setminus \{i, j\}$:*

- $\dfrac{\mathcal{C}_i}{\mathcal{C}_k} \geq 1 + \dfrac{\epsilon(i,j) - \epsilon(k,j)}{\mathcal{C}_k}$
- $\dfrac{\mathcal{C}_j}{\mathcal{C}_k} \geq 1 + \dfrac{\epsilon(j,i) - \epsilon(k,i)}{\mathcal{C}_k}$
- $\dfrac{1}{2} + \dfrac{\epsilon(i,j)}{\mathcal{C}_j} \leq \dfrac{\mathcal{C}_i}{\mathcal{C}_j} \leq 2 - 2 \cdot \left(\dfrac{\epsilon(j,i)}{\mathcal{C}_j}\right)$

in which case there exists a $\sigma^* = (i, j)$ (and $\sigma^* = (j, i)$ by symmetry) PSNE.

Budget multiplier

Definition

For arbitrary integer budgets $\mathcal{B}_A$ and $\mathcal{B}_B$, the budget multiplier is defined as:

$$\text{BM}(\Gamma) := \max_{\sigma \in \Sigma^*} \frac{\pi_A(\sigma)/\pi_B(\sigma)}{\mathcal{B}_A/\mathcal{B}_B}$$

Theorem

For any Γ that admits at least one PSNE,

$$1 \leq BM < 2$$

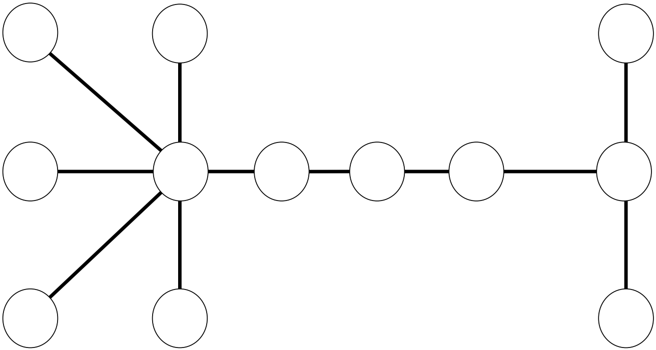

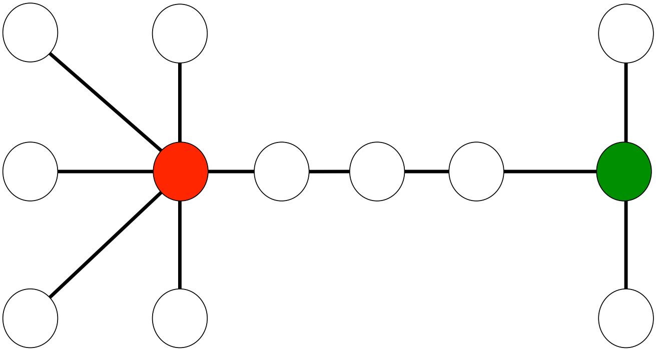

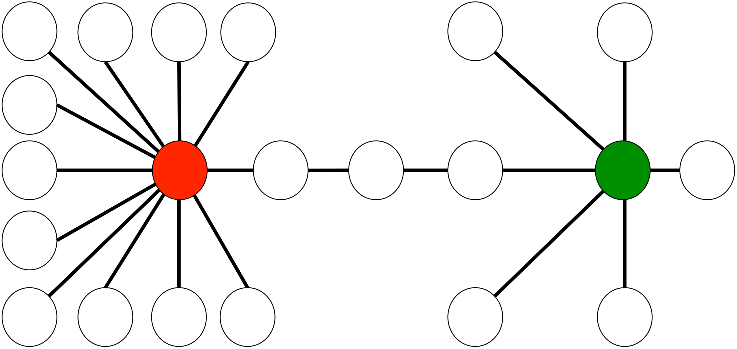

Price of anarchy

Social planner's objective: $Y(\sigma) := \pi_A(\sigma) + \pi_B(\sigma)$ (i.e. firms' total payoffs).

Definition

Price of Anarchy is defined as:

$$\mathrm{PoA}(\Gamma) = \frac{\max_{\sigma \in \Sigma} Y(\sigma)}{\min_{\sigma \in \Sigma^*} Y(\sigma)}$$

Theorem

For any Γ that admits at least one PSNE,

$$1 \leq PoA(\Gamma) < 1.5$$

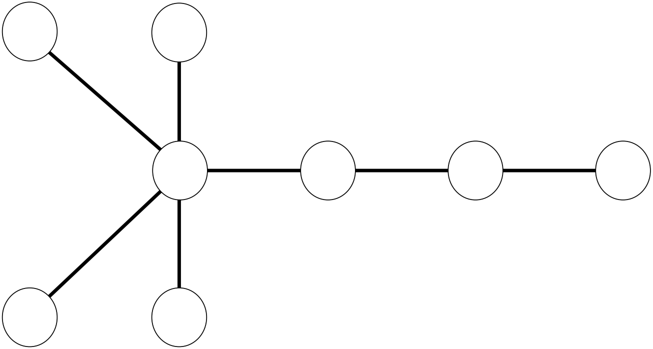

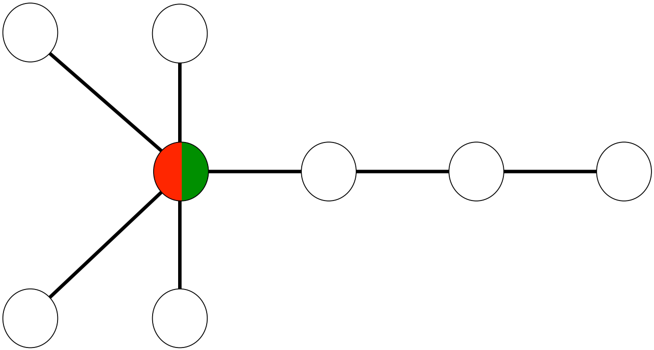

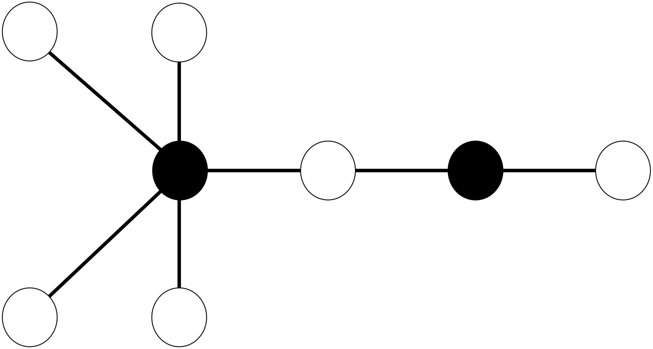

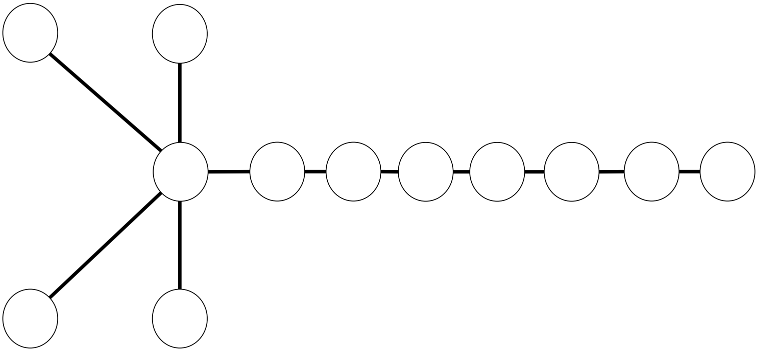

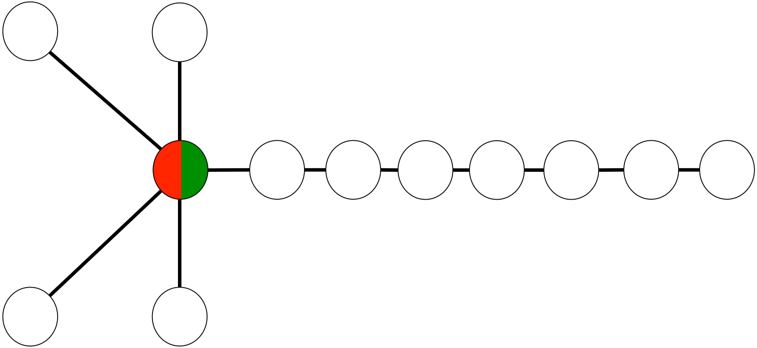

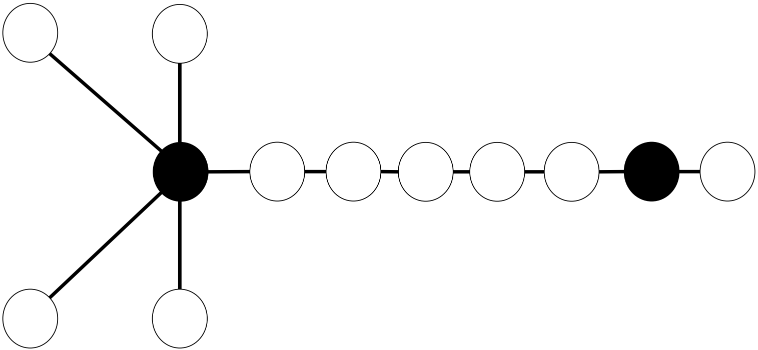

Trees

Proposition

Suppose G is a tree. Then Γ admits at least one PSNE.

Trees

- We provide necessary and sufficient conditions only on the largest and second largest degree of the trees such that:
  - all PSNEs are efficient
  - no PSNEs are efficient
  - at least one PSNE is efficient
  - at least one PSNE is inefficient
  - there is at least one efficient and one inefficient PSNE

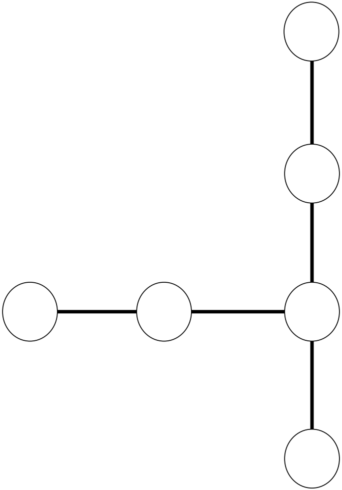

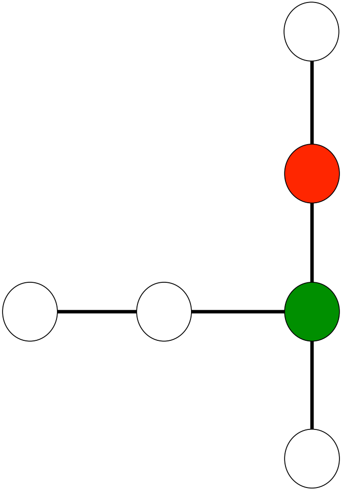

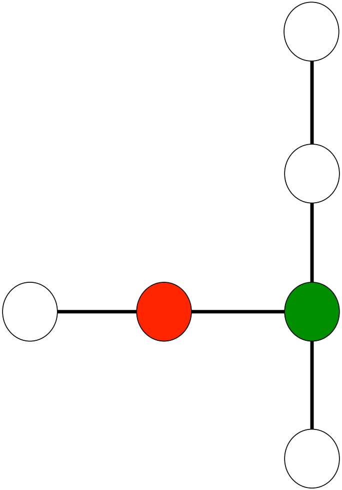

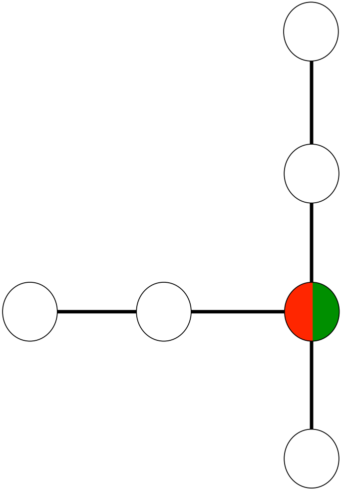

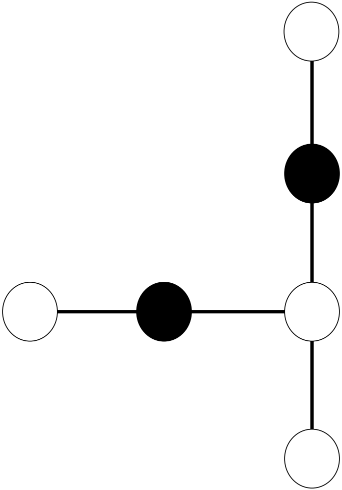

Conclusions

- Using a new notion of *cascade centrality*, we analyzed a tractable cascade process on general networks.
- We applied these insights to studying competitive diffusion.
- The competition model can be extended in a bunch of ways (different quality of products, sequential entry, multiple seeds).

Conclusions

- Using a new notion of *cascade centrality*, we analyzed a tractable cascade process on general networks.
- We applied these insights to studying competitive diffusion.
- The competition model can be extended in a bunch of ways (different quality of products, sequential entry, multiple seeds).

Conclusions

- Using a new notion of *cascade centrality*, we analyzed a tractable cascade process on general networks.
- We applied these insights to studying competitive diffusion.
- The competition model can be extended in a bunch of ways (different quality of products, sequential entry, multiple seeds).

Main conclusion

References I

Noga Alon, Michal Feldman, Ariel D Procaccia, and Moshe Tennenholtz. A note on competitive diffusion through social networks. *Information Processing Letters*, 110(6):221–225, 2010.

Krzysztof R Apt and Evangelos Markakis. Diffusion in social networks with competing products. In *Proceedings of the 4th International Symposium on Algorithmic Game Theory*, pages 212–223. Springer, 2011.

Krzysztof R Apt and Evangelos Markakis. Social networks with competing products. *Fundamenta Informaticae*, 129(3): 225–250, 2014.

Shishir Bharathi, David Kempe, and Mahyar Salek. Competitive influence maximization in social networks. In *Internet and Network Economics*, pages 306–311. Springer, 2007.

Kostas Bimpikis, A. Ozdaglar, and E. Yildiz. Competing over networks. Technical report, Stanford GSB, 2014.

Allan Borodin, Mark Braverman, Brendan Lucier, and Joel Oren. Strategyproof mechanisms for competitive influence in networks. In *Proceedings of the 22nd International Conference on World Wide Web*, pages 141–150, 2013.

Wei Chen, Yifei Yuan, and Li Zhang. Scalable influence maximization in social networks under the linear threshold model. In *Data Mining (ICDM), 2010 IEEE 10th International Conference on*, pages 88–97. IEEE, 2010.

Arastoo Fazeli and Ali Jadbabaie. Targeted marketing and seeding products with positive externality. In *50th Annual Allerton Conference on Communication, Control, and Computing*, pages 1111–1117, 2012a.

Arastoo Fazeli and Ali Jadbabaie. Duopoly pricing game in networks with local coordination effects. In *Proceedings of the 51th IEEE Conference on Decision and Control*, pages 2684–2689, 2012b.

Arastoo Fazeli and Ali Jadbabaie. Game theoretic analysis of a strategic model of competitive contagion and product adoption in social networks. In *Proceedings of the 51th IEEE Conference on Decision and Control*, pages 74–79, 2012c.

Arastoo Fazeli, Amir Ajorlou, and Ali Jadbabaie. Optimal budget allocation in social networks: Quality or seeding. Preprint 1404.1405, arXiv, April 2014.

Sanjeev Goyal and Michael Kearns. Competitive contagion in networks. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 759–774, 2012.

Harold Hotelling. Stability in competition. *The Economic Journal*, 39(153):41–57, 1929.

References II

David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *SIGKDD'03 Proceedings of the ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 137–146, 2003.

Wenjun Mei and Francesco Bullo. Modeling and analysis of competitive propagation with social conversion. In *IEEE Conference on Decision and Control*, 2014.

Sunil Simon and Krzysztof R Apt. Choosing products in social networks. In *Internet and Network Economics*, pages 100–113, 2012.

Vasileios Tzoumas, Christos Amanatidis, and Evangelos Markakis. A game-theoretic analysis of a competitive diffusion process over social networks. In *Internet and Network Economics*, pages 1–14. Springer, 2012.