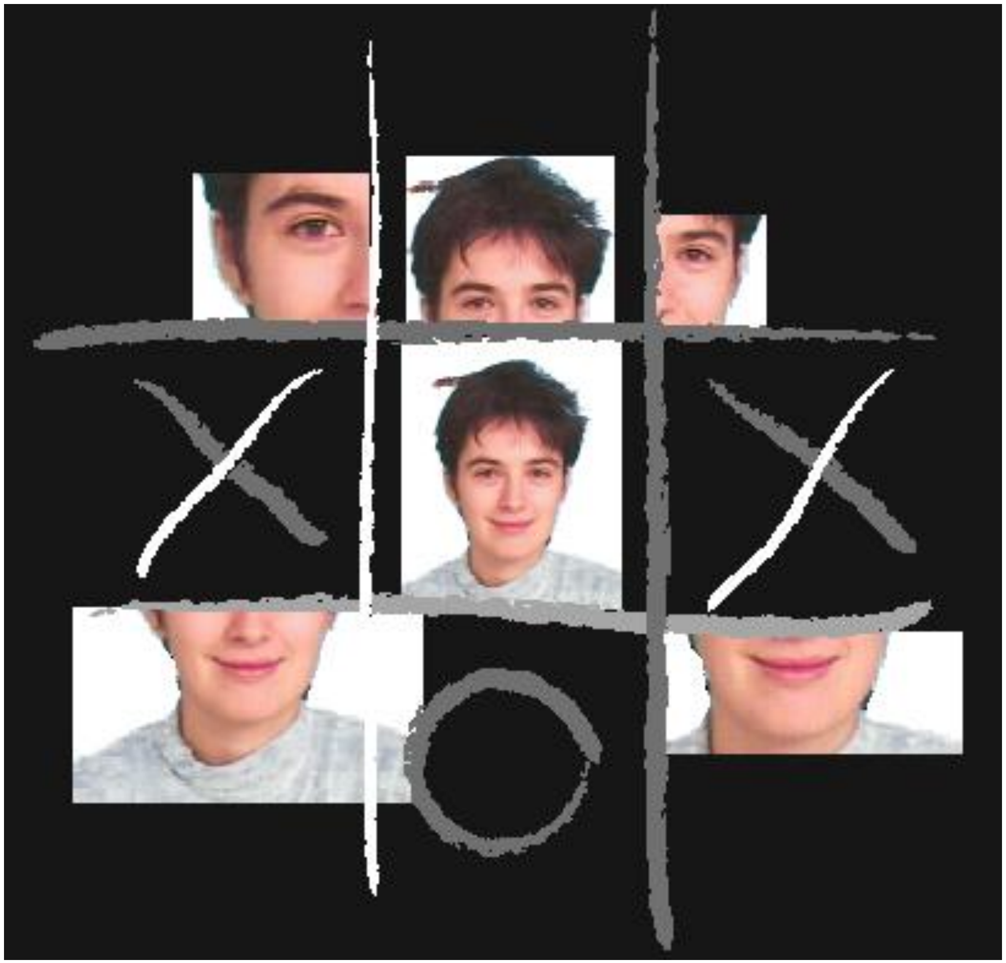

Delay-critical systems and networks

- Multimedia compression & proc
- Rigorous cross-layer design
- Mission-critical networks & systems
- Energy-efficient multimedia sys
- HW/SW co-design for media apps
- Real-time stream mining

Designs for networked communities

- New classes of games & learning
- Microeconomics of grid computers
- Policing in networks and systems
- Design and incentives in Social Nets (current talk)

Delay-critical networking and processing

Multimedia compression, processing and networking

MPEG, Philips

Rigorous methods for cross-layer design (dynamic environments)

NSF, Intel, HP, Microsoft

Delay-critical Networking and Online Learning

NSF, ONR, Intel, Cisco

Real-time Stream Mining

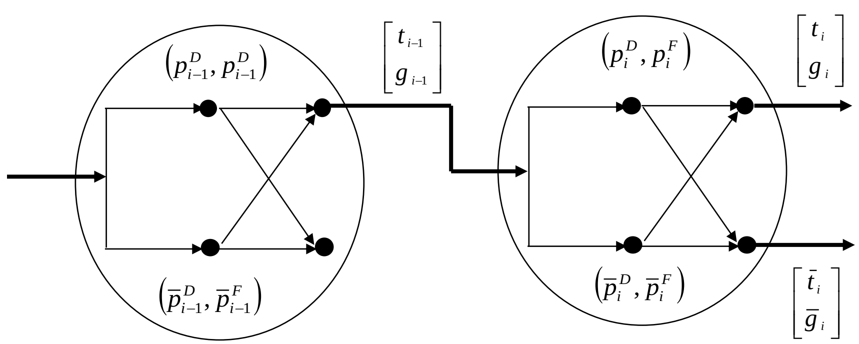

IBM, NSF

Information processing and economics

New Classes of Engineering Games

NSF

Network economics

Policing in networks (Intervention)

Design and Incentives in Networked Communities

NSF, IBM

Design and incentives in Social Networks

Mihaela van der Schaar

Electrical Engineering, UCLA

Web: medianetlab.ee.ucla.edu

Networked Communities – Emergence

- Rapid expansion of social networks, social computing, P2P networks, grid computing etc.
- Networked communities allow individuals and organizations to get connected and build relationships.

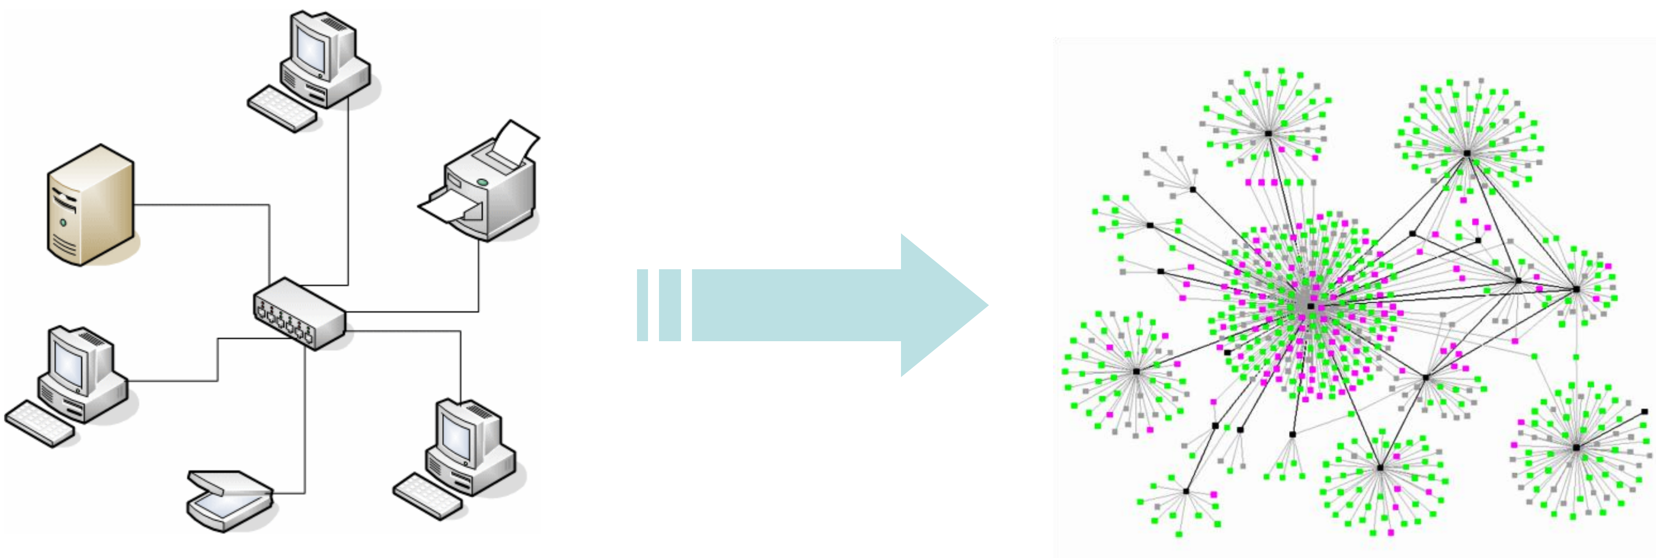

Networked communities = collection of self-interested, learning agents (people, machines, software …)

Promises and Challenges

- **Promises**
  - Network communities open up many opportunities for users to gain from various types of cooperative interactions (e.g., share information or content, and provide resources such as bandwidth, storage space, and computing power).
- **Goals**
  - Design networks and protocols that are socially optimal and induce compliance by agents
- **Challenges**
  - Self-interested, rational, intelligent, and heterogeneous users
  - Large-scale
  - Ongoing interactions
  - Limited (truthful?) information
  - Privacy and security

Existing Approaches

- **Networking**
  - Studies network architecture and design, communication protocols, network operations and management, etc.
  - Example: Network utility maximization (NUM)
  - ☹ Mostly, network users are assumed to be obedient (machines following given rules), but social networks consist of self-interested, self-organizing users.
- **Signal Processing**
  - Studies how to collect, learn, and process information from signals.
  - Users in social networks create and process a lot of information and signals.
  - ☹ Processing techniques do not consider the social context, the self-interest of users etc.

Existing Approaches

- **Control Theory**
  - Studies distributed control of autonomous agents to achieve the system objective.
  - ☹ Again, agents are assumed to be obedient.

- **Classical Game Theory**
  - Studies interactions of self-interested strategic players.
  - Descriptive, often not constructive
  - ☹ Not prescriptive - The issues of system design are neglected (The game in which players interact is taken as given, rather than being considered as something that can be designed.)

Where are we coming from and where are we going?

Classical Engineering

- Nodes: Cooperative
- System designer has a high degree of control: prescribes decision rules for nodes
- Systems assume compliance
- Social and individual goals coincide, e.g. utility maximization
- Truthful information revelation assumed
- Mostly single-agent learning, prescriptive

Next-generation Engineering

- Agents: Self-interested, strategic
- System designer can control only a playground on which agents interact, but the agents choose how to play
- System compliance not guaranteed - Strategy-proof protocols needed
- Social and individual goals in conflict, e.g. system collapse
- Agents may lie/hide information
- Multi-agent learning

Our contributions

- Paradigm shift: new theories to capture features unique to networked communities.
- Develop new principles for analyzing and designing distributed systems composed of self-interested agents.
  - What information and decision-making rules to provide for agents so that they can optimize their decisions given their local information?
  - Given network and application constraints, how to design the system?
  - If agents can themselves determine the connectivity, how to design the rules that influence which agents should produce information, process information and to whom they should distribute the information?
  - What new applications can emerge in such networks?
  - How to design effective incentive schemes that lead to socially-efficient system designs?

Classification of Incentive Schemes

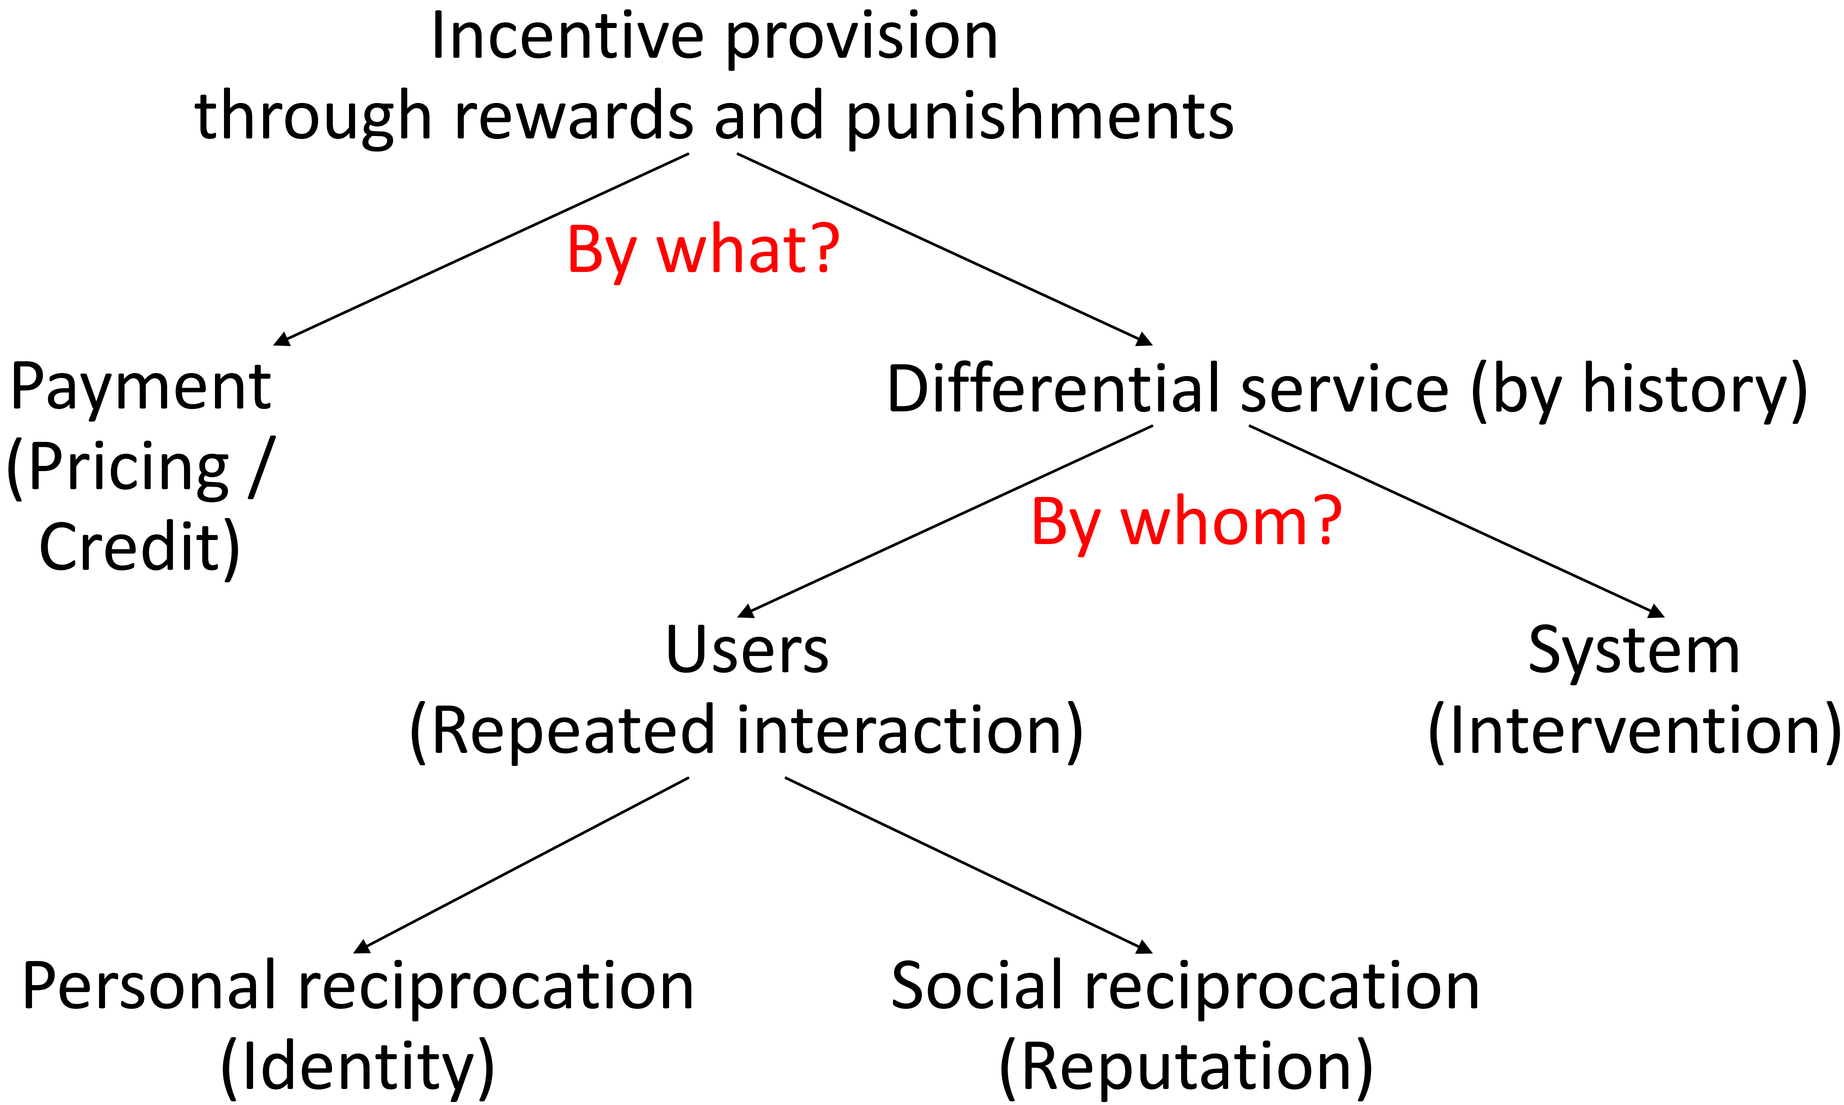

What can be achieved?

Approach	Performance	Requirements for protocol designer	
		Enforcement on peers	Knowledge about the system
Non-cooperative without any incentive scheme	Inefficient	None	None
Cooperative	Socially optimal	Actions	Complete knowledge
Payment scheme	Can be socially optimal	Payments	Complete knowledge (can be replaced with learning)
Reciprocation scheme	Can be socially optimal	None (self-enforcing)	Complete knowledge (can be replaced with learning)
Intervention scheme	Can be socially optimal	Intervention by the system	Complete knowledge (can be replaced with learning)

Jaeok Park and Mihaela van der Schaar, "A Game Theoretic Analysis of Incentives in Content Production and Sharing over Peer-to-Peer Networks," *IEEE Journal of Selected Topics in Signal Processing – Special issue on Social Networks*, vol. 4, no. 4, pp. 704-717, August 2010.

Who interacts with whom?

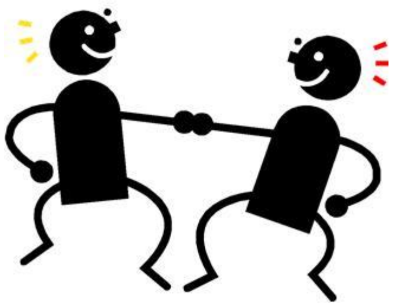

- Random partners: Design of social norms to sustain cooperation
- Choosing partners: Design of dynamic personal reciprocation policies
- Community formation: Link formation and information production, sharing and consumption in networked communities

Random partners

Design of Social Norms to Sustain Cooperation

Yu Zhang, Jaeok Park, and Mihaela van der Schaar,
“Social Networks Protocol Designs Based on Social Norms,”
Gamenets 2011, ICASSP 2011, ITA 2011, submitted for Journal Publication

Related Work

Approach	Works	Limitations
Personal reciprocation	Buragohain, Agrawal, and Suri (2003) - *CS* Ma, Lee, Lui, and Yau (2006) - *CS* Cohen (2003) - *CS*	Not effective in a large-scale network with a high turnover rate and asymmetry of interests
Social reciprocation (reputation)	Feldman, Lai, Stoica, and Chuang (2006) - *CS* Kamvar, Schlosser, and Molina (2006) - *EE*	Focuses on practical techniques to aggregate, process, and disseminate information generated in a system
Social reciprocation (social norm)	Kandori (1992) - *Econ* Blanc, Liu, and Vahdat (2007) - *CS*	K: Design missing K: Ignores practical complications such as reputation update errors, turnover of population, and whitewashing attempts K&BLV: Consider only the limiting case as the discount factor goes to 1

Setup

- We consider a large-scale networked community where users can help each other by providing a sort of service.
- In each period, a user generates exactly one service request, which is sent to another user chosen randomly.
- There is no connectivity constraint, and thus a service request from a user can be sent to any other user in the network.
- We assume uniform random matching where each user receives exactly one service request per period.
- In other words, each user is involved in two matches in each period, once as a sender of a request and the other as a receiver of a request.

Game Played by a Pair of Matched Peers

- Players
  - Client: peer requesting a file
  - Server: peer receiving an upload request
- Actions
  - Client: no action to choose
  - Server: $a \in \mathcal{A} = \{S, NS\}$
    - S ("Serve"): provide service (i.e., upload the requested file)
    - NS ("Not Serve"): refuse to provide service (i.e., do nothing)
- Payoffs
  - When a file is uploaded, the client receives a benefit of b>0 while the server suffers a cost of c>0.
  - We assume that b>c so that the net social benefit from a service is positive.

	Server: S	Server: NS
Client	b, -c	0, 0

Applications

Social Welfare

- Social welfare in a period = the average payoff of peers.

	Server S	Server NS
Client	b, -c	0, 0

- Social welfare is maximized when the server chooses S in every match, maximum = b-c.
- When the server chooses its action to myopically maximize its payoff, the optimum/equilibrium behavior is NS => suboptimal social welfare = 0.
- Therefore, incentive schemes are necessary to induce self-interested servers to choose S.

Social Norms

- Social norm is defined as the rules for appropriate and inappropriate behaviors
  - Compliance
  - Rewards (present and future)
  - Punishments (present and future)
- We consider a social norm using reputation.
  - Each peer is tagged a reputation label.

Formal Representation of a Social Norm

- A social norm is represented by $\kappa = (\Theta, \theta_0, \tau, \sigma)$.
  - Θ : set of reputation labels
  - $\theta_0 \in \Theta$: initial reputation
  - $\tau : \Theta \times \Theta \times \mathcal{A} \rightarrow \Theta$: reputation update rule
    - $\tau(\theta, \tilde{\theta}, a_R)$ is the new reputation for a server with current reputation θ when it is matched with a client with reputation $\tilde{\theta}$ and its action is reported as a_R.
  - $\sigma : \Theta \times \Theta \rightarrow \mathcal{A}$: social strategy
    - $\sigma(\theta, \tilde{\theta})$ is the approved action for a server with reputation θ that is matched with a client with reputation $\tilde{\theta}$.

reputation scheme

What do agents know? What choices do they have?

Design Choice

- The design choice in the protocol design problem is a social norm $(\Theta, \theta_0, \tau, \sigma)$.
- Starting point- we impose the following restrictions on the reputation scheme (Θ, θ_0, τ).
  - Θ is finite, i.e., $\Theta = \{0, 1, \ldots, L\}$ for some integer L.
  - $\theta_0 = L$.
- We call the above reputation scheme the maximum punishment reputation scheme with punishment length L.
- The design choice is reduced to (L, σ).

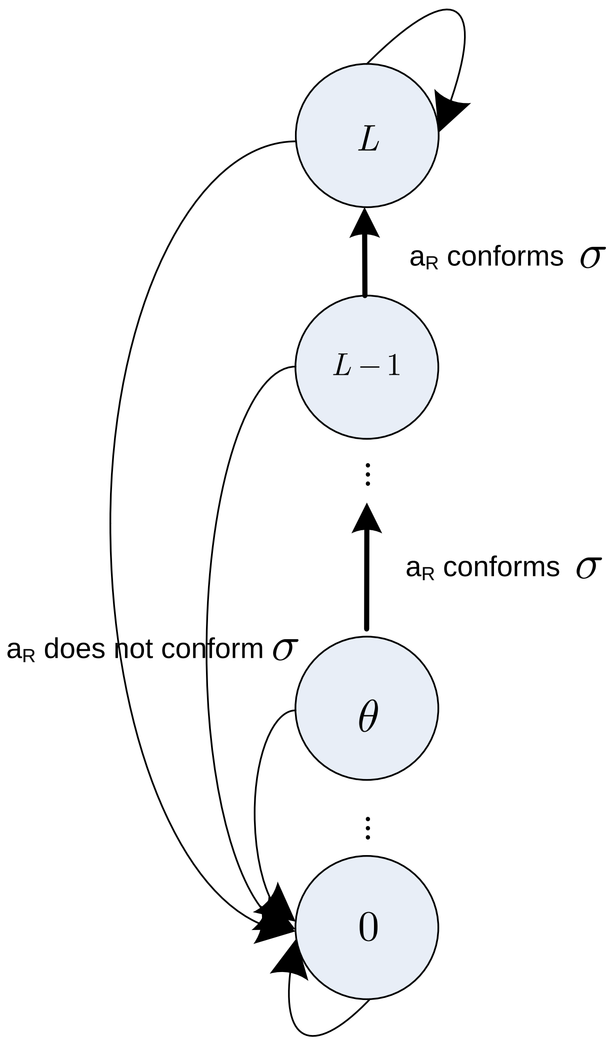

System Parameters

- There are five system parameters that affect the solution of the protocol design problem, $(b, c, \alpha, \beta, \varepsilon)$.
  - b: benefit from service
  - c: cost of service
  - $\alpha \in [0,1]$: turnover rate (the fraction of peers that leave and join the network between two consecutive periods)
  - $\beta \in [0,1)$: time discount factor (patience)
  - $\varepsilon \in [0, 0.5]$: report error probability (the probability that the action of a server is misreported)

Focus on long-run/steady-state

- Lemma 1: For any $\alpha \in [0,1]$ and $\varepsilon \in [0,0.5]$, there exists a unique stationary distribution of reputations:

$$\eta_\tau(\theta) = (1-\alpha)^{\theta+1}(1-\varepsilon)^{\theta}\varepsilon, \text{ for } 0 \le \theta \le L-1,$$

$$\eta_\tau(L) = \frac{(1-\alpha)^{L+1}(1-\varepsilon)^{L}\varepsilon + \alpha}{1-(1-\alpha)(1-\varepsilon)}.$$

Moreover, the distribution converges to the stationary distribution starting from any initial distribution.

Agents' Payoffs and Social Welfare

- The expected period payoff of the agent following the social strategy, evaluated before it is matched:

$$v_\kappa(\theta) = \sum_{\tilde{\theta}\in\Theta} \eta_\tau(\tilde{\theta}) b_\sigma(\tilde{\theta},\theta) - \sum_{\tilde{\theta}\in\Theta} \eta_\tau(\tilde{\theta}) c_\sigma(\theta,\tilde{\theta})$$

stationary distribution of reputations

download benefit

upload cost

- The expected long-term payoff of the agent:

$$v_\kappa^\infty(\theta) = v_\kappa(\theta) + \beta(1-\alpha)\sum_{\theta'} p_\kappa(\theta' \mid \theta) v_\kappa^\infty(\theta')$$

time discount factor

probability of remaining in the network

transition probability of reputations

- System - Social welfare: average period payoff of peers in the stationary distribution

$$U_\kappa = \sum_\theta \eta_\tau \;\; \theta \;\; v_\kappa \;\; \theta$$

Sustainable Social Norms

We do not just want to have rules.

We *must* have rules that agents *will follow*!

- A reputation scheme τ sustains a social strategy σ, or a social norm $\kappa = (\tau, \sigma)$ is sustainable if

$$-c_{\sigma}(\theta,\tilde{\theta}) + \delta \sum_{\theta'} p_{\kappa}(\theta' \mid \theta) v_{\kappa}^{\infty}(\theta') \geq -c_{\sigma'}(\theta,\tilde{\theta}) + \delta \sum_{\theta'} p_{\kappa,\sigma'}(\theta' \mid \theta,\tilde{\theta}) v_{\kappa,\sigma'}^{\infty}(\theta')$$

for all σ', for all $(\theta,\tilde{\theta})$. ($\delta = \beta(1-\alpha)$)

Lemma 2: τ sustains σ if and only if

$$\underbrace{\delta(1-2\varepsilon) v_{\kappa}^{\infty}(\min\{\theta+1, L\})}_{\text{Follow}} \geq \underbrace{c + \delta(1-2\varepsilon) v_{\kappa}^{\infty}(0)}_{\text{Deviate}}$$

Protocol Design Problem

- An optimal social norm is a social norm that maximizes social welfare *among sustainable social norms.*
- The protocol design problem can be expressed as

$$\underset{(L,\sigma)}{\text{maximize}} \quad U_{\kappa} = \sum_{\theta} \eta_{\tau}(\theta) v_{\kappa}(\theta)$$

$$\text{subject to} \quad \delta(1-2\varepsilon)\left[v_{\kappa}^{\infty}(\min\{\theta+1, L\}) - v_{\kappa}^{\infty}(0)\right] \geq c,$$

$$\forall \theta \text{ such that } \exists \tilde{\theta} \text{ such that } \sigma(\theta, \tilde{\theta}) = S.$$

Optimal Value

- U^* = optimal value of the protocol design problem.

  $\cancel{\beta \to 1}$ $\cancel{\varepsilon \to 0}$

- Theorem 1:

  (i) $0 \le U^* \le b - c$.

  (ii) $U^* = 0$ if $\frac{c}{b} > \frac{\beta(1-\alpha)(1-2\varepsilon)}{1-\beta(1-\alpha)(2-3\varepsilon)}$.

  $\frac{c}{b} \le \beta(1-\alpha)$

  (iv) $U^* < b - c$ if $\varepsilon > 0$

  (v) $U^* \ge [1 - \ 1 - \alpha \ \ \varepsilon] \ b - c$ if $\frac{c}{b} \le \beta(1-\alpha)(1-2\varepsilon)$

- Corollary 1: For any (b, c) such that b>c,

  (i) U^* converges to $b - c$ as $\beta \to 1$, $\alpha \to 0$, and $\varepsilon \to 0$.

  (ii) U^* converges to 0 as $\beta \to 0$, $\alpha \to 1$, or $\varepsilon \to 0.5$.

How big does L have to be?

- Let U_L^* and σ_L^* be the optimal value and the optimal social strategy, respectively, of the protocol design problem given L.
- Proposition 2: $U_L^* \geq U_{L'}^*$ for all L and L' such that $L \geq L'$.

$$U^* = \lim_{L\to\infty} U_L^* = \sup_L U_L^*$$

Conjecture: The supremum is achieved by some finite L.

Cost – benefit tradeoff when choosing L

Optimal Social Strategy For Fixed L

- We now study the structure of σ_L^*.
- Theorem 3: (i) If $\sigma_L^*(0,\tilde{\theta}) = S$ for some $\tilde{\theta}$, then there is a θ^* such that $\sigma_L^*(0,\hat{\theta}) = S$ for all $\hat{\theta} \geq \theta^*$.

  (ii) There is a $\theta^{**} \in \{1,\ldots,L-1\}$, such that $\sigma_L^*(\theta, L) = S$ when $\theta \geq \theta^{**}$.

  (iii) If $\sigma_L^*(L,\tilde{\theta}) = S$ for some $\tilde{\theta}$, then $\sigma_L^*(L,L) = S$.

Optimal Social Strategy For Fixed L

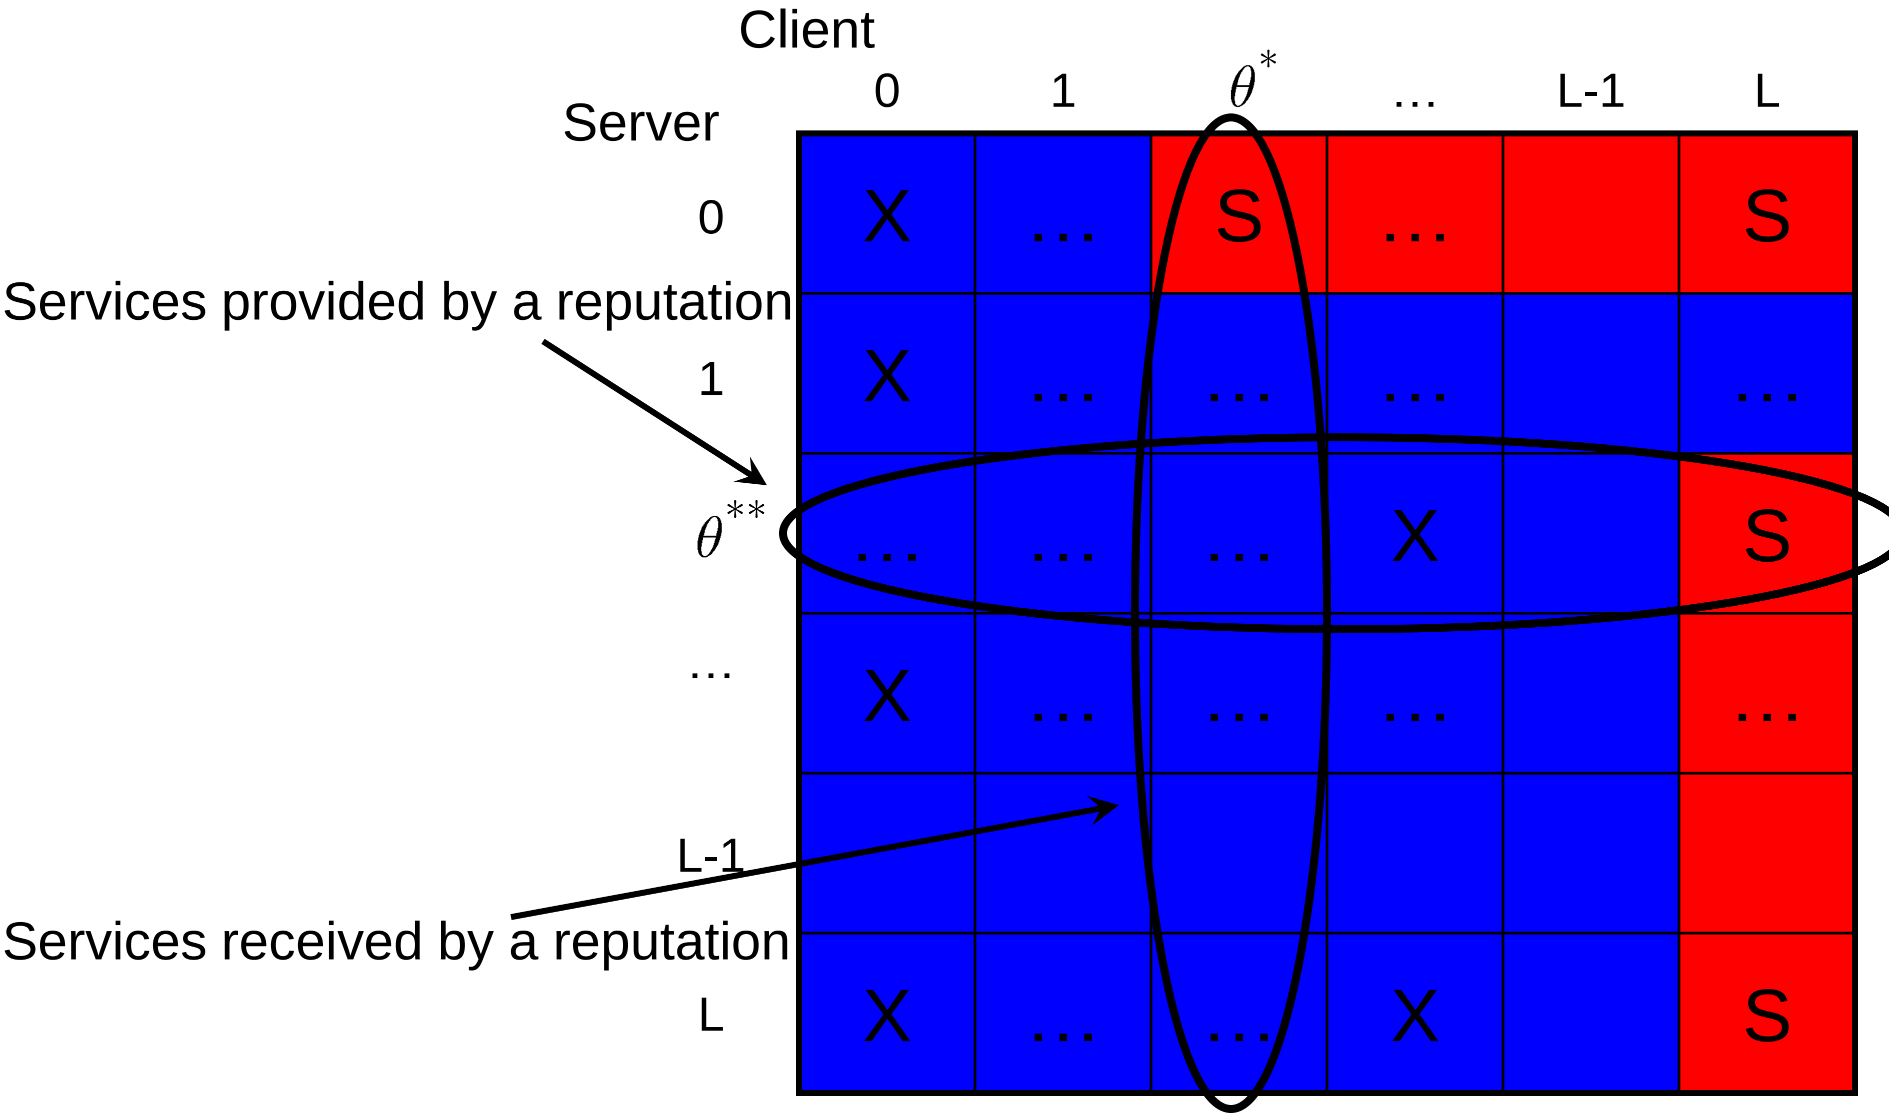

Illustration with L=1

- 16 possible social strategies
- only 4 of these strategies *can possibly be optimal*
- 12 of these strategies *can never be optimal*

- By Theorem 3, there are four social strategies that can be an optimal social strategy when $\varepsilon > 0$:

$$\sigma_1^1 = \begin{bmatrix} NS & S \\ S & S \end{bmatrix},\ \sigma_1^2 = \begin{bmatrix} S & S \\ NS & S \end{bmatrix},\ \sigma_1^3 = \begin{bmatrix} NS & S \\ NS & S \end{bmatrix},\ \sigma_1^4 = \begin{bmatrix} NS & NS \\ NS & NS \end{bmatrix}.$$

Illustration with L=1

- Proposition 4: Suppose that $0 < (1-\alpha)\varepsilon < 1/2$. Then

$$\sigma_1^* = \begin{cases} \sigma_1^1 & \text{if } 0 < \frac{c}{b} \le \frac{\beta(1-\alpha)^2(1-2\varepsilon)\varepsilon}{1+\beta(1-\alpha)^2(1-2\varepsilon)\varepsilon}, \\ \sigma_1^2 & \text{if } \frac{\beta(1-\alpha)^2(1-2\varepsilon)\varepsilon}{1+\beta(1-\alpha)^2(1-2\varepsilon)\varepsilon} < \frac{c}{b} \le \frac{\beta(1-\alpha)(1-2\varepsilon)[1-(1-\alpha)\varepsilon]}{1-\beta(1-\alpha)^2(1-2\varepsilon)\varepsilon}, \\ \sigma_1^3 & \text{if } \frac{\beta(1-\alpha)(1-2\varepsilon)[1-(1-\alpha)\varepsilon]}{1-\beta(1-\alpha)^2(1-2\varepsilon)\varepsilon} < \frac{c}{b} \le \beta(1-\alpha)(1-2\varepsilon), \\ \sigma_1^4 & \text{if } \beta(1-\alpha)(1-2\varepsilon) < \frac{c}{b} < 1. \end{cases}$$

Illustration with L=1

$$b = 10, \alpha = 0.1, \beta = 0.8, \varepsilon = 0.2$$

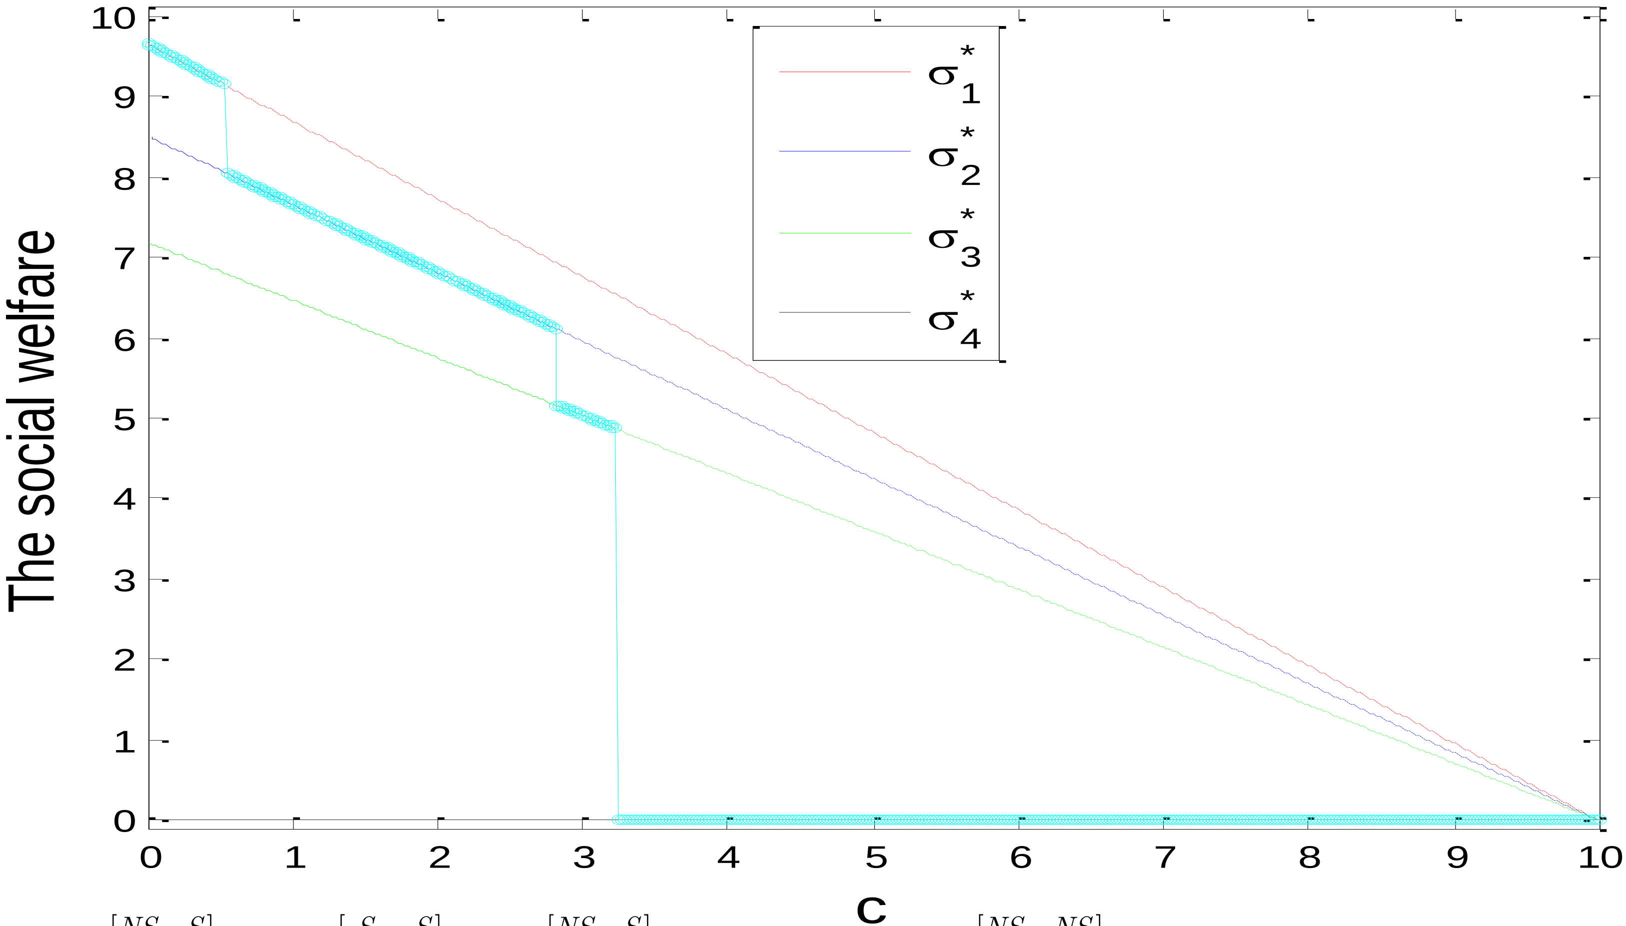

$\sigma_1^*:$ $\begin{bmatrix} NS & S \\ S & S \end{bmatrix}$ $\begin{bmatrix} S & S \\ NS & S \end{bmatrix}$ $\begin{bmatrix} NS & S \\ NS & S \end{bmatrix}$ $\begin{bmatrix} NS & NS \\ NS & NS \end{bmatrix}$

Illustration with L=1

$$b = 10, \alpha = 0.1, \beta = 0.8, \varepsilon = 0.2$$

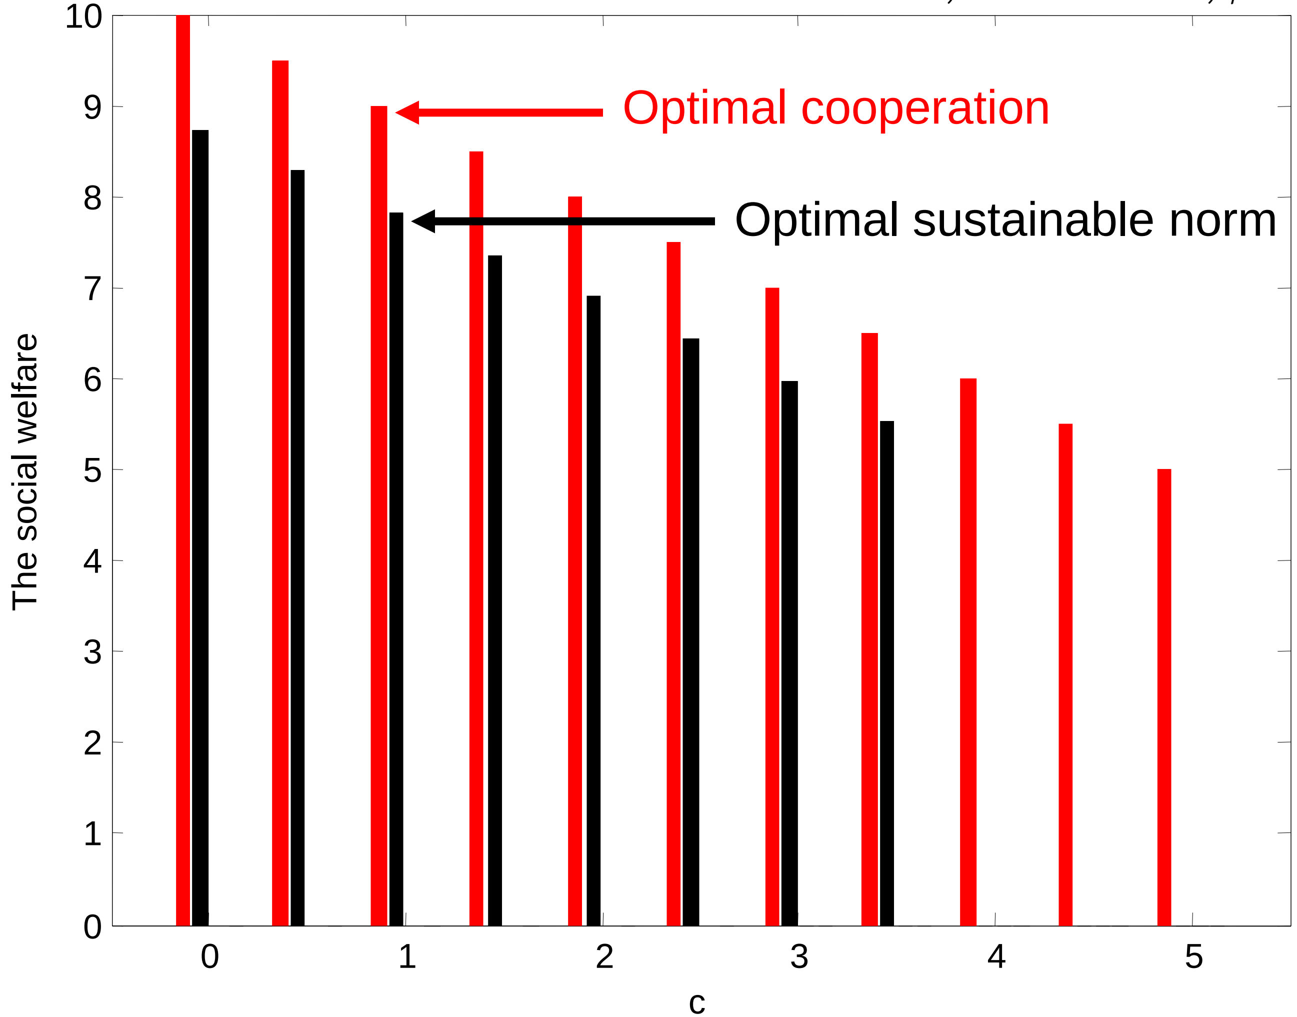

Network collapse

Illustration with L=1

$$b = 10, \alpha = 0.1, \beta = 0.8, \varepsilon = 0.4$$

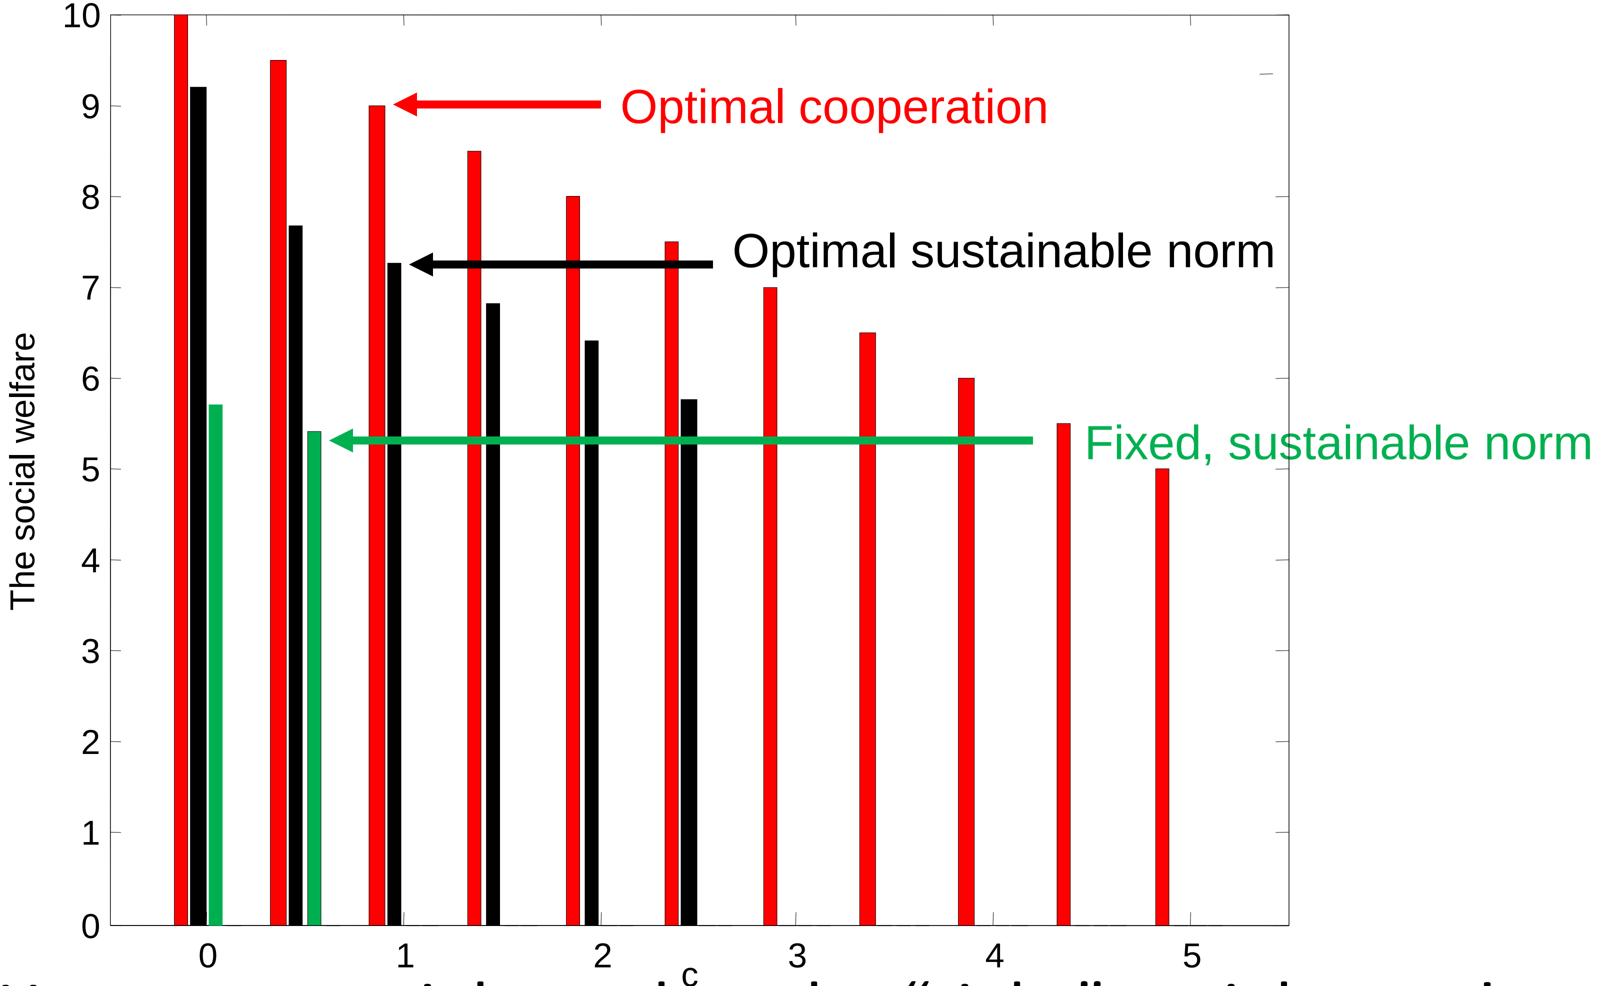

Hence, essential to select the "right" social norm!

Illustration with L=2

$$b = 10, \alpha = 0.1, \beta = 0.8, \varepsilon = 0.2$$

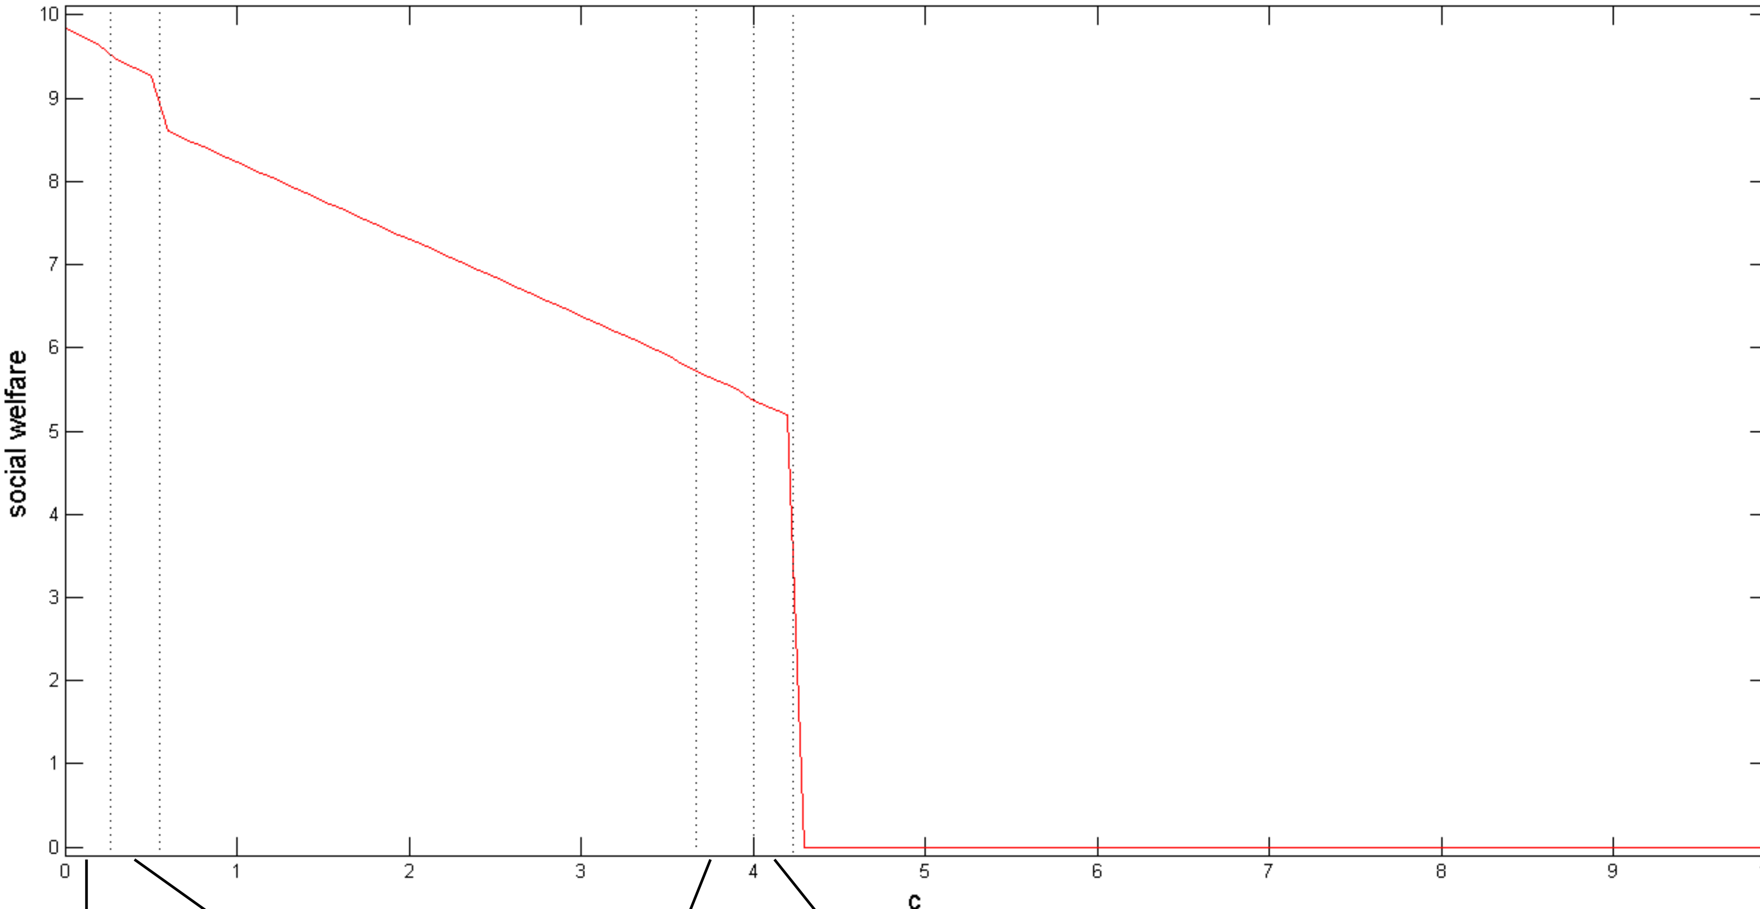

$$\sigma_2^* : \begin{bmatrix} S & S & S \\ NS & S & S \\ S & S & S \end{bmatrix} \begin{bmatrix} NS & S & S \\ NS & S & S \\ S & S & S \end{bmatrix} \quad \begin{bmatrix} S & S & S \\ S & S & S \\ NS & S & S \end{bmatrix} \quad \begin{bmatrix} S & S & S \\ NS & S & S \\ NS & S & S \end{bmatrix} \begin{bmatrix} NS & S & S \\ NS & S & S \\ NS & S & S \end{bmatrix} \quad \begin{bmatrix} NS & NS & NS \\ NS & NS & NS \\ NS & NS & NS \end{bmatrix}$$

Illustration with L=2

$$b = 10, \alpha = 0.1, \beta = 0.8, \varepsilon = 0.4$$

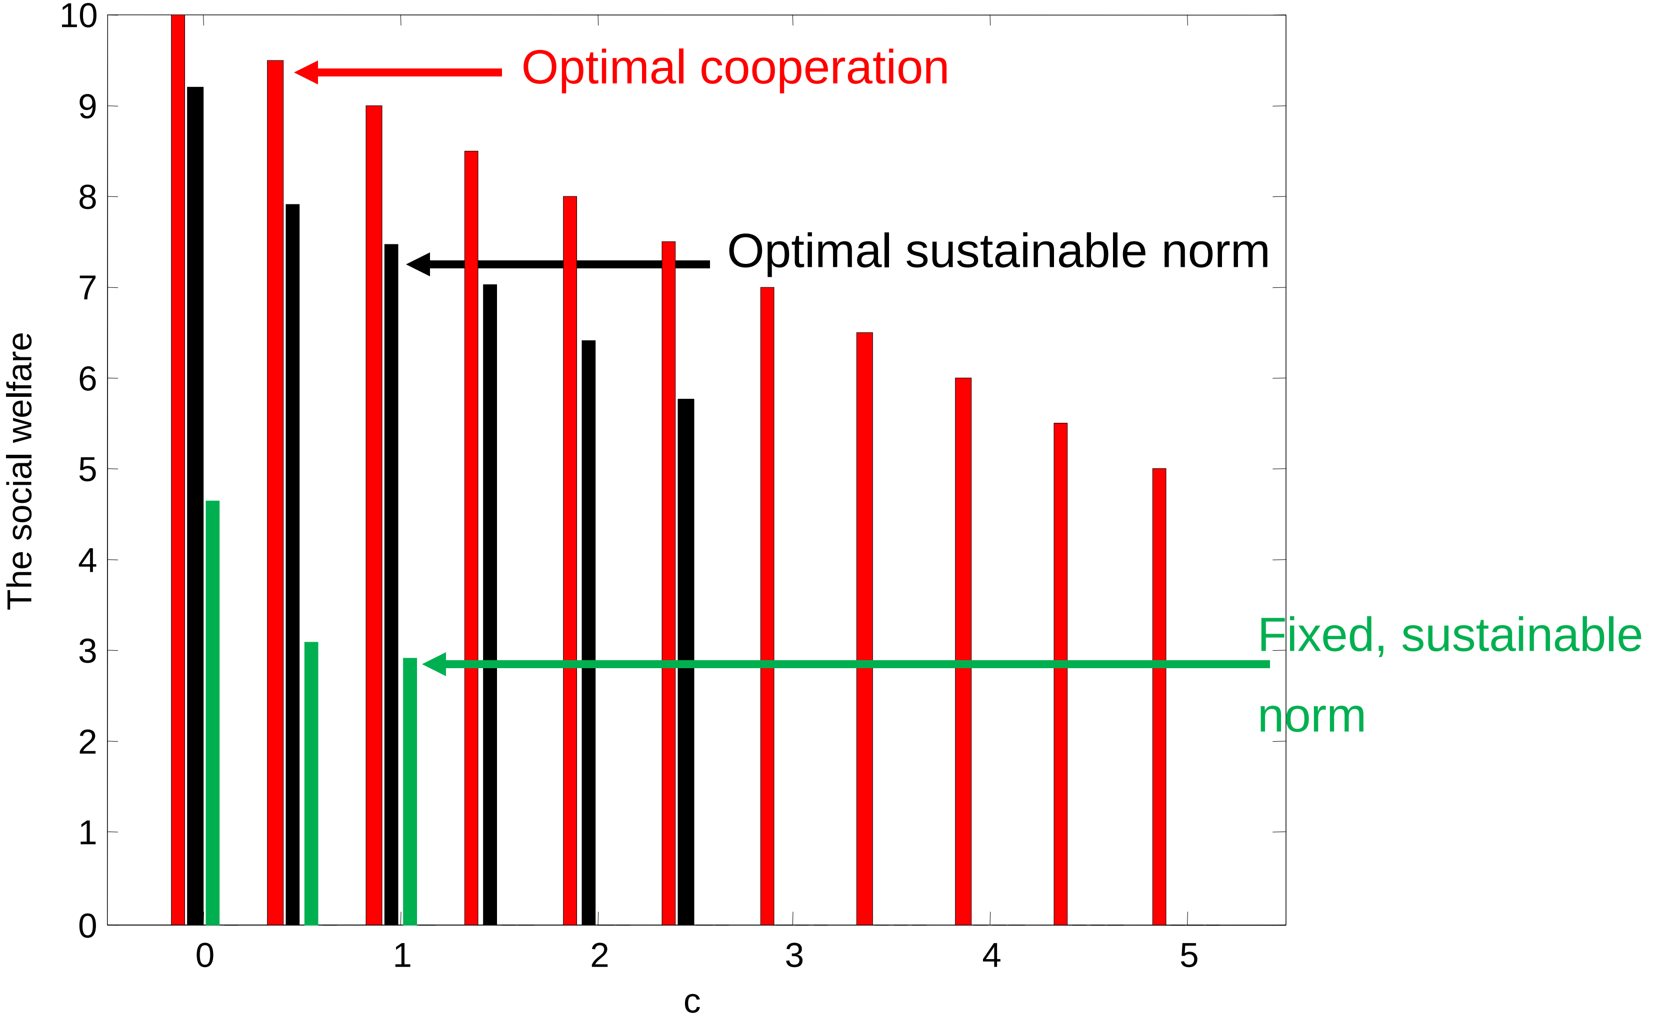

Reputation Schemes with Less Severe Punishment

- We generalize maximum punishment reputation schemes by considering reputation schemes described by

$$\tau(\theta, \tilde{\theta}, a_R) = \begin{cases} \min\{\theta + 1, L\} & \text{if } a_R = \sigma(\theta, \tilde{\theta}), \\ \max\{\theta - M, 0\} & \text{if } a_R \neq \sigma(\theta, \tilde{\theta}). \end{cases}$$

- For fixed (L, σ), increasing M has two counteracting effects when there are report errors:

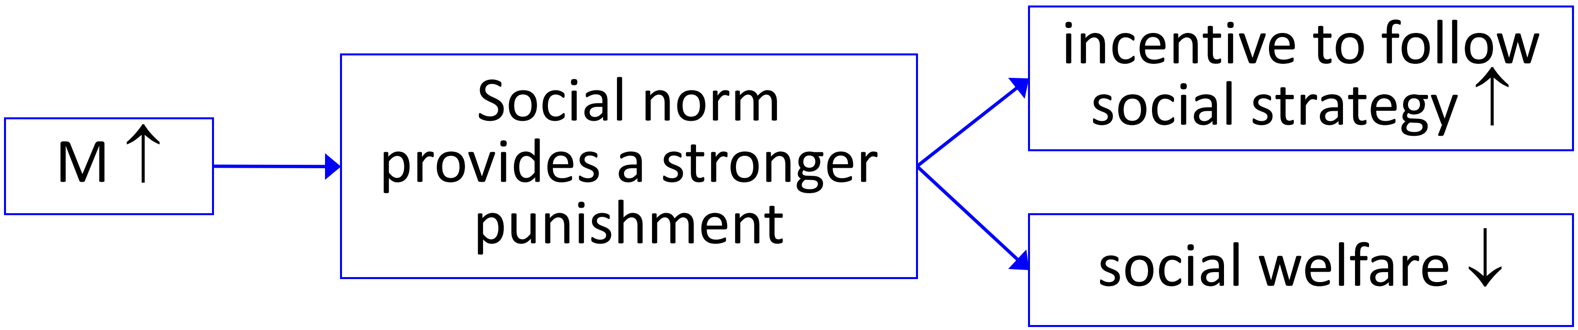

Illustration with Fixed L=3

$$b = 10, \alpha = 0.1, \beta = 0.8, \varepsilon = 0.3$$

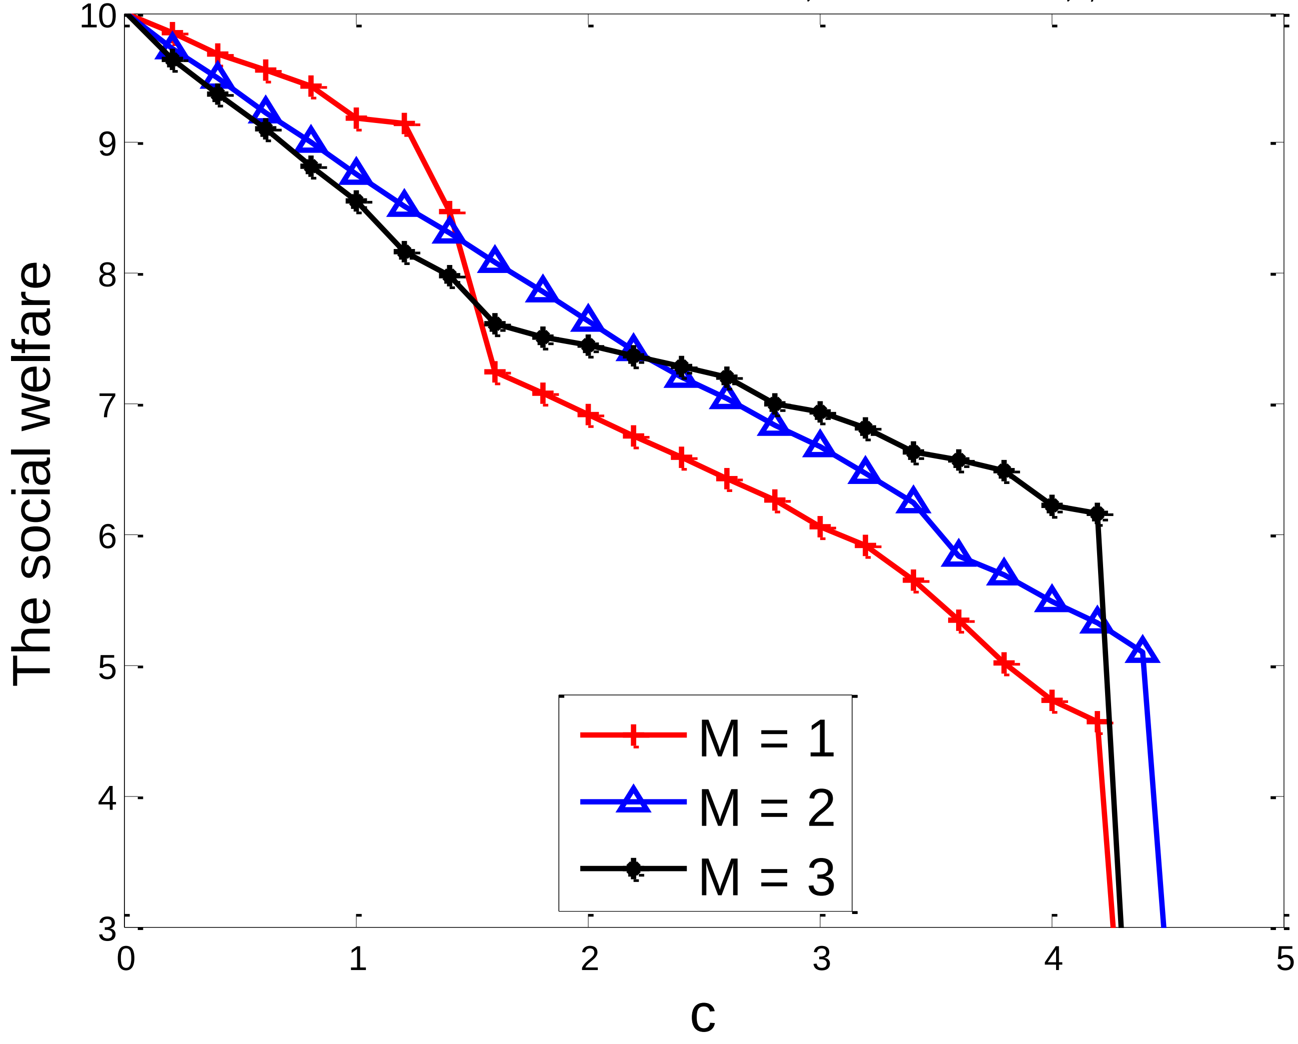

Whitewash-Proof Social Norms

- We now relax the restriction that the initial reputation is L and consider the case where a peer can whitewash its identity at the cost of $c_w > 0$ after observing its new reputation.
- Let $K \in \{0, \ldots, L\}$ be the initial reputation label.
- Then a social norm κ is whitewash-proof if and only if $v_\kappa^\infty(K) \leq v_\kappa^\infty(\theta) + c_w$ for all $\theta = 0, \ldots, L$.
- The protocol design problem is to choose a social norm (L, K, σ) that maximizes social welfare among sustainable and whitewash-proof social norms.

Whitewash-Proof Social Norms

- What about the initial reputation label K?

- For fixed L, increasing K has two counteracting effects:

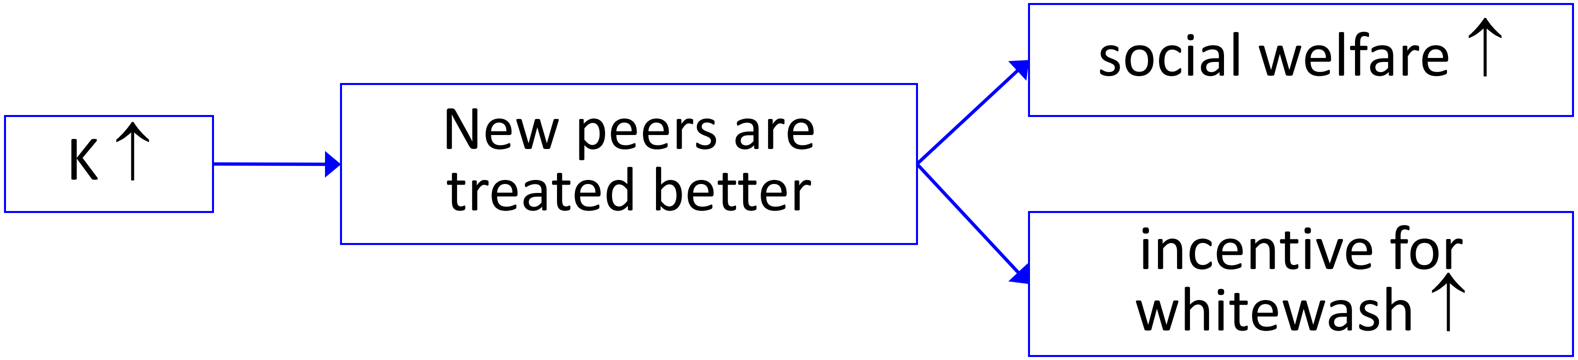

Planetlab performance analysis

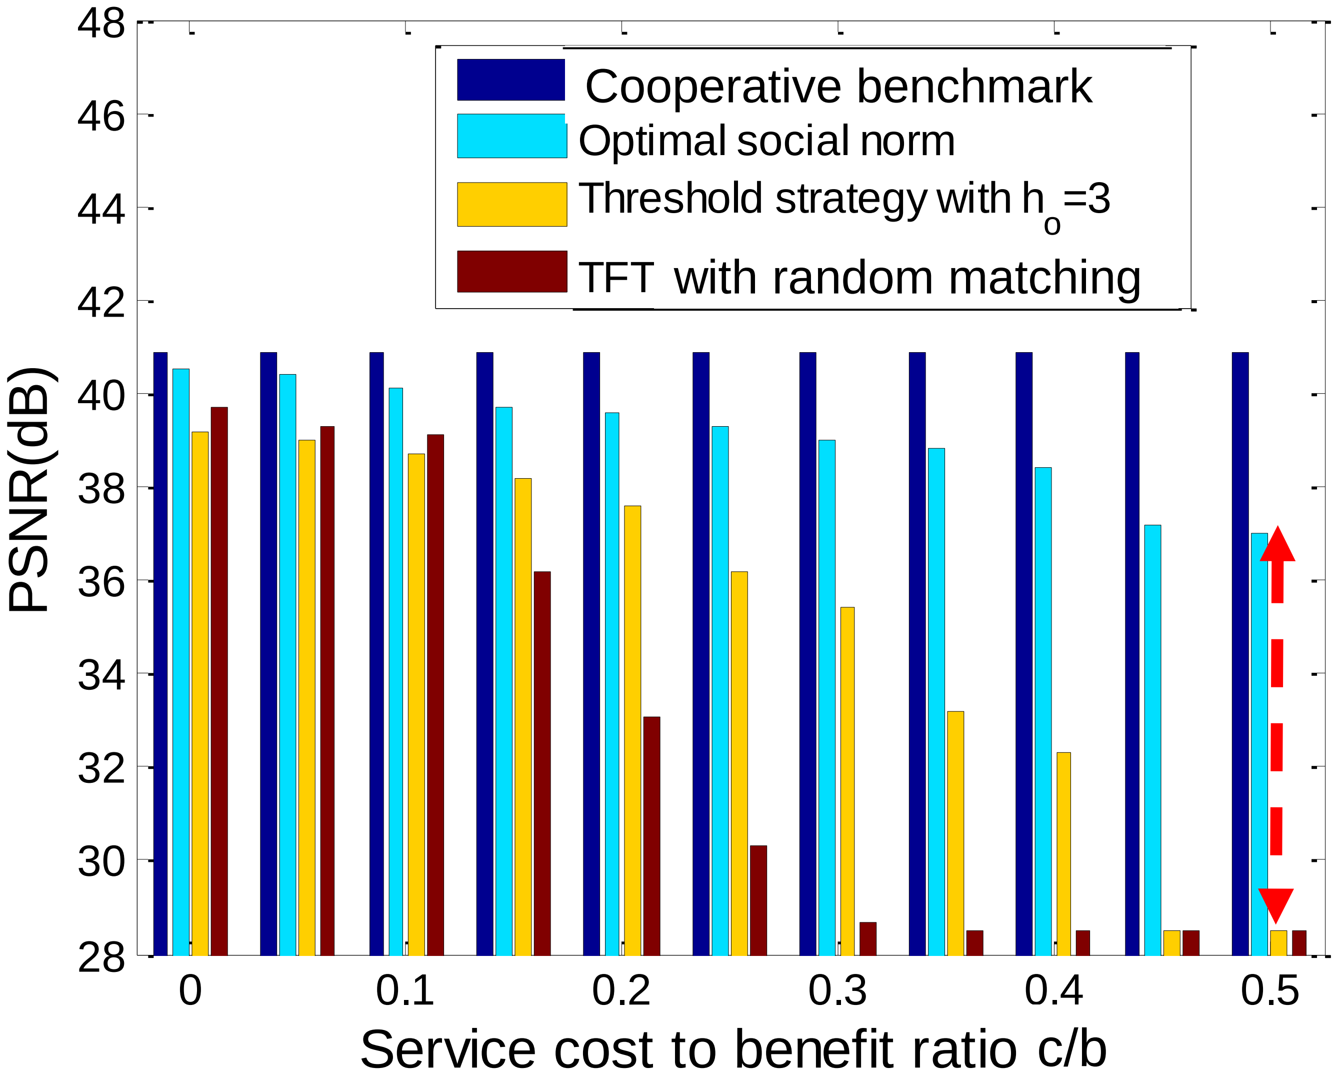

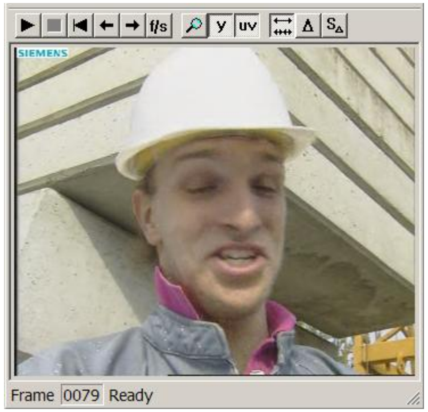

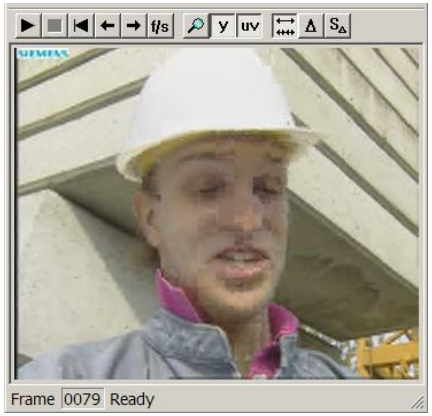

Extensions

- Network with helpers
  - Incentives to follow go down
  - Network collapse occurs sooner
  - Social welfare goes down
- Other types of altruism
- Learning
- Social norms vs. token-exchanges

Social Norms		Tokens
Central memory	Memory	No central memory (tokens as memory)
Reputation ↓	Punishments	
High	Informational requirements	Low
Does not limit effectiveness of design	Impatience	Limits effectiveness of design (nobody chooses to build a large treasury)
Initial reputation	Whitewashing	Initial endowment

Part II:

Design of Dynamic Personal Reciprocation Policies

- Hyunggon Park and Mihaela van der Schaar, “A Framework for Foresighted Resource Reciprocation in P2P Networks,” *IEEE Trans. Multimedia*, vol. 11, no. 1, pp. 101-116, Jan. 2009.
- Hyunggon Park and Mihaela van der Schaar, “Evolution of Resource Reciprocation Strategies in P2P Networks,” *IEEE Trans. Signal Process.*, vol. 58, no. 3, pp. 1205-1218, Mar. 2010.
- Rafit Izhak-Ratzin, Hyunggon Park and Mihaela van der Schaar, “Reinforcement Learning in BitTorrent Systems,” accepted *Infocom 2011.*

Part II: Dynamic P2P systems

- As before
  - Users interact repeatedly
  - Users are heterogeneous
  - Information is decentralized

- New
  - Choose ***partners*** and ***level*** of cooperation
  - Environment ***changing***

→ **No previous solutions for rigorously designing and evaluating protocols for P2P systems in dynamic environments**

Our approach – central issues

a) What reciprocation policy (protocol) to adopt while environment is known and stationary?

b) How to change the policy when environment changes?

A) Markov strategies – use Markov Decision Processes (MDPs) to determine policies

B) Online learning –reinforcement learning or model-based

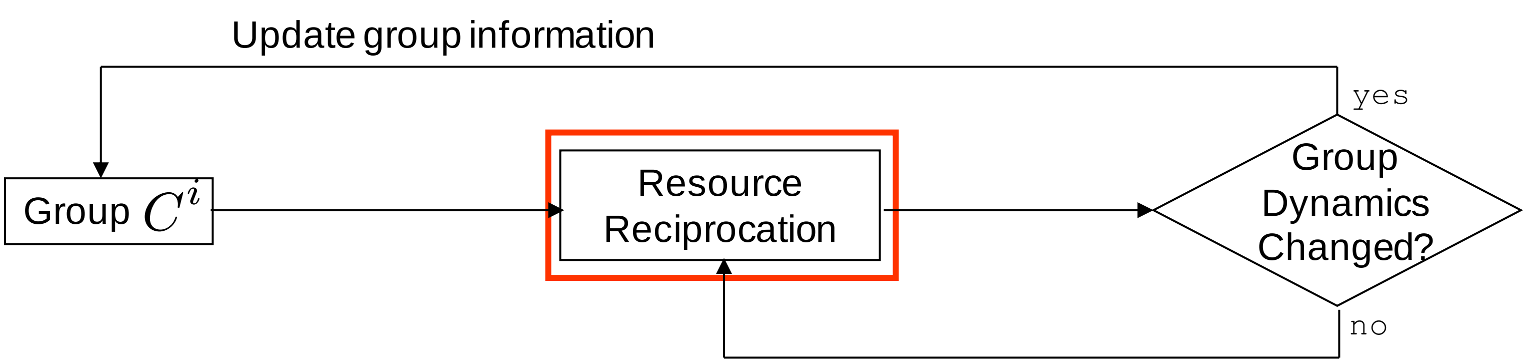

Resource Reciprocation

- A finite set of agents (peers)
- Actions: upload bandwidth allocations
- Policy: actions selected today are based on yesterday's reciprocation levels = states
- Utility: download rates, video quality, etc.
- Foresighted peers worry about long-term utility

State descriptions =>

Peers' intelligence

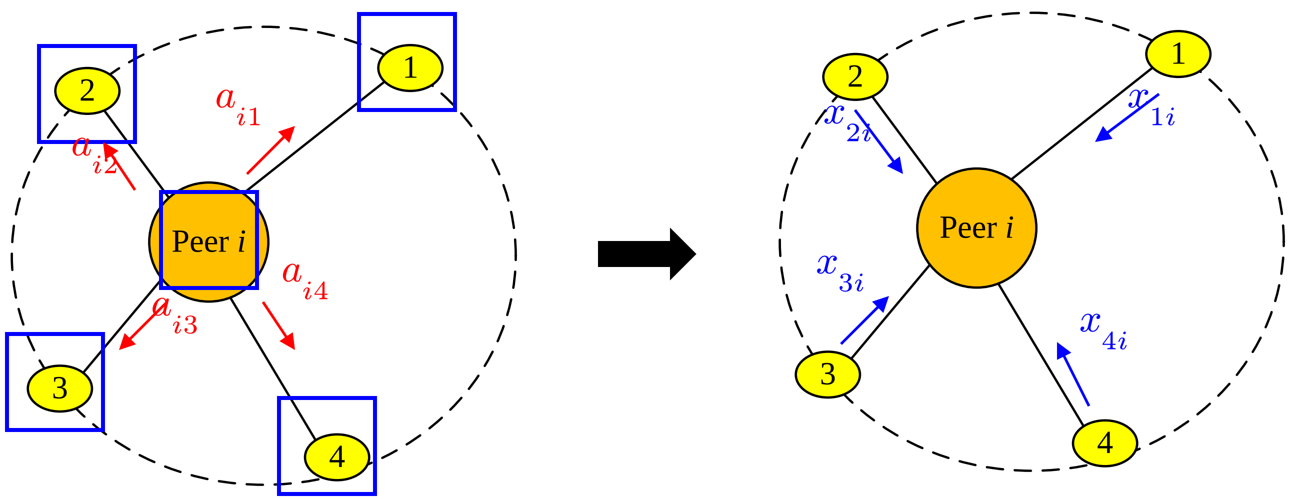

! Policy determines optimal *level* of cooperation, unlike "all or nothing" solution in BitTorrent (Tit-For-Tat)

Discrimination among peers - How?

- We prove assortative matching
  - Richer peers (=peers with higher bandwidth) match with richer peers
    - Generosity prompts generosity
  - Smarter peers (= peers with more refined states) match with smarter peers
    - Careful monitoring prompts careful monitoring
    - Better to cooperate with smarter peers than to steal from stupid peers ☺

Evolution of Mutual Resource Reciprocation

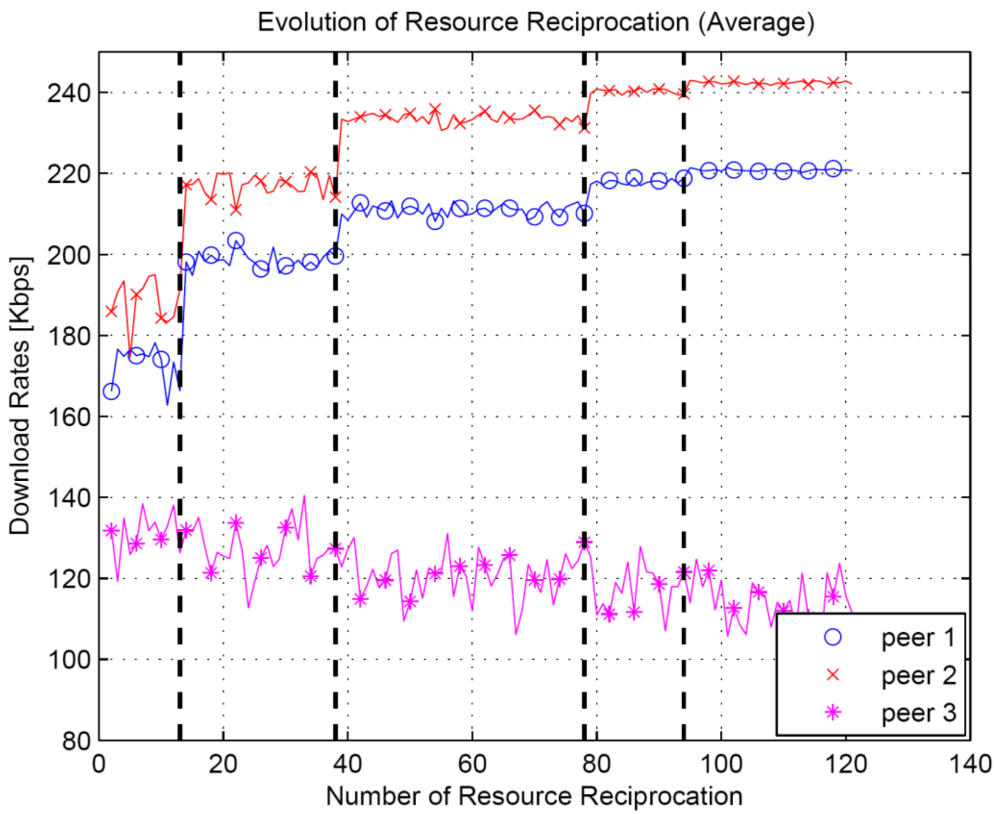

→Peer 1 and peer 2 improve their average download rates – Improvement is bounded by initial number of state description

→ Peer 3 is penalized

Clustering for Heterogeneous Peers

Different state refinement ability, same available bandwidth

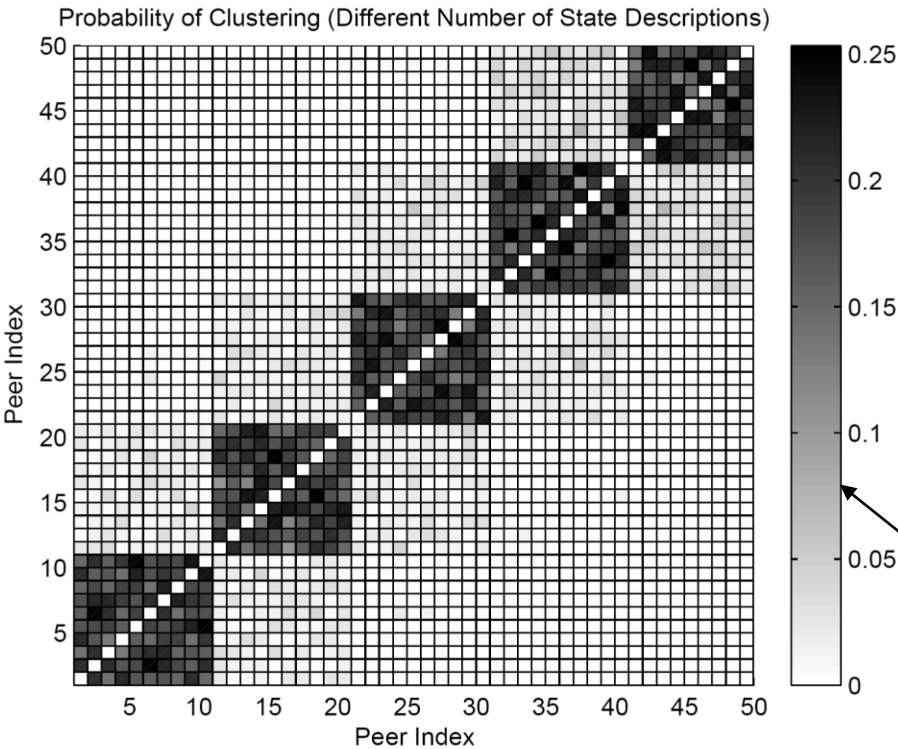

Probability

5 different classes: different ability to refine states

→ Peers prefer to form a group with peers having similar ability to refine states

Implementation and real-world experiments in Planetlab (Infocom 2011)

Part III: Community Formation

Information production, sharing and consumption and link formation in networked communities

- Jaeok Park and Mihaela van der Schaar, “A Game Theoretic Analysis of Incentives in Content Production and Sharing over Peer-to-Peer Networks,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 4, pp. 704-717, August 2010.
- Jaeok Park and Mihaela van der Schaar, “Content Pricing in Peer-to-Peer Networks,” *NetEcon '10.*
- Jaeok Park and Mihaela van der Schaar, “Pricing and Incentives in Peer-to-Peer Networks,” *INFOCOM 2010.*

Current EE/CS/Econ Literature		Our research
Fixed	<-------- Who produces? -------->	Choice
Fixed	<-------- What/how much is produced? -------->	Choice
Fixed	<-------- What/how much is shared? -------->	Choice
Fixed	<-------- Who connects to whom? -------->	Choice

Challenges

Introduction

- In today's Internet - the emergence of user-generated content in the form of videos, multi-modal information, customer reviews, etc.
- P2P networks can offer a useful platform for sharing user-generated content, but the free-riding problem may hinder the efficient utilization of P2P networks.
- Our contributions
  - Propose a model of content production and sharing, and use content pricing to overcome the free-riding problem.
  - We also consider the problem of network topology design.

Existing Work

- Existing Work on Pricing in P2P Networks:
  - Golle, Leyton-Brown, Mironov, and Lillibridge (2001) construct a game theoretic model and propose a micro-payment mechanism to provide an incentive for sharing.
  - Antoniadis, Courcoubetis, and Mason (2004) compare different pricing schemes and their informational requirements in the context of a simple file-sharing game.
  - Adler, Kumar, Ross, Rubenstein, Turner, and Yao (2004) investigate the problem of selecting multiple server peers given the prices of service and a budget constraint.
- Limitations: The models of the above papers capture only a partial picture of a content production and sharing scenario.
  - No production decisions, no explanation for how prices and budgets are determined, etc.

Model

- We consider a P2P network consisting of N peers, which produce content using their own production technologies and distribute produced content using the P2P network.
- $\mathcal{N} \triangleq \{1, \ldots, N\}$: set of peers in the P2P network
- $D(i)$: set of peers that peer *i* can download from
- $U(i)$: set of peers that peer *i* can upload to
- We model content production and sharing in the P2P network as a three-stage sequential game, called the content production and sharing (CPS) game.

Description of the CPS Game

- *Stage One (Production):* Each peer determines its level of production. $x_i \in \mathbb{R}_+$ represents the amount of content produced by peer *i* and is known only to peer *i.*
- *Stage Two (Sharing):* Each peer specifies its level of sharing. $y_i \in [0, x_i]$ represents the amount of content that peer *i* makes available to other peers. Peer *i* observes $(y_j)_{j \in D(i)}$ at the end of stage two.
- *Stage Three (Transfer):* Each peer determines the amount of content that it downloads from other peers. Peer *i* serves all the requests it receives from any other peer in $U(i)$ up to y_i. $z_{ij} \in [0, y_j]$ represents the amount of content that peer *i* downloads from peer $j \in D(i)$, or equivalently peer *j* uploads for peer *i*.

Allocation and Payoff

- An allocation of the CPS game is represented by $(\mathbf{x}, \mathbf{y}, \mathbf{Z})$, where $\mathbf{x} \triangleq (x_1, \ldots, x_N)$, $\mathbf{y} \triangleq (y_1, \ldots, y_N)$, $\mathbf{z}_i \triangleq (z_{ij})_{j \in D(i)}$, for each $i \in \mathcal{N}$, and $\mathbf{Z} \triangleq (\mathbf{z}_1, \ldots, \mathbf{z}_N)$.

- The payoff function of peer *i* in the CPS game is

$$v_i(\mathbf{x}, \mathbf{y}, \mathbf{Z}) = \underbrace{f_i(x_i, \mathbf{z}_i)}_{\substack{\text{utility from}\\ \text{consumption}\\ \text{(diff., concave)}}} - \underbrace{k_i(x_i)}_{\substack{\text{production}\\ \text{cost}\\ \text{(diff., convex)}}} - \underbrace{\sum_{j \in D(i)} \delta_{ij} z_{ij}}_{\substack{\text{download}\\ \text{cost}}} - \underbrace{\sum_{j \in U(i)} \sigma_{ji} z_{ji}}_{\substack{\text{upload}\\ \text{cost}}}.$$

Nash Equilibrium

- A strategy for peer *i* in the CPS game is its complete contingent plan over the three stages, which can be represented by $(x_i, y_i(x_i), \mathbf{z}_i(x_i, y_i, (y_j)_{j\in D(i)}))$.
- Nash equilibrium (NE) is defined as a strategy profile such that no peer can improve its payoff by a unilateral deviation
- NE describes the outcome when peers behave selfishly.

Nash Equilibrium

- Proposition: Suppose that, for each $i \in \mathcal{N}$, a solution to $\max_{x \geq 0}\{f_i(x,0) - k_i(x)\}$ exists, and denote it as x_i^e. Unique NE outcome of the CPS game has $x_i = x_i^e$ and $z_{ij} = 0$ for all $j \in D(i)$, for all $i \in \mathcal{N}$.
- Idea of Proof: If $z_{ij} > 0$ for some $i \in \mathcal{N}$ and $j \in D(i)$, peer j can increase its payoff by deviating to $y_j = 0$. Therefore, $z_{ij} = 0$ for all $i \in \mathcal{N}$ and $j \in D(i)$ at any NE outcome.
- Given that there is no transfer of content, peers choose an autarkic optimal level of production.

Pricing as an Incentive Scheme

- We introduce a pricing scheme in the CPS game as a solution to overcome the network collapse.
- p_{ij}: unit price of content that peer *j* provides to peer *i*.
- A pricing scheme can be represented by $\mathbf{p} \triangleq (p_{ij})_{i \in \mathcal{N}, j \in D(i)}$.
- The payoff function of peer *i* in the CPS game with pricing scheme $\mathbf{p}$ is given by

$$\pi_i(\mathbf{x}, \mathbf{y}, \mathbf{Z}; \mathbf{p}) = v_i(\mathbf{x}, \mathbf{y}, \mathbf{Z}) - \sum_{j \in D(i)} p_{ij} z_{ij} + \sum_{j \in U(i)} p_{ji} z_{ji}$$

$$v_i(\mathbf{x}, \mathbf{y}, \mathbf{Z}) = f_i(x_i, \mathbf{z}_i) - k_i(x_i) - \sum_{j \in D(i)} \delta_{ij} z_{ij} - \sum_{j \in U(i)} \sigma_{ji} z_{ji}.$$

What can we accomplish with prices?

- We measure social welfare by the sum of the payoffs of peers, $\sum_{i=1}^{N} v_i(\mathbf{x}, \mathbf{y}, \mathbf{Z})$.

- A socially optimal (SO) allocation is an allocation that maximizes social welfare among feasible allocations.

- Using Karush-Kuhn-Tucker (KKT) conditions, we can characterize SO allocations.

Content Pricing

- Proposition: Let $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{Z}^*)$ be an SO allocation. There is a pricing scheme $\mathbf{p}^* = (p_{ij}^*)_{i \in \mathcal{N}, j \in D(i)}$ such that $(\mathbf{x}^*, \mathbf{y}^*, \mathbf{Z}^*)$ is a NE outcome of the CPS game with this pricing scheme.
- Prices

  In the expression $p_{ij}^* = \lambda_{ij} + \sigma_{ij}$, we can see that peer *i* compensates peer *j* for the upload cost, σ_{ij}, as well as the shadow price, λ_{ij}, of content supplied from peer *j* to peer *i*

Idea: We can construct an NE strategy such that the KKT conditions for the NE outcome with the proposed pricing scheme coincide with those for social optimum.

- Pricing schemes can always get us to the social optimum
- Important questions:
  - Who produces and how much?
  - Who shares with whom?
- The answers to these questions depend sensitively on the topologies
- Network topology may be a *design* parameter

* Next: We work out some surprising examples

Assumptions

- (Perfectly substitutable content) The utility from consumption depends only on the total amount of content.
- (Linear production cost) The production cost is linear in the amount of content produced.
- (Socially valuable P2P network) Obtaining a unit of content through the P2P network costs less to peers than producing it privately.

Production costs

- β_i is the per capita cost of peer *i* producing one unit of content and supplying it to every other peer to which peer *i* can upload, and we call it the cost parameter of peer *i*.

$$\beta_i \triangleq [\kappa_i + \sum_{j \in D(i)} (\delta_{ji} + \sigma_{ji})] / (1 + |D(i)|)$$

$$\beta \triangleq \min\{\beta_1, \ldots, \beta_N\}$$

Focus on efficient production

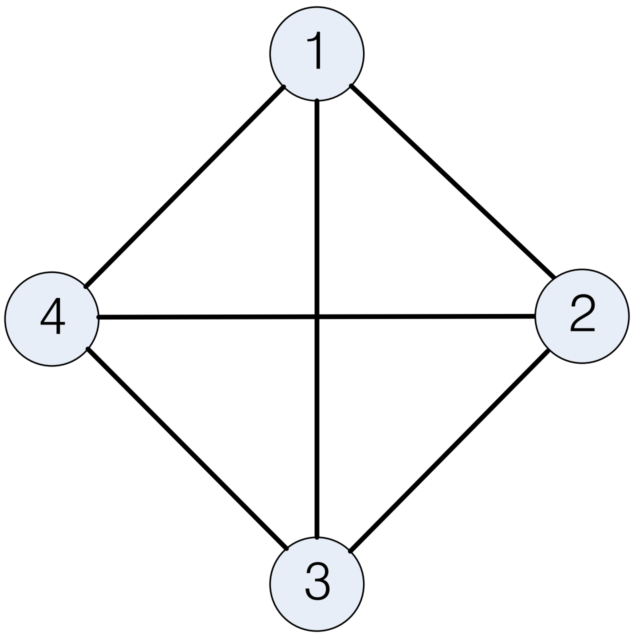

- In a fully connected P2P network, we have $D(i) = U(i) = \mathcal{N} \setminus \{i\}$ for all $i \in \mathcal{N}$.
- It is SO to have only the most "cost-efficient" peers (i.e., peers with the smallest cost parameter in the network) produce a positive amount, where the total amount of production is given by $\hat{x}_{\beta}$.

Focus on efficient production

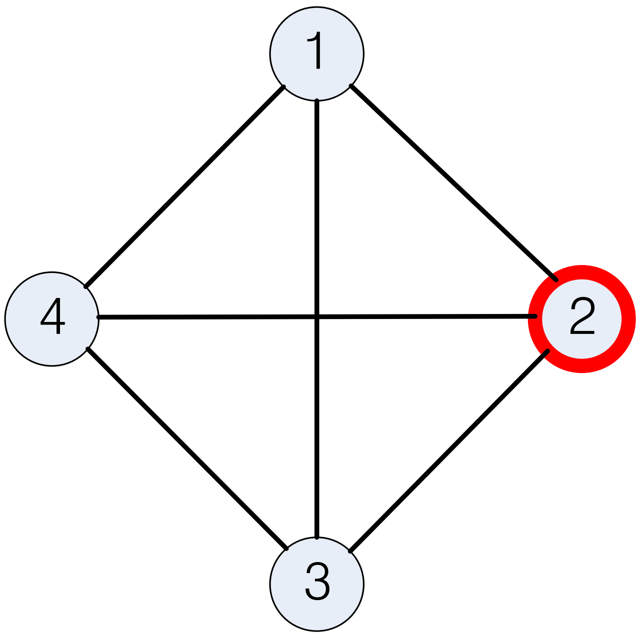

- In a fully connected P2P network, we have $D(i) = U(i) = \mathcal{N} \setminus \{i\}$ for all $i \in \mathcal{N}$.
- It is SO to have only the most “cost-efficient” peers (i.e., peers with the smallest cost parameter in the network) produce a positive amount, where the total amount of production is given by $\hat{x}_{\beta}$.

Focus on efficient production

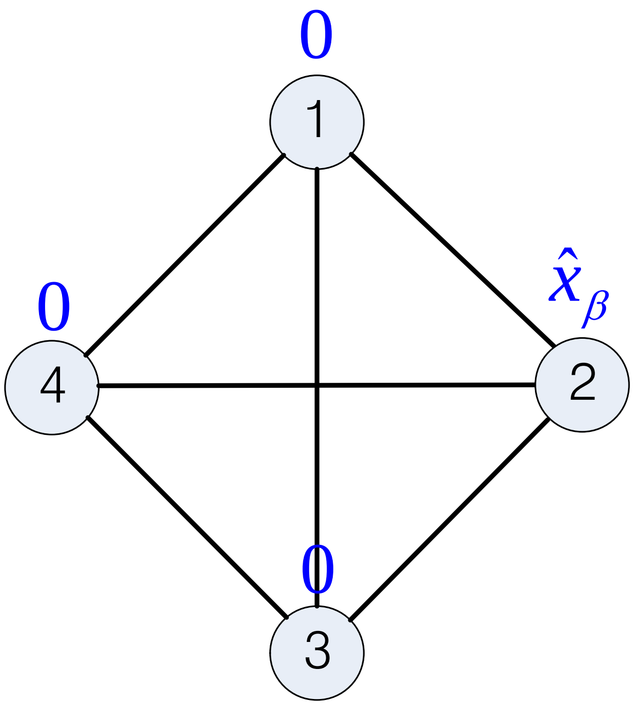

- In a fully connected P2P network, we have $D(i) = U(i) = \mathcal{N} \setminus \{i\}$ for all $i \in \mathcal{N}$.
- It is SO to have only the most "cost-efficient" peers (i.e., peers with the smallest cost parameter in the network) produce a positive amount, where the total amount of production is given by $\hat{x}_\beta$.

Focus on efficient production

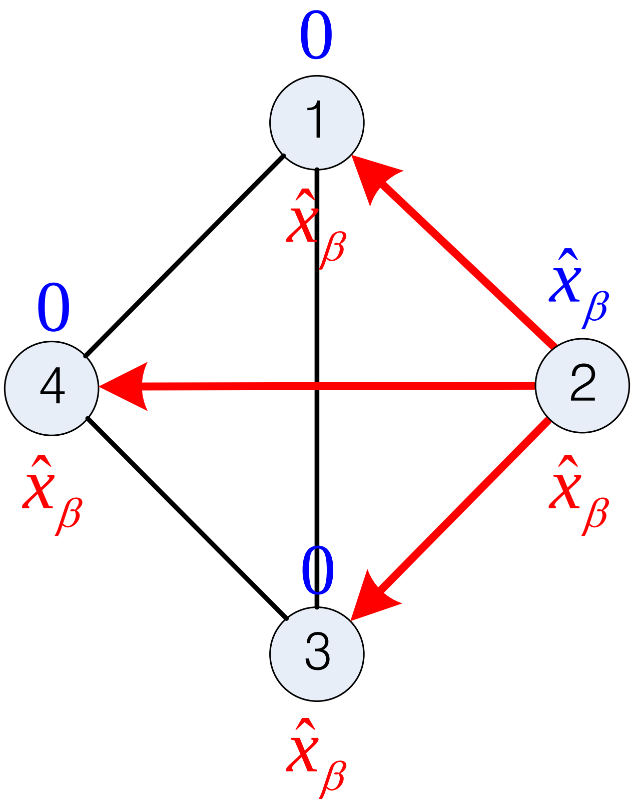

- In a fully connected P2P network, we have $D(i) = U(i) = \mathcal{N} \setminus \{i\}$ for all $i \in \mathcal{N}$.
- It is SO to have only the most "cost-efficient" peers (i.e., peers with the smallest cost parameter in the network) produce a positive amount, where the total amount of production is given by $\hat{x}_\beta$.
- The most efficient producer does all the producing.

Focus on efficient production

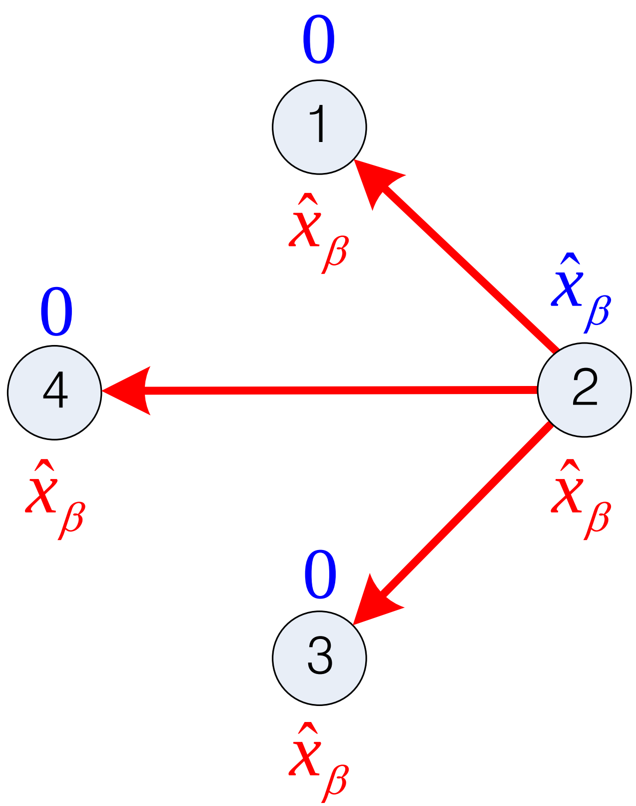

- Suppose that the network designer can choose connectivity between peers, assuming that the optimal pricing scheme is implemented given the connectivity topology.
- Then the star topology with the most cost-efficient peer as the center will achieve the maximum social welfare with the minimum number of links.

Focus on impact of topology

- We consider homogeneous peers – same benefit functions and cost parameters.
- We consider three stylized network topologies: a star topology, a ring topology, and a line topology.

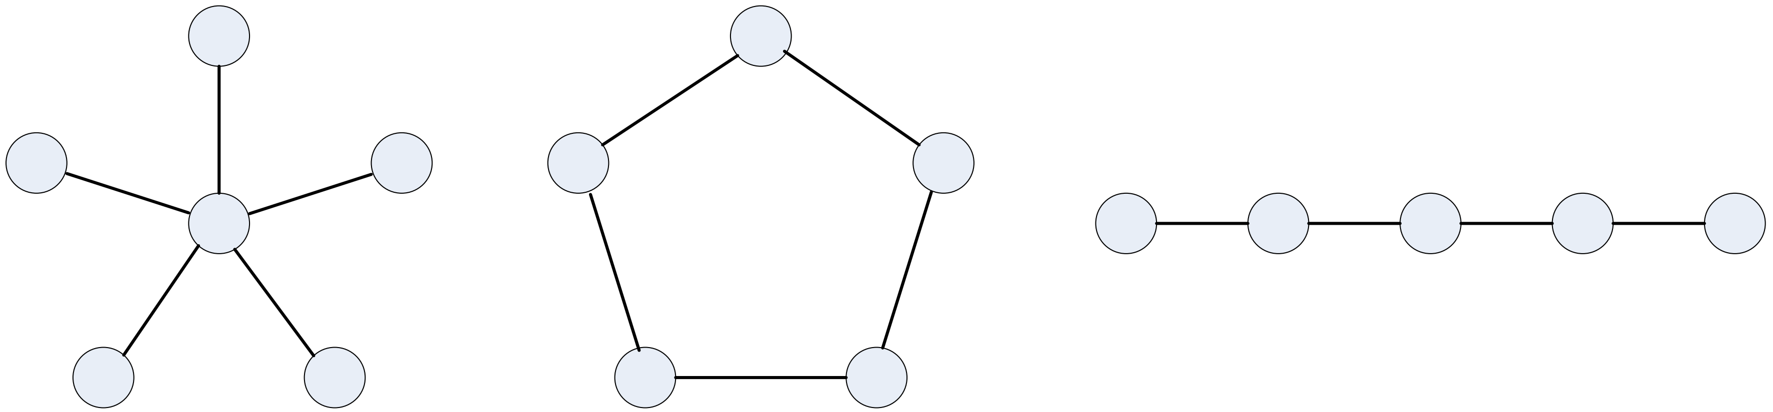

Star Topology

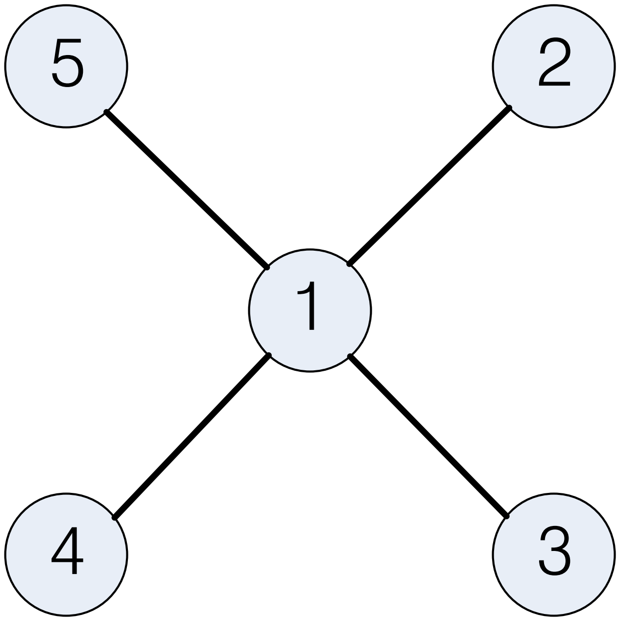

- $\beta_1 = [\kappa + (N-1)(\delta + \sigma)]/N = \beta$ and $\beta_j = (\kappa + \delta + \sigma)/2$ for $j \neq 1$.
- Since peer 1 is more connected than other peers, it is more cost-efficient (i.e., $\beta_1 < \beta_j$ for all $j \neq 1$).

Star Topology

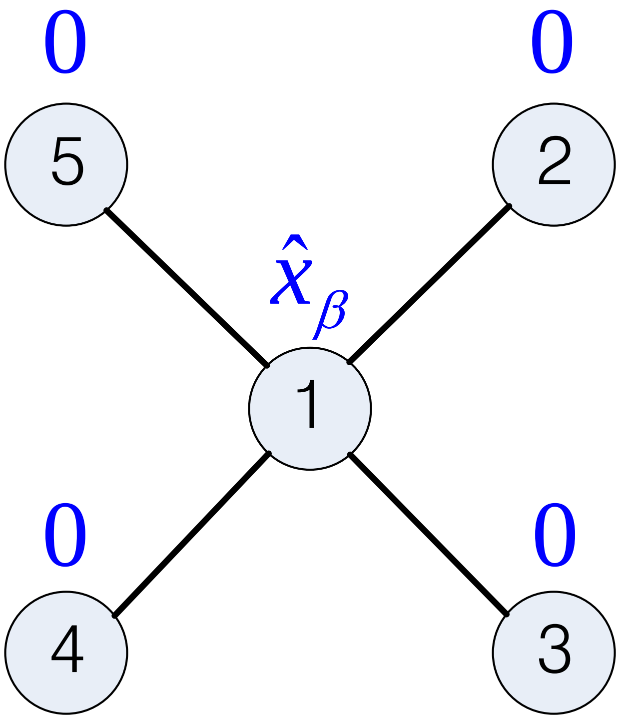

- $\beta_1 = [\kappa + (N-1)(\delta + \sigma)]/N = \beta$ and $\beta_j = (\kappa + \delta + \sigma)/2$ for $j \neq 1$.
- Since peer 1 is more connected than other peers, it is more cost-efficient (i.e., $\beta_1 < \beta_j$ for all $j \neq 1$).
- Only peer 1 produces a positive amount of content $\hat{x}_\beta$ and uploads it to every other peer at the SO allocation.

Star Topology

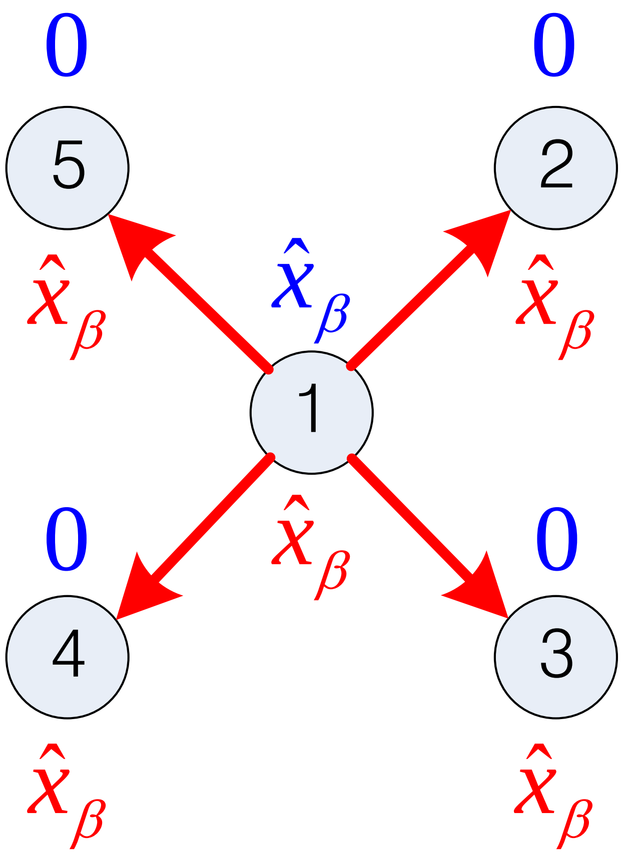

- $\beta_1 = [\kappa + (N-1)(\delta+\sigma)]/N = \beta$ and $\beta_j = (\kappa+\delta+\sigma)/2$ for $j \neq 1$.
- Since peer 1 is more connected than other peers, it is more cost-efficient (i.e., $\beta_1 < \beta_j$ for all $j \neq 1$).
- Only peer 1 produces a positive amount of content $\hat{x}_\beta$ and uploads it to every other peer at the SO allocation.
- The most efficient shipper does all the producing.

Ring Topology

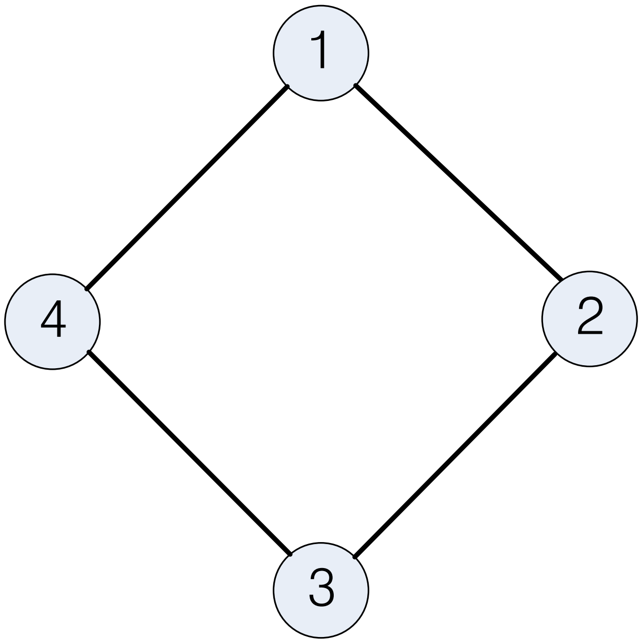

- Every peer is connected to two neighboring peers, and thus peers have the same cost parameter $\tilde{\beta} \triangleq [\kappa + 2(\delta + \sigma)] / 3$.

Ring Topology

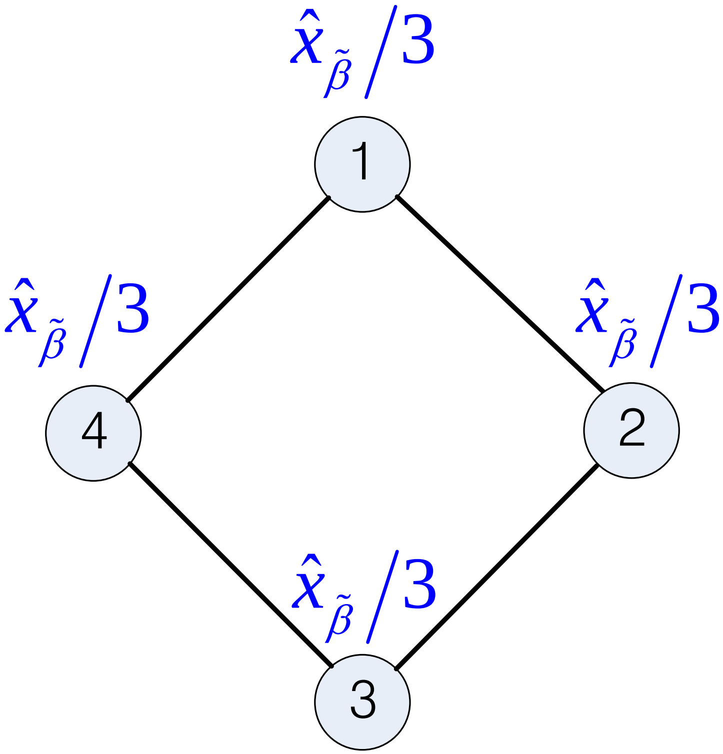

- Every peer is connected to two neighboring peers, and thus peers have the same cost parameter $\tilde{\beta} \triangleq [\kappa + 2(\delta + \sigma)]/3$.
- Each peer produces the amount $\hat{x}_{\tilde{\beta}}/3$ while consuming $\hat{x}_{\tilde{\beta}}$ at the SO allocation

Ring Topology

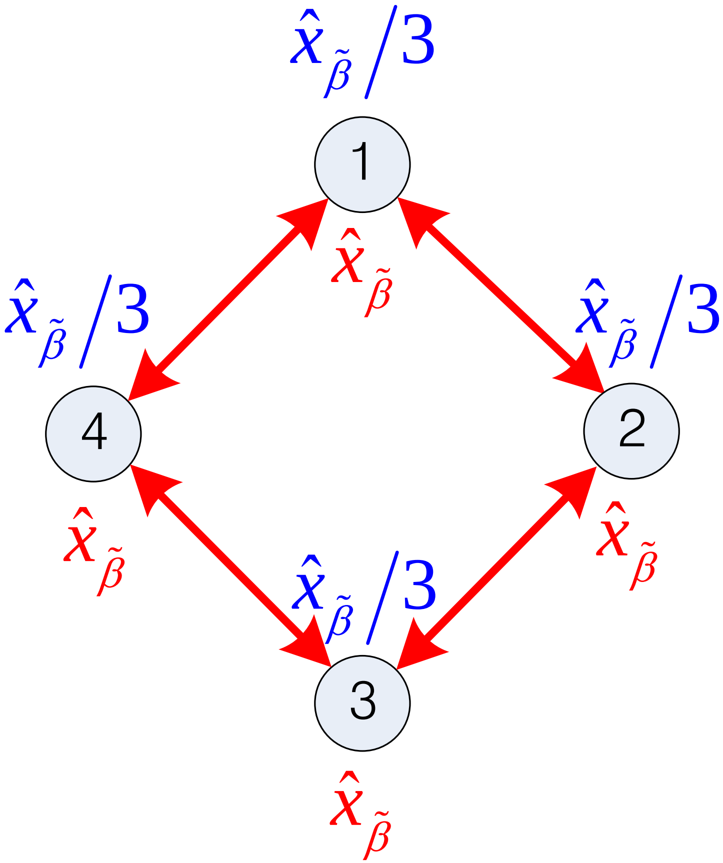

- Every peer is connected to two neighboring peers, and thus peers have the same cost parameter $\tilde{\beta} \triangleq [\kappa + 2(\delta + \sigma)]/3$.
- Each peer produces the amount $\hat{x}_{\tilde{\beta}}/3$ while consuming $\hat{x}_{\tilde{\beta}}$ at the SO allocation
- Interesting: The SO amounts of production and consumption and the maximum per capita social welfare are independent of N.

Line Topology

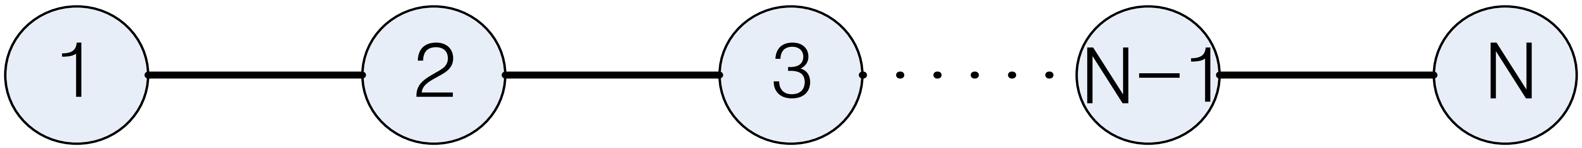

- $\beta_1 = \beta_N = (\kappa + \delta + \sigma)/2$ and $\beta_i = \tilde{\beta}$ for all $i \neq 1, N$.
- Peers at the ends (peers 1 and N) have high shipping costs, hence do not ship, hence do not produce
- Behavior of peers in the middle will turn out to depend on N.
- The level and structure of SO productions and consumptions depends on N and on locations.
- Prices depend both on N and on locations.

Line Topology

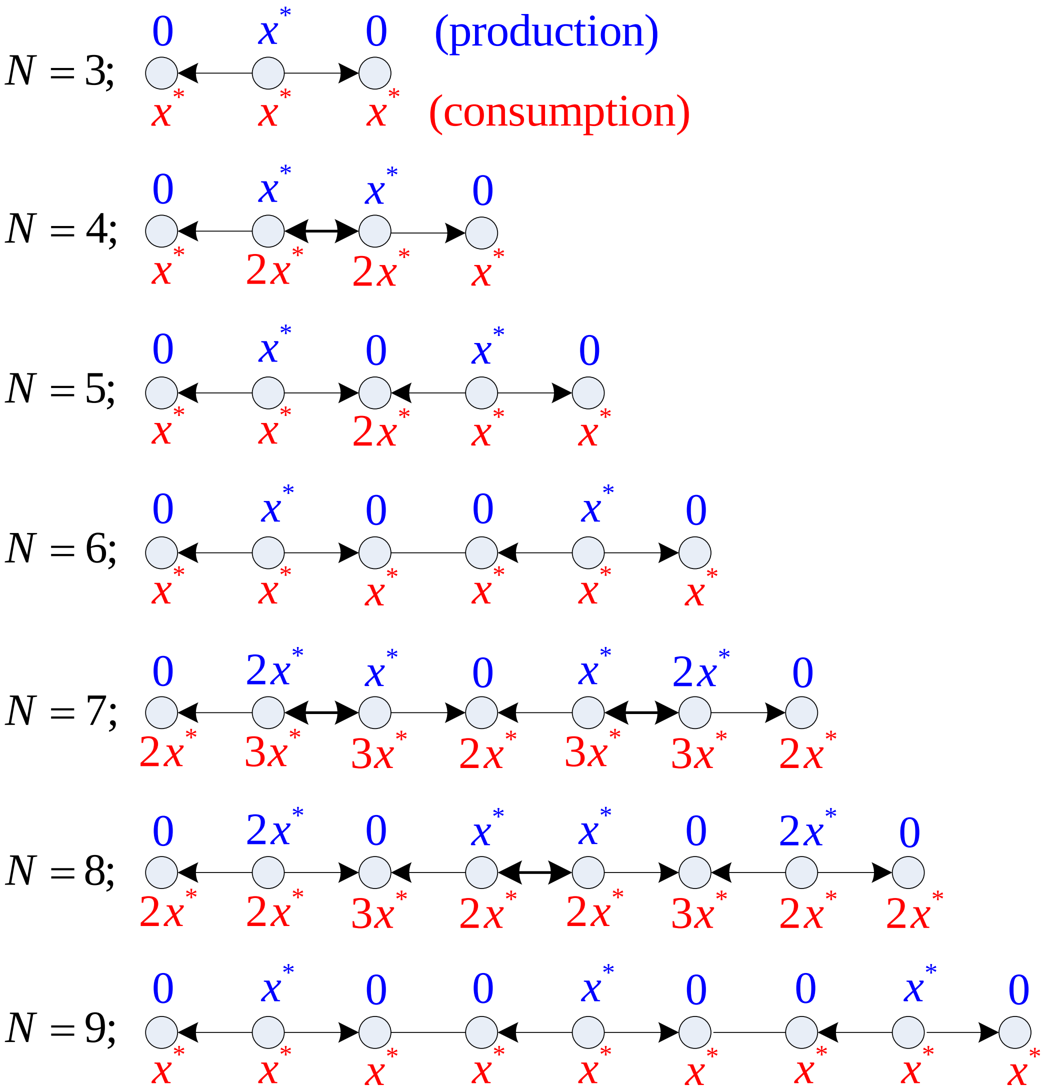

Research agenda: Next generation designs

Current and ongoing work in our research lab

- Content production and sharing in distributed systems
- Social norms/Pricing/Direct reciprocity for social systems
- Intervention mechanisms (threats/punishments)
- Strategic security in social systems
- Network formation/growth
- New classes of system/networking games
- Stochastic games for dynamic multimedia systems
- Conjectural equilibrium and conjecture-based learning
- Multi-agent learning in social systems

Our research: http://medianetlab.ee.ucla.edu