

# Sensing And Decision Making In Social Networks

Vikram Krishnamurthy  
Cornell University  
Electrical & Computer Engineering

from a statistical signal processing/stochastic  
control viewpoint



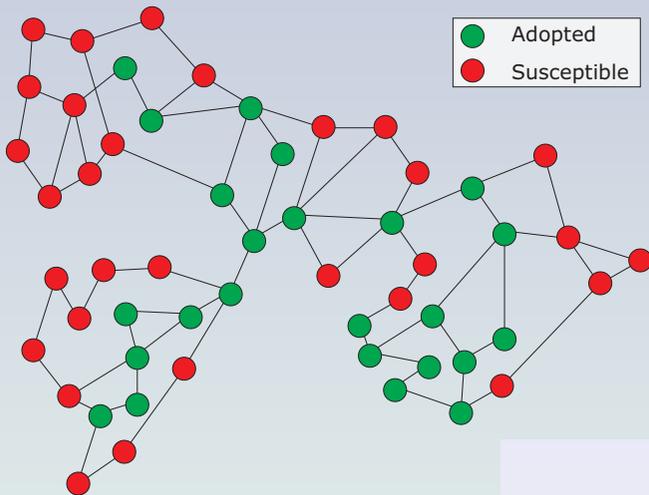
# SOCIAL (HUMAN) SENSORS

- Social sensor: Provides information about environment to a social network after interaction with other agents.
- Examples: Twitter posts, Facebook status updates, ratings on online reputation systems (Yelp, Tripadvisor)

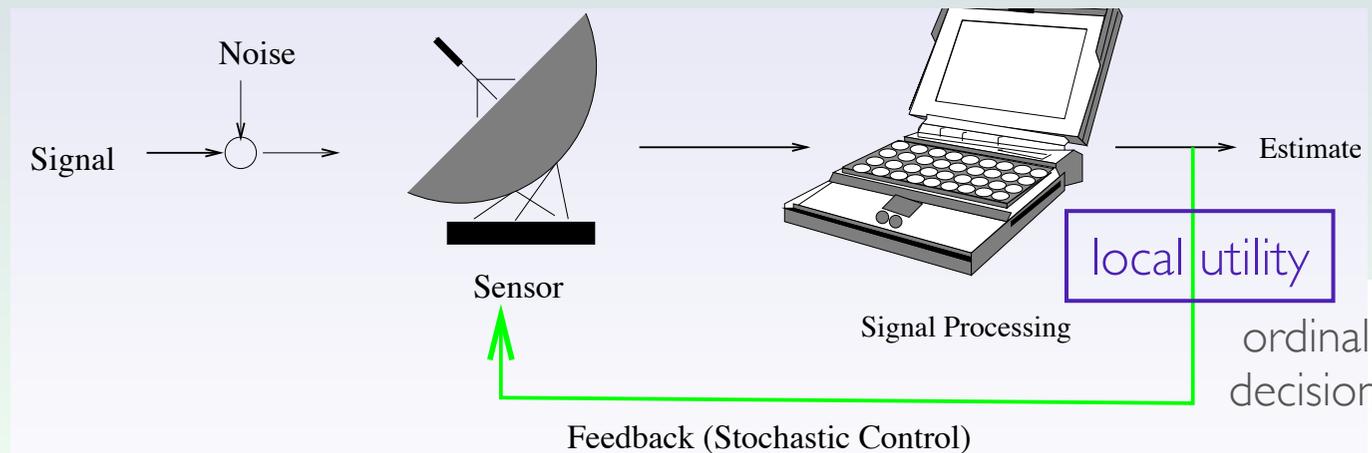


# SOCIAL (HUMAN) SENSORS

- Social sensor: Provides information about environment to a social network after interaction with other agents.
- Examples: Twitter posts, Facebook status updates, ratings on online reputation systems (Yelp, Tripadvisor)



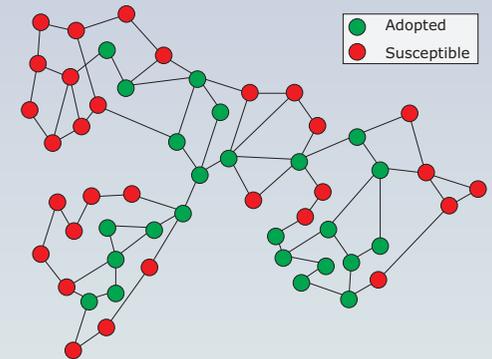
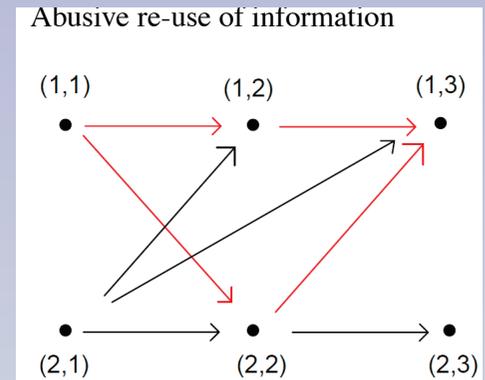
1. Social Sensors influence each other over a network
2. Social Sensors have dynamics: learn from past decisions and decisions of others
3. Social Sensors reveal quantized decisions (privacy) and are ordinal.
4. Social sensors go beyond physical sensors.



# OUTLINE

**Part 1.** How do agents learn from observations and actions of other agents?  
*Social Learning:* Herds and Data incest occur.

**Part 2.** How does information propagate in a large scale social network? Mean Field Dynamics of sensing over a random graph.

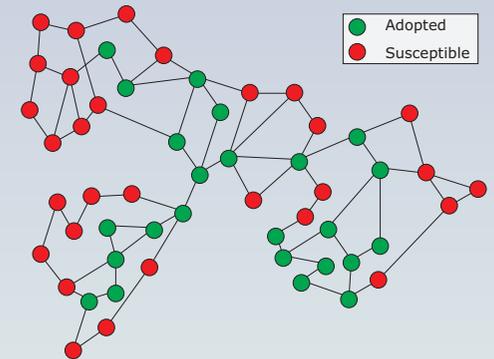
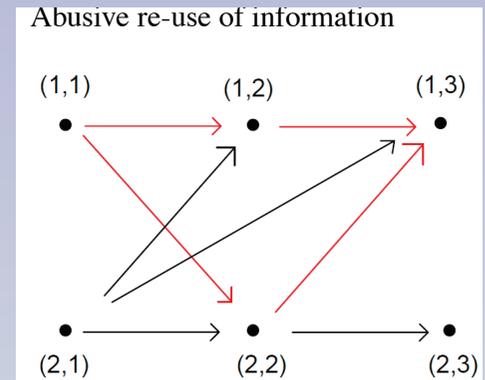


# OUTLINE

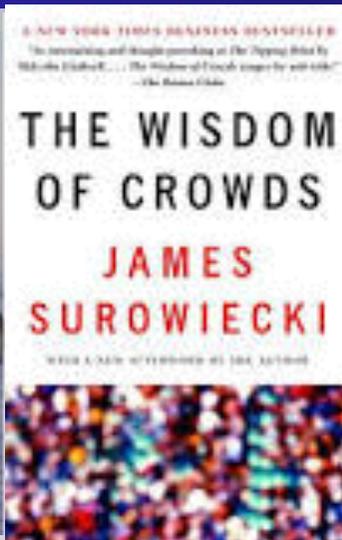
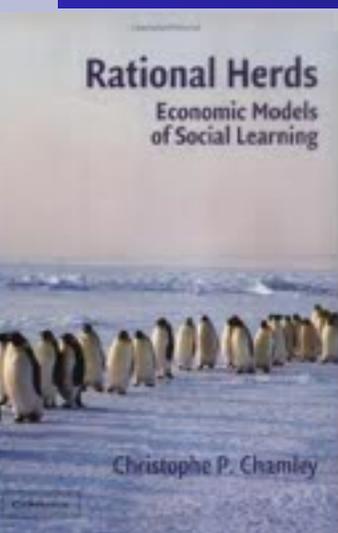
**Part 1.** How do agents learn from observations and actions of other agents?  
*Social Learning:* Herds and Data incest occur.

**Part 2.** How does information propagate in a large scale social network? Mean Field Dynamics of sensing over a random graph.

Unifying Theme: Interaction of dynamical sensors  
Interaction of local and global decision makers.  
Non-standard information patterns.



# PART 1: SOCIAL LEARNING FOR SENSING



## Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades

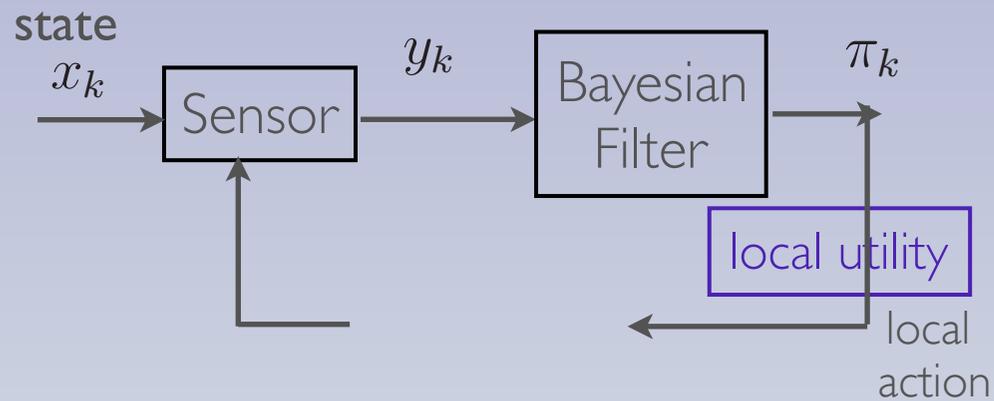
Sushil Bikhchandani, David Hirshleifer and Ivo Welch  
*The Journal of Economic Perspectives*  
Vol. 12, No. 3 (Summer, 1998), pp. 151-170

psychology, economics, sociology (groupthink),  
computer science, signal processing

## OUTLINE for Part I

- Bayesian model for Social Learning
- Global Sensing with Social Learning
- Data Incest in Social Learning

# 1.1 GLOBAL SENSING WITH SOCIAL LEARNING

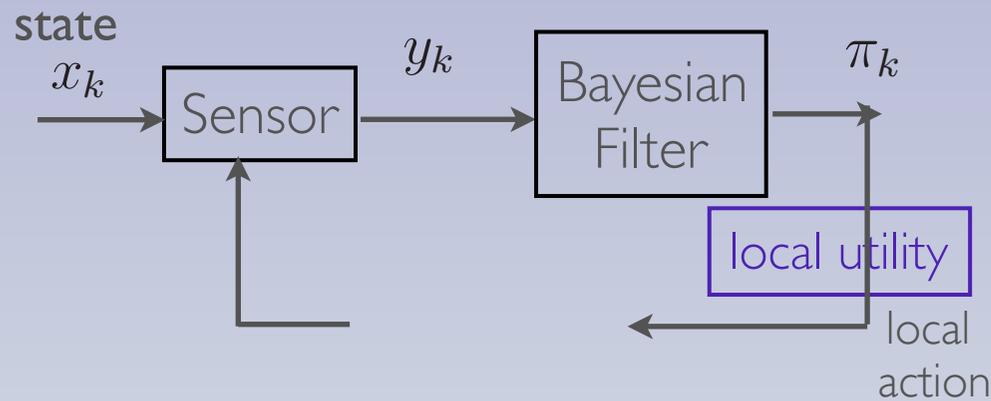


## *The social sensor*

### *Examples:*

- sentiment sensing in microblogs
- High Frequency Trading [Quant Finance]

# 1.1 GLOBAL SENSING WITH SOCIAL LEARNING



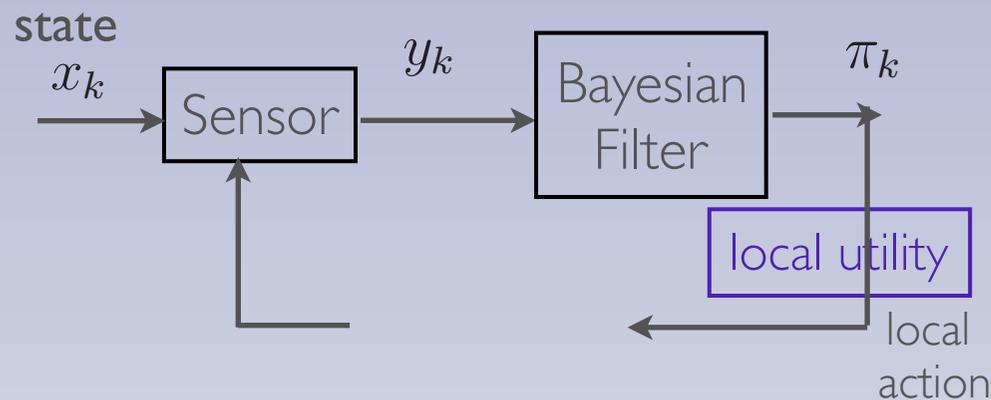
## *The social sensor*

### *Examples:*

- sentiment sensing in microblogs
- High Frequency Trading [Quant Finance]

*Some perspective on “vanilla” social learning ...*

# 1.1 GLOBAL SENSING WITH SOCIAL LEARNING



## The social sensor

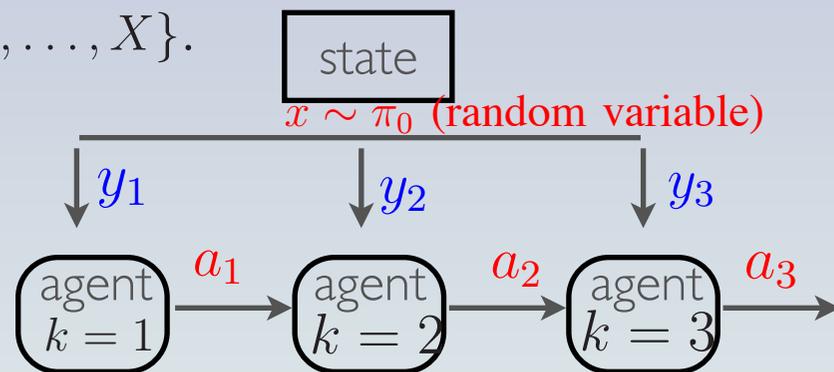
### Examples:

- sentiment sensing in microblogs
- High Frequency Trading [Quant Finance]

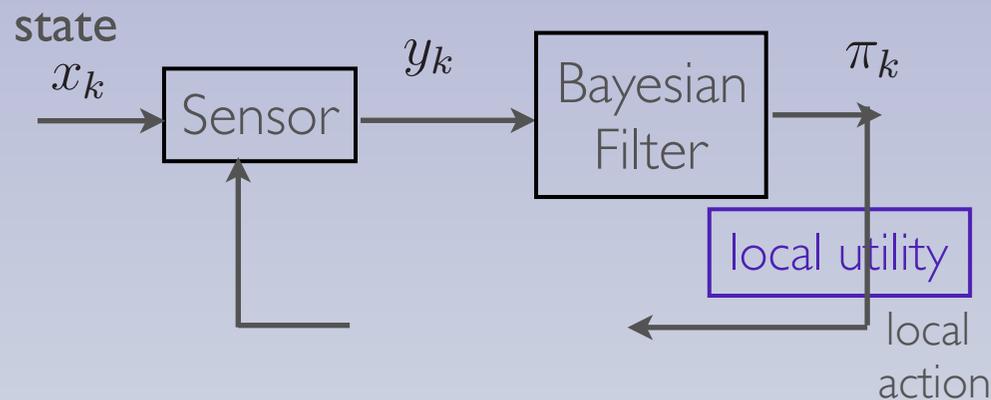
Agents  $k = 1, 2, \dots$  act sequentially to estimate  $x \in \{1, 2, \dots, X\}$ .

**Social Learning:** Given prior  $\pi_k = P(x|a_1, \dots, a_k)$  and observation  $y_{k+1}$  from a finite set

- Agent  $k + 1$ : picks local action  $a_{k+1} = f(y_{k+1}, \pi_k)$   
 $a_{k+1} = \operatorname{argmin}_{a \in \mathbb{A}} \mathbb{E}\{c(x, a) | a_1, \dots, a_k, y_{k+1}\}$
- Broadcast local action (ordinal decision)
- Other agents update public belief  
 $\pi_{k+1} = P(x|a_1, \dots, a_{k+1}) \propto \sum_y P(a_{k+1}|y, \pi_k) P(y|x) \pi_k$



# 1.1 GLOBAL SENSING WITH SOCIAL LEARNING



## The social sensor

### Examples:

- sentiment sensing in microblogs
- High Frequency Trading [Quant Finance]

Agents  $k = 1, 2, \dots$  act sequentially to estimate  $x \in \{1, 2, \dots, X\}$ .

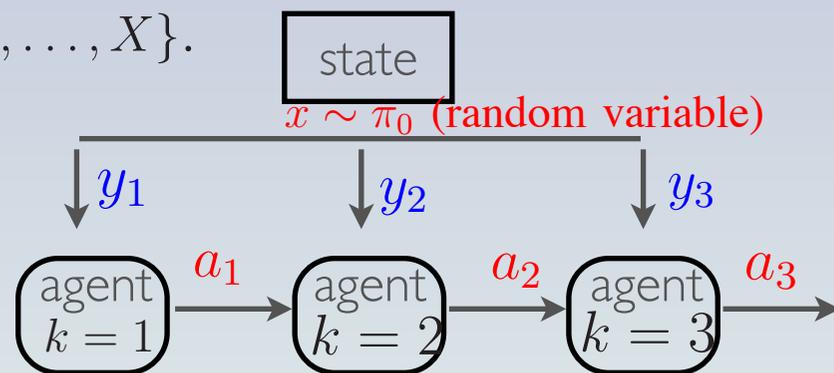
**Social Learning:** Given prior  $\pi_k = P(x|a_1, \dots, a_k)$  and observation  $y_{k+1}$  from a finite set

- Agent  $k + 1$ : picks local action  $a_{k+1} = f(y_{k+1}, \pi_k)$   
 $a_{k+1} = \operatorname{argmin}_{a \in \mathbb{A}} \mathbb{E}\{c(x, a) | a_1, \dots, a_k, y_{k+1}\}$
- Broadcast local action (ordinal decision)

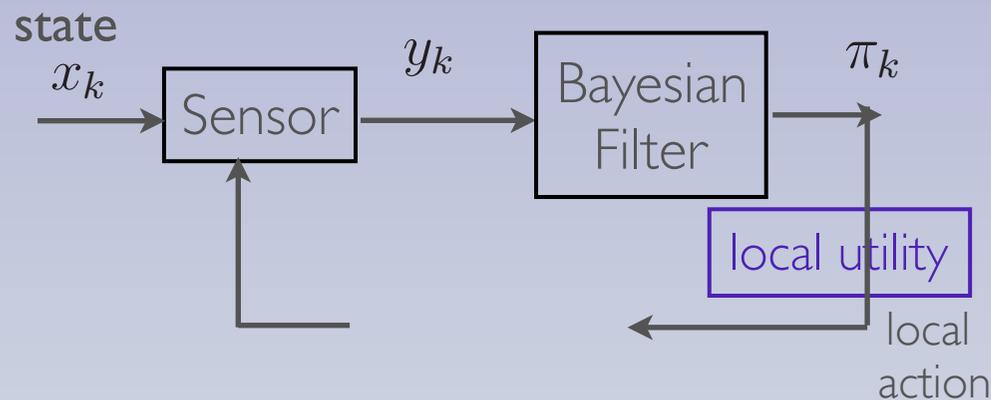
- Other agents update public belief

$$\pi_{k+1} = P(x|a_1, \dots, a_{k+1}) \propto \sum_y P(a_{k+1}|y, \pi_k) P(y|x) \pi_k$$

likelihood



# 1.1 GLOBAL SENSING WITH SOCIAL LEARNING



## The social sensor

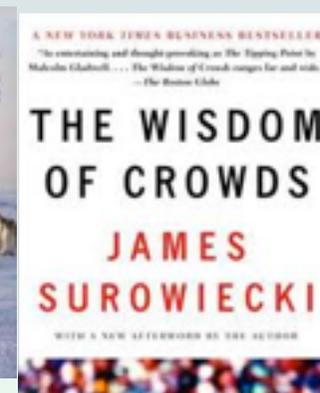
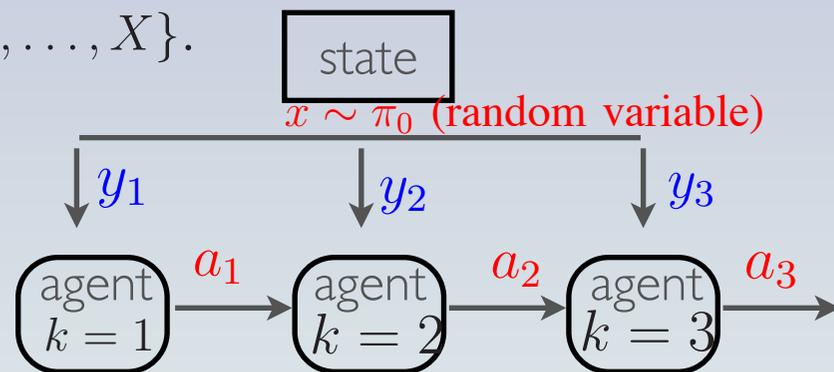
### Examples:

- sentiment sensing in microblogs
- High Frequency Trading [Quant Finance]

Agents  $k = 1, 2, \dots$  act sequentially to estimate  $x \in \{1, 2, \dots, X\}$ .

**Social Learning:** Given prior  $\pi_k = P(x|a_1, \dots, a_k)$  and observation  $y_{k+1}$  from a finite set

- Agent  $k + 1$ : picks local action  $a_{k+1} = f(y_{k+1}, \pi_k)$   
 $a_{k+1} = \operatorname{argmin}_{a \in \mathbb{A}} \mathbb{E}\{c(x, a) | a_1, \dots, a_k, y_{k+1}\}$
- Broadcast local action (ordinal decision)
- Other agents update public belief  
 $\pi_{k+1} = P(x|a_1, \dots, a_{k+1}) \propto \sum_y P(a_{k+1}|y, \pi_k) P(y|x) \pi_k$



**Theorem:** [Bhikchandani, J. Political Economy, 1992; Cover & Hellman, 1970]

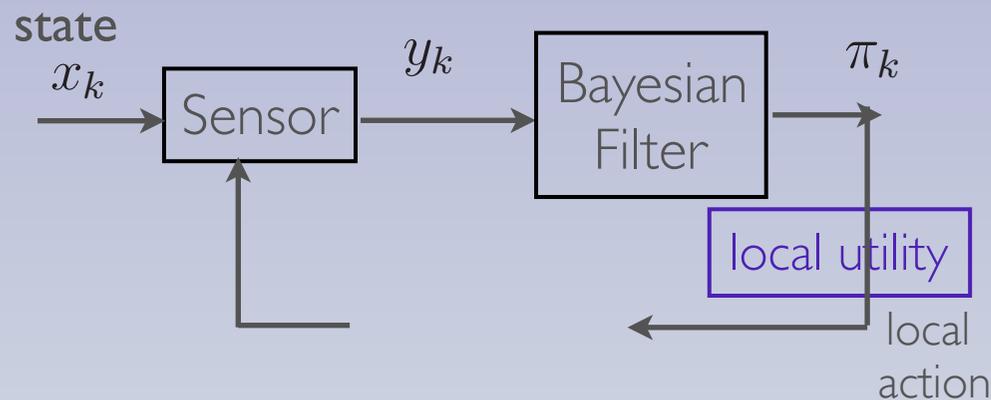
Agents eventually choose same action (information cascade, herd).

Social learning stops w.p.1 for finite  $k$ .

Acemoglu & Ozdaglar [2010,...]: General communication graphs.



# 1.1 GLOBAL SENSING WITH SOCIAL LEARNING



## The social sensor

### Examples:

- sentiment sensing in microblogs
- High Frequency Trading [Quant Finance]

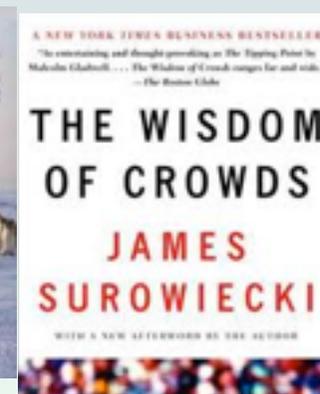
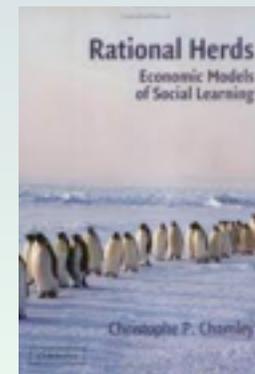
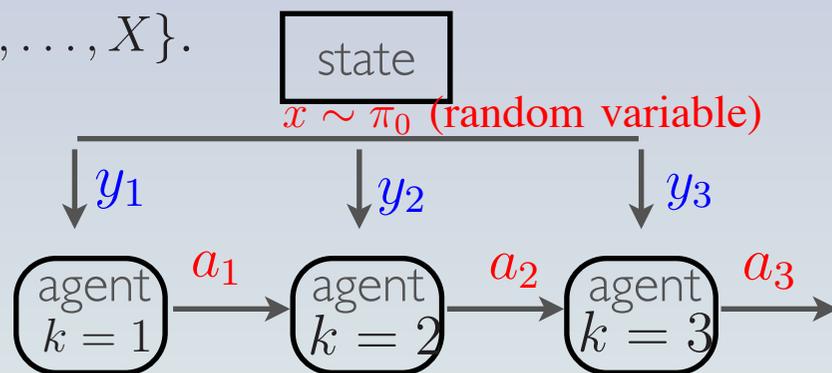
Agents  $k = 1, 2, \dots$  act sequentially to estimate  $x \in \{1, 2, \dots, X\}$ .

**Social Learning:** Given prior  $\pi_k = P(x|a_1, \dots, a_k)$  and observation  $y_{k+1}$  from a finite set

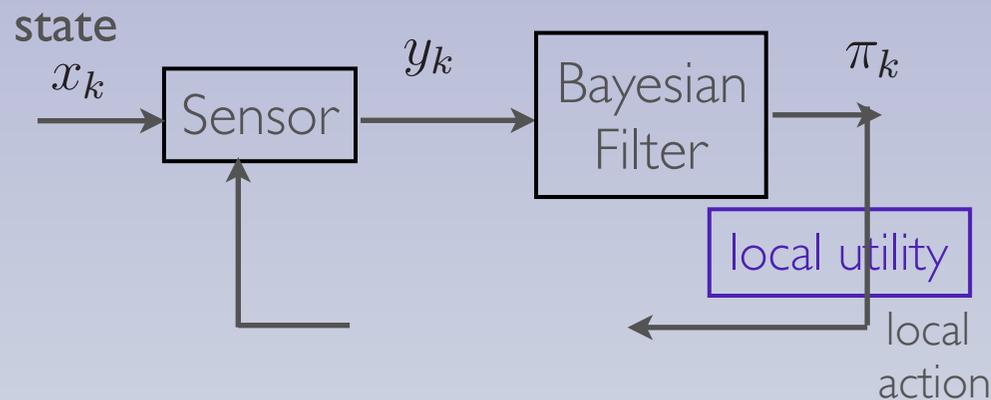
- Agent  $k + 1$ : picks local action  $a_{k+1} = f(y_{k+1}, \pi_k)$   
 $a_{k+1} = \operatorname{argmin}_{a \in \mathbb{A}} \mathbb{E}\{c(x, a) | a_1, \dots, a_k, y_{k+1}\}$
- Broadcast local action (ordinal decision)
- Other agents update public belief  
 $\pi_{k+1} = P(x|a_1, \dots, a_{k+1}) \propto \sum_y P(a_{k+1}|y, \pi_k) P(y|x) \pi_k$

• When I see others taking umbrellas, I take an umbrella without checking the weather forecast. I assume their private info is accurate.

. *Rational Herds*



# 1.1 GLOBAL SENSING WITH SOCIAL LEARNING



## The social sensor

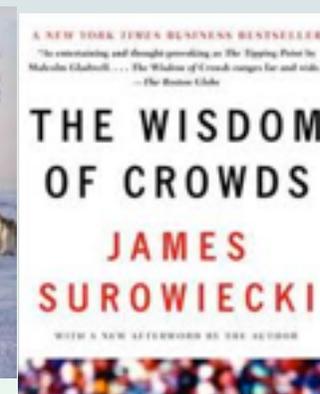
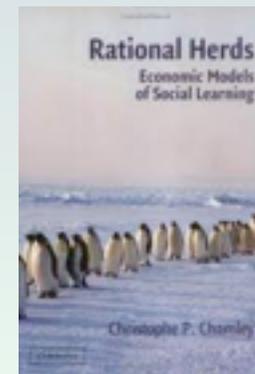
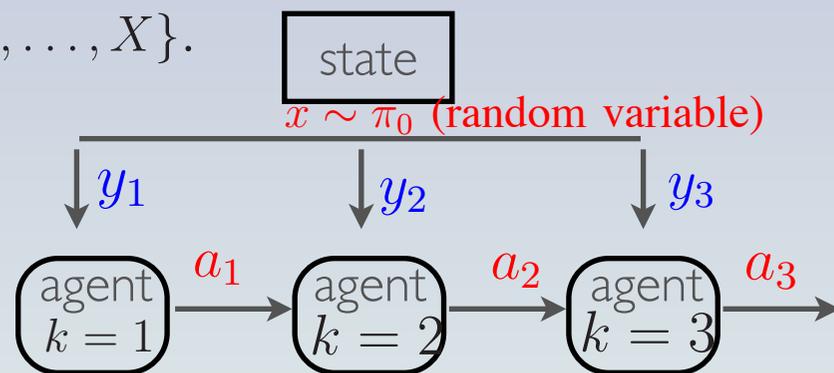
### Examples:

- sentiment sensing in microblogs
- High Frequency Trading [Quant Finance]

Agents  $k = 1, 2, \dots$  act sequentially to estimate  $x \in \{1, 2, \dots, X\}$ .

**Social Learning:** Given prior  $\pi_k = P(x|a_1, \dots, a_k)$  and observation  $y_{k+1}$  from a finite set

- Agent  $k + 1$ : picks local action  $a_{k+1} = f(y_{k+1}, \pi_k)$   
 $a_{k+1} = \operatorname{argmin}_{a \in \mathbb{A}} \mathbb{E}\{c(x, a) | a_1, \dots, a_k, y_{k+1}\}$
- Broadcast local action (ordinal decision)
- Other agents update public belief  
 $\pi_{k+1} = P(x|a_1, \dots, a_{k+1}) \propto \sum_y P(a_{k+1}|y, \pi_k) P(y|x) \pi_k$



In 1995, management gurus Treacy & Wiersema secretly bought 50,000 copies of their own book. Made NY times best seller list. How to cope with malicious agents?

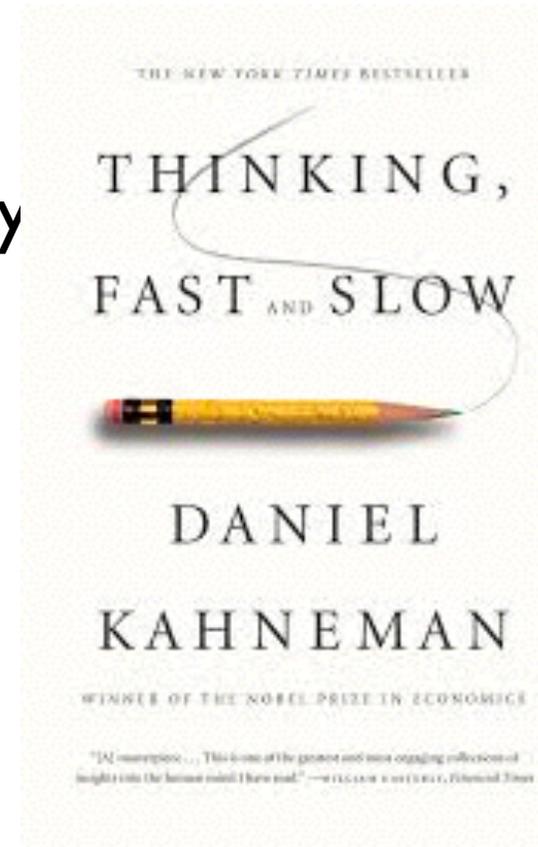
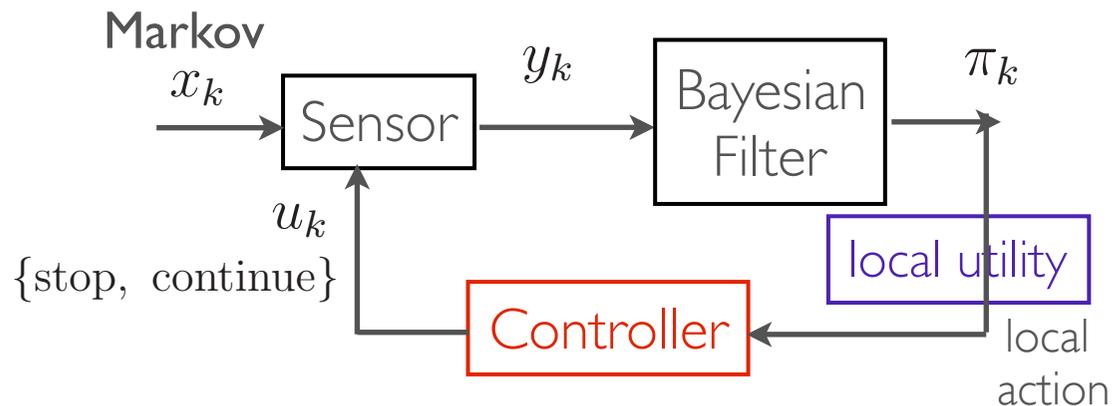


Social learning results in herding.

Q1. How do Local and Global Agents Interact in decision making?

Q2. How to optimize Social Learning to delay herding?

Q3. How to price a product?



# Q1. HOW DO LOCAL & GLOBAL DECISION MAKERS INTERACT?

**Example: Multiagent Quickest Change Detection**



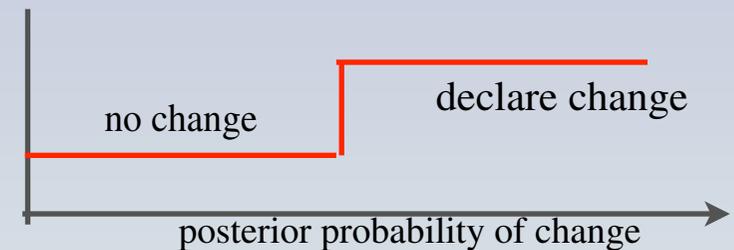
# Q1. HOW DO LOCAL & GLOBAL DECISION MAKERS INTERACT?

## Example: Multiagent Quickest Change Detection

Observations  $y_k \sim \begin{cases} B_1(\cdot) & k \leq \tau^0 \\ B_2(\cdot) & k > \tau^0 \end{cases}$ , where  $\tau^0 = \text{change time (usually geometric)}$

**Aim:** Compute time  $\tau$  to announce change: **Minimize**  $\mathbb{E}_{\pi_0}^{\mu} \left\{ \underbrace{d|\tau - \tau^0|^+}_{\text{delay}} + \underbrace{f I(\tau < \tau^0)}_{\text{false-alarm}} \right\}$

**Classical:** Given posterior  $\pi_k = P(\text{change} | y_1, \dots, y_k)$ :  
Optimal decision policy is threshold. [Shiriyayev 1950s].



# Q1. HOW DO LOCAL & GLOBAL DECISION MAKERS INTERACT?

## Example: Multiagent Quickest Change Detection

Observations  $y_k \sim \begin{cases} B_1(\cdot) & k \leq \tau^0 \\ B_2(\cdot) & k > \tau^0 \end{cases}$ , where  $\tau^0 = \text{change time (usually geometric)}$

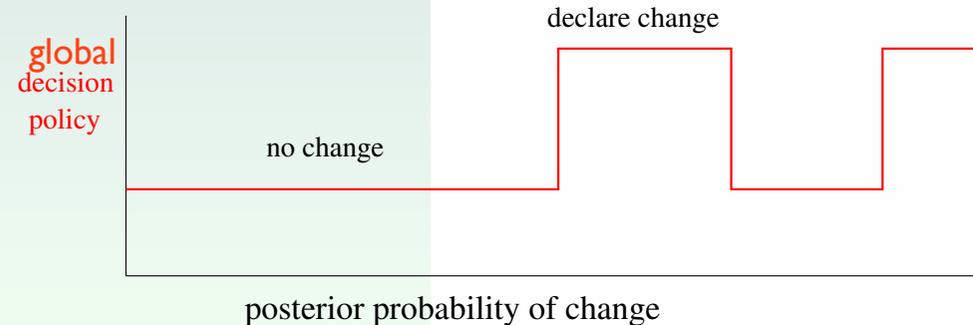
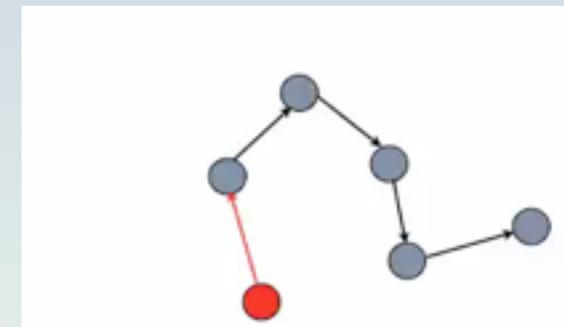
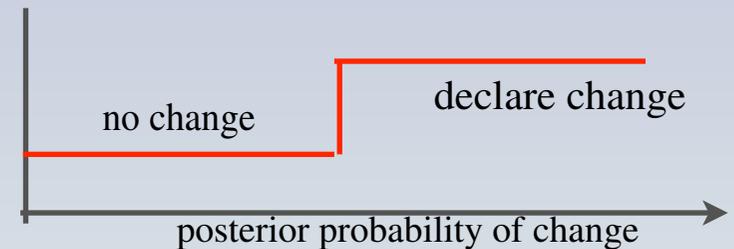
**Aim:** Compute time  $\tau$  to announce change: **Minimize**  $\mathbb{E}_{\pi_0}^{\mu} \left\{ \underbrace{d|\tau - \tau^0|^+}_{\text{delay}} + f \underbrace{I(\tau < \tau^0)}_{\text{false-alarm}} \right\}$

**Classical:** Given posterior  $\pi_k = P(\text{change} | y_1, \dots, y_k)$ :  
Optimal decision policy is threshold. [Shiriyayev 1950s].

**Social Learning:** Public belief  $\pi_{k-1} = P(\text{change} | a_1, \dots, a_{k-1})$ .

- Agent  $k$ : Observes  $y_k \sim P(y|x)$
- Broadcasts action  $a_k = \arg \min_a \mathbb{E}\{c(\text{state}, a) | a_1, \dots, a_{k-1}, y_k\}$ .
- Other agents update public belief (social learning filter)

$$\pi_k = P(\text{change} | a_1, \dots, a_k) \propto \sum P(a_k | y, \pi_{k-1}) P(y|x) \times (\text{predictor update})$$



# Q1. HOW DO LOCAL & GLOBAL DECISION MAKERS INTERACT?

## Example: Multiagent Quickest Change Detection

Observations  $y_k \sim \begin{cases} B_1(\cdot) & k \leq \tau^0 \\ B_2(\cdot) & k > \tau^0 \end{cases}$ , where  $\tau^0 = \text{change time (usually geometric)}$

**Aim:** Compute time  $\tau$  to announce change: **Minimize**  $\mathbb{E}_{\pi_0}^{\mu} \left\{ \underbrace{d|\tau - \tau^0|^+}_{\text{delay}} + \underbrace{f I(\tau < \tau^0)}_{\text{false-alarm}} \right\}$

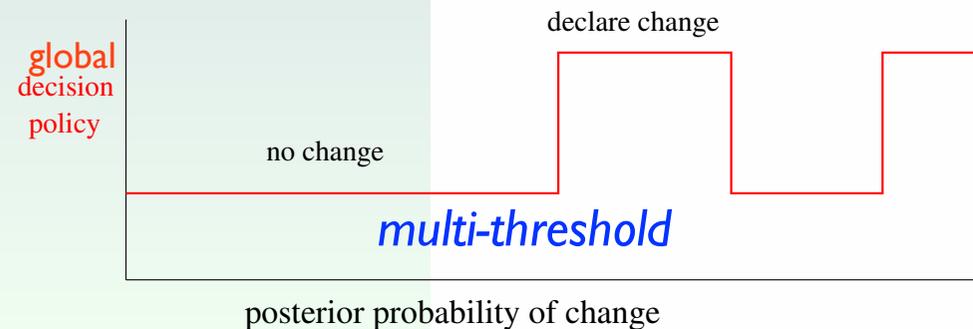
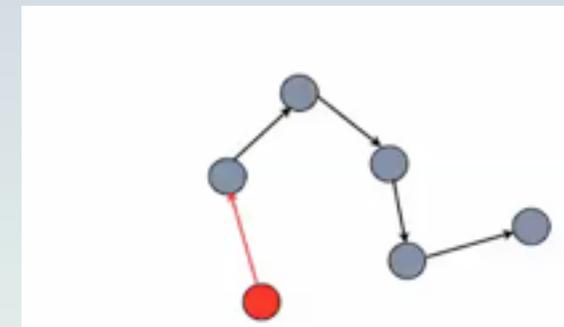
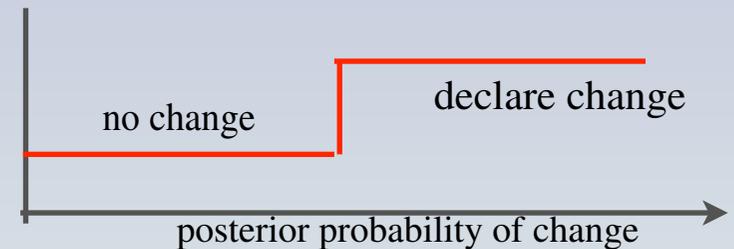
**Classical:** Given posterior  $\pi_k = P(\text{change} | y_1, \dots, y_k)$ :  
Optimal decision policy is threshold. [Shiriyayev 1950s].

**Social Learning:** Public belief  $\pi_{k-1} = P(\text{change} | a_1, \dots, a_{k-1})$ .

- Agent  $k$ : Observes  $y_k \sim P(y|x)$
- Broadcasts action  $a_k = \arg \min_a \mathbb{E}\{c(\text{state}, a) | a_1, \dots, a_{k-1}, y_k\}$ .
- Other agents update public belief (**social learning filter**)

$$\pi_k = P(\text{change} | a_1, \dots, a_k) \propto \sum P(a_k | y, \pi_{k-1}) P(y|x) \times (\text{predictor update})$$

When should global decision-maker declare change?



# Q1. HOW DO LOCAL & GLOBAL DECISION MAKERS INTERACT?

## Example: Multiagent Quickest Change Detection

Observations  $y_k \sim \begin{cases} B_1(\cdot) & k \leq \tau^0 \\ B_2(\cdot) & k > \tau^0 \end{cases}$ , where  $\tau^0 =$  change time (usually geometric)

**Aim:** Compute time  $\tau$  to announce change: **Minimize**  $\mathbb{E}_{\pi_0}^{\mu} \{ \underbrace{d|\tau - \tau^0|^+}_{\text{delay}} + f \underbrace{I(\tau < \tau^0)}_{\text{false-alarm}} \}$

**Classical:** Given posterior  $\pi_k = P(\text{change} | y_1, \dots, y_k)$ :  
Optimal decision policy is threshold. [Shiriyayev 1950s].

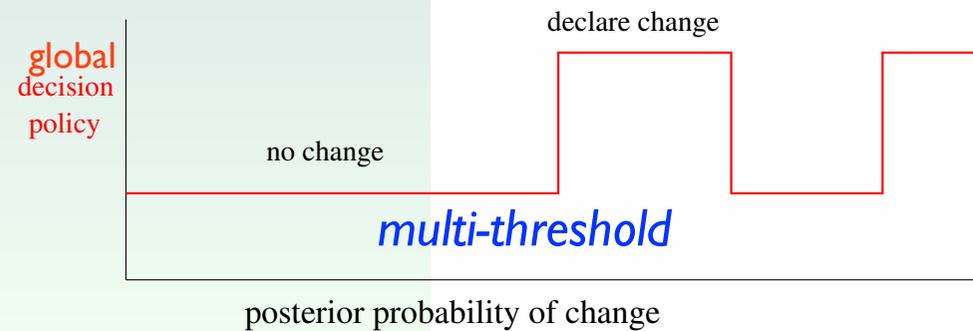
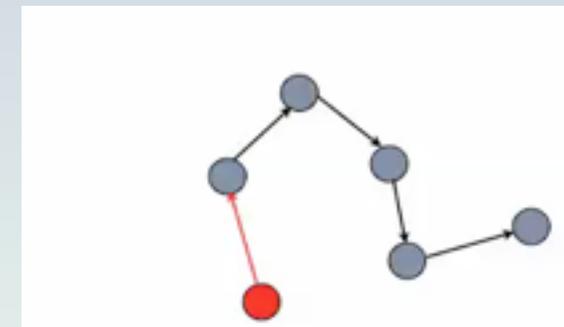
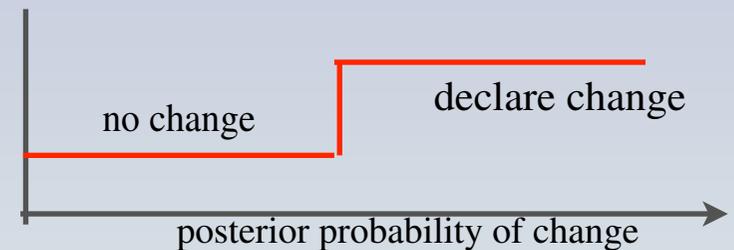
**Social Learning:** Public belief  $\pi_{k-1} = P(\text{change} | a_1, \dots, a_{k-1})$ .

- Agent  $k$ : Observes  $y_k \sim P(y|x)$
- Broadcasts action  $a_k = \arg \min_a \mathbb{E}\{c(\text{state}, a) | a_1, \dots, a_{k-1}, y_k\}$ .
- Other agents update public belief (social learning filter)

$$\pi_k = P(\text{change} | a_1, \dots, a_k) \propto \sum P(a_k | y, \pi_{k-1}) P(y|x) \times (\text{predictor update})$$

When should global decision-maker declare change?

.Stopping set is non-convex



# Q1. HOW DO LOCAL & GLOBAL DECISION MAKERS INTERACT?

## Example: Multiagent Quickest Change Detection

Observations  $y_k \sim \begin{cases} B_1(\cdot) & k \leq \tau^0 \\ B_2(\cdot) & k > \tau^0 \end{cases}$ , where  $\tau^0 = \text{change time (usually geometric)}$

**Aim:** Compute time  $\tau$  to announce change: **Minimize**  $\mathbb{E}_{\pi_0}^{\mu} \left\{ \underbrace{d|\tau - \tau^0|^+}_{\text{delay}} + f \underbrace{I(\tau < \tau^0)}_{\text{false-alarm}} \right\}$

**Classical:** Given posterior  $\pi_k = P(\text{change} | y_1, \dots, y_k)$ :  
Optimal decision policy is threshold. [Shiriyayev 1950s].

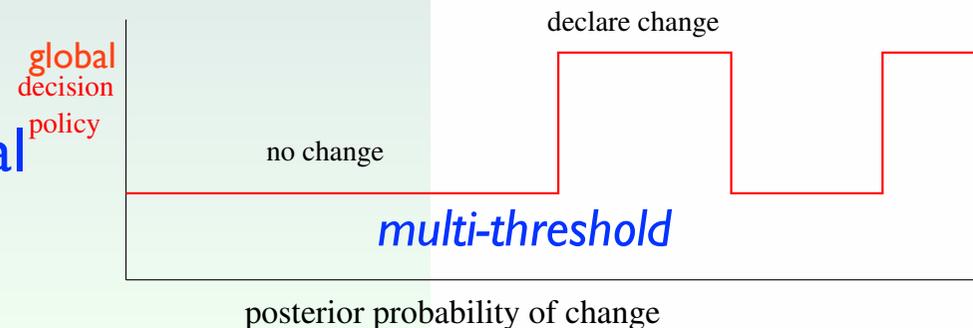
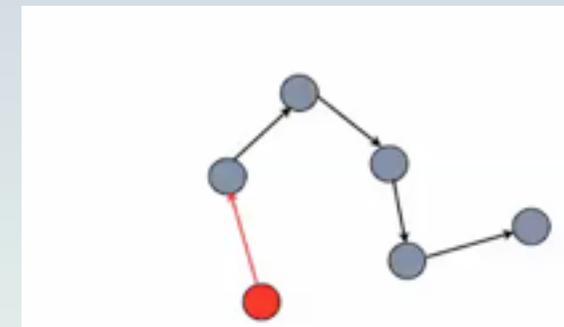
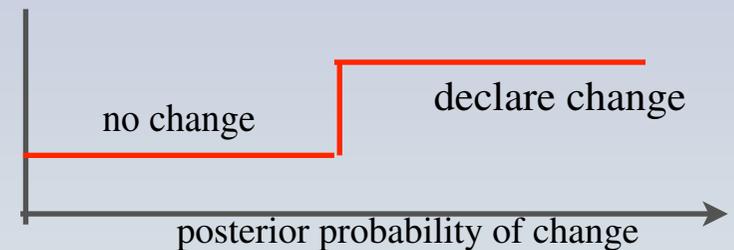
**Social Learning:** Public belief  $\pi_{k-1} = P(\text{change} | a_1, \dots, a_{k-1})$ .

- Agent  $k$ : Observes  $y_k \sim P(y|x)$
- Broadcasts action  $a_k = \arg \min_a \mathbb{E}\{c(\text{state}, a) | a_1, \dots, a_{k-1}, y_k\}$ .
- Other agents update public belief (social learning filter)

$$\pi_k = P(\text{change} | a_1, \dots, a_k) \propto \sum P(a_k | y, \pi_{k-1}) P(y|x) \times (\text{predictor update})$$

When should global decision-maker declare change?

**Summary: Global Decision making using local decisions is non-monotone!**



## Q2: HOW TO OPTIMIZE SOCIAL LEARNING?

**Social Learning:** Choose local decision greedily:  $a_k = \min_a \mathbb{E}_{\pi_{k-1}, y_k} \{c(x, a)\}$ .  
Results in **herding**. Posterior  $\pi_k = P(x|a_1, \dots, a_k)$  freezes.



## Q2: HOW TO OPTIMIZE SOCIAL LEARNING?

**Social Learning:** Choose local decision greedily:  $a_k = \min_a \mathbb{E}_{\pi_{k-1}, y_k} \{c(x, a)\}$ .  
Results in **herding**. Posterior  $\pi_k = P(x|a_1, \dots, a_k)$  freezes.

**Socialistic Learning:** [More Sophisticated Protocol] To estimate  $x \sim \pi_0$

*I choose my local decision to sacrifice my local utility so that my action provides useful information to subsequent agents*

Benevolent agents choose **local decision** by minimizing *social welfare* cost:



# Q2: HOW TO OPTIMIZE SOCIAL LEARNING?

**Social Learning:** Choose local decision greedily:  $a_k = \min_a \mathbb{E}_{\pi_{k-1}, y_k} \{c(x, a)\}$ .  
Results in **herding**. Posterior  $\pi_k = P(x|a_1, \dots, a_k)$  freezes.

**Socialistic Learning:** [More Sophisticated Protocol] To estimate  $x \sim \pi_0$

*I choose my local decision to sacrifice my local utility so that my action provides useful information to subsequent agents*

Benevolent agents choose **local decision** by minimizing *social welfare* cost:

$$a_k = \mu^*(\pi_{k-1}, y_k) \in \{y_k(\text{reveal}), a_{k-1}(\text{herd})\}, \quad \mu^* = \underset{\mu}{\arg \min} \mathbb{E}_{\pi_0}^{\mu} \left\{ \sum_{k=1}^N c(x, a_k) \right\}$$

global  
decision



## Q2: HOW TO OPTIMIZE SOCIAL LEARNING?

**Social Learning:** Choose local decision greedily:  $a_k = \min_a \mathbb{E}_{\pi_{k-1}, y_k} \{c(x, a)\}$ .  
Results in **herding**. Posterior  $\pi_k = P(x|a_1, \dots, a_k)$  freezes.

**Socialistic Learning:** [More Sophisticated Protocol] To estimate  $x \sim \pi_0$

*I choose my local decision to sacrifice my local utility so that my action provides useful information to subsequent agents*

Benevolent agents choose **local decision** by minimizing *social welfare* cost:

$$a_k = \mu^*(\pi_{k-1}, y_k) \in \{y_k(\text{reveal}), a_{k-1}(\text{herd})\}, \quad \mu^* = \arg \min_{\mu} \mathbb{E}_{\pi_0}^{\mu} \left\{ \sum_{k=1}^N c(x, a_k) \right\}$$

Partially Observed Stochastic Control Problem.



# Q2: HOW TO OPTIMIZE SOCIAL LEARNING?

**Social Learning:** Choose local decision greedily:  $a_k = \min_a \mathbb{E}_{\pi_{k-1}, y_k} \{c(x, a)\}$ .  
Results in **herding**. Posterior  $\pi_k = P(x|a_1, \dots, a_k)$  freezes.

**Socialistic Learning:** [More Sophisticated Protocol] To estimate  $x \sim \pi_0$

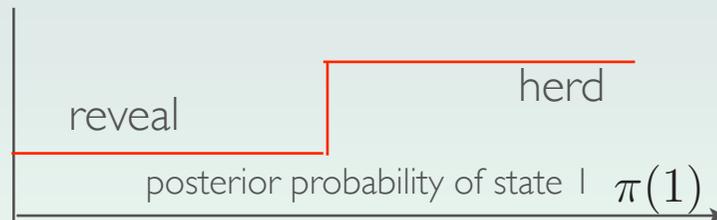
*I choose my local decision to sacrifice my local utility so that my action provides useful information to subsequent agents*

Benevolent agents choose **local decision** by minimizing *social welfare* cost:

$$a_k = \mu^*(\pi_{k-1}, y_k) \in \{y_k(\text{reveal}), a_{k-1}(\text{herd})\}, \quad \mu^* = \arg \min_{\mu} \mathbb{E}_{\pi_0}^{\mu} \left\{ \sum_{k=1}^N c(x, a_k) \right\}$$

Can show: [IEEE Trans. Info. Theory, 2011]

- Under supermodular assumptions global decision policy is threshold.



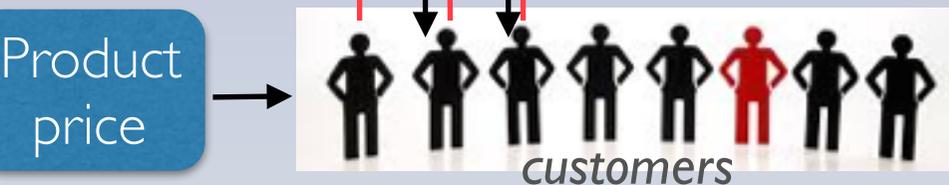
- Global decision policy: Initially socialistic then capitalistic.
- Privacy vs Reputation



# Q3: PRICE HIGH OR LOW?



# Q3: PRICE HIGH OR LOW?

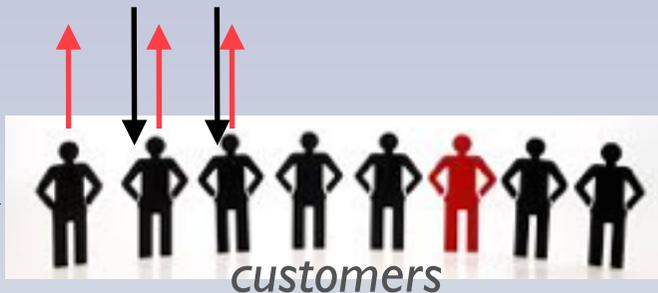


Each Sale to a customer

- Brings in money
- Customers perform social learning on the product

*How should the product price be chosen over time?*

# Q3: PRICE HIGH OR LOW?



Each Sale to a customer

- Brings in money
- Customers perform social learning on the product

*How should the product price be chosen over time?*

**Optimal Pricing for Monopolist.**

Choose price  $u_k = \mu(\pi_{k-1})$

**Monopolist reward:**

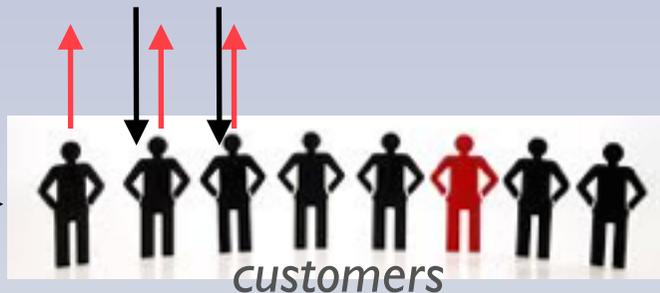
$$\text{Compute: } \sup_{\mu} J_{\mu}(\pi) = \mathbb{E}_{\pi}^{\mu} \left\{ \sum_{k=1}^{\infty} \rho^{k-1} I(a_k = \text{buy}) u_k \right\}$$

**Agent:**

$$y_k \sim p(\cdot|x), x \sim \pi_0.$$
$$a_k = \operatorname{argmin}_a c'_{a,u_k,y_k} \pi_{k-1}$$

**Public belief:**  $\pi_k = T(\pi_{k-1}, a_k)$

# Q3: PRICE HIGH OR LOW?



Each Sale to a customer

- Brings in money
- Customers perform social learning on the product

*How should the product price be chosen over time?*

**Optimal Pricing for Monopolist.**

Choose price  $u_k = \mu(\pi_{k-1})$

**Monopolist reward:**

$$\text{Compute: } \sup_{\mu} J_{\mu}(\pi) = \mathbb{E}_{\pi}^{\mu} \left\{ \sum_{k=1}^{\infty} \rho^{k-1} I(a_k = \text{buy}) u_k \right\}$$

**Agent:**  $y_k \sim p(\cdot|x), x \sim \pi_0.$

$$a_k = \operatorname{argmin}_a c'_{a,u_k,y_k} \pi_{k-1}$$

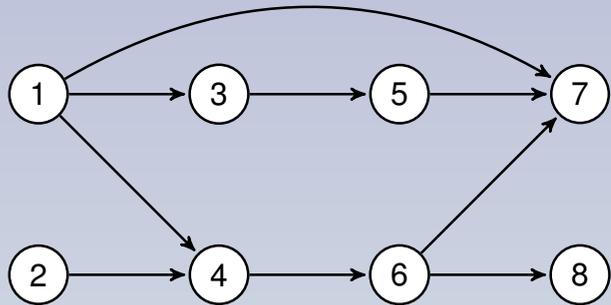
**Public belief:**  $\pi_k = T(\pi_{k-1}, a_k)$

*“Optimal soln”:* Price high initially; build elite customer base; then gradually decrease prices.

# DISCUSSION FOR PART 1: SOCIAL LEARNING

## 1. Data Incest + Herding in social learning over general graphs

Borkar & Varaiya: 1982



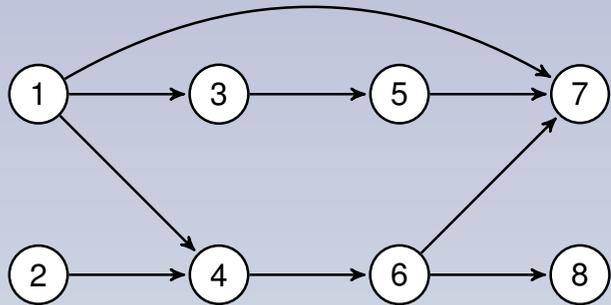
*Data incest* results in overconfidence in estimate

- **How to build unbiased reputation system?**
- One-star increase in the Yelp rating maps to 5-9 % revenue increase. [Harvard Business School, 2011]

# DISCUSSION FOR PART 1: SOCIAL LEARNING

## 1. Data Incest + Herding in social learning over general graphs

Borkar & Varaiya: 1982



*Data incest* results in overconfidence in estimate

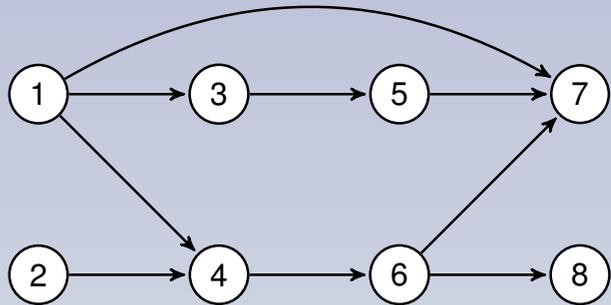
- **How to build unbiased reputation system?**
- One-star increase in the Yelp rating maps to 5-9 % revenue increase. [Harvard Business School, 2011]

## 2. Coherent Risk Measures instead of expected value.

# DISCUSSION FOR PART 1: SOCIAL LEARNING

## 1. Data Incest + Herding in social learning over general graphs

Borkar & Varaiya: 1982



*Data incest* results in overconfidence in estimate

- **How to build unbiased reputation system?**
- One-star increase in the Yelp rating maps to 5-9 % revenue increase. [Harvard Business School, 2011]

## 2. Coherent Risk Measures instead of expected value.

$$a_k = \arg \min_a \mathbb{E}_{\pi_{k-1}, y_k} \{c(x, a)\} \quad \longrightarrow \quad a_k = \arg \min \mathcal{R}_{\pi_{k-1}, a_k} \{c(x, a)\}$$

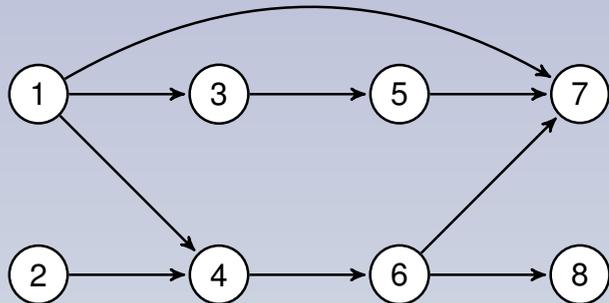
*Example:*  $a_k = \operatorname{argmin}_{a \in \mathcal{A}} \{\operatorname{CVaR}_\alpha(c(x_k, a))\}$  (Rockafellar & Uryasev, 2000)

- **Decisions are ordinal in belief and observation**
- **Risk averse agents herd more frequently at cheaper costs and therefore compromise state estimate.**

# DISCUSSION FOR PART 1: SOCIAL LEARNING

## 1. Data Incest + Herding in social learning over general graphs

Borkar & Varaiya: 1982



*Data incest* results in overconfidence in estimate

- **How to build unbiased reputation system?**
- One-star increase in the Yelp rating maps to 5-9 % revenue increase. [Harvard Business School, 2011]

2. Coherent Risk Measures instead of expected value.

3. **Experimental Data:** Collaboration with Department of Psychology UBC. *In perceptual tasks, data incest patterns occurred 79% and caused individuals to modify actions 21% of the time.*

4. **Human Interpretation of Data:** Should two security guards look at one TV monitor and then discuss suspicious behavior? (with data incest and herding)

5. **In which order should a panel of experts be polled?**  
[Ottaviani & Sorenson, 2001]

# Extension: Panel of Experts

High reputation



Good reputation



In which order to poll agents? If senior agents talk first, they unduly affect junior agents.

No reputation



Low reputation



# Extension: Panel of Experts

High reputation



Good reputation



Seniority Rule?



In which order to poll agents? If senior agents talk first, they unduly affect junior agents.



No reputation



Low reputation



# Extension: Panel of Experts

High reputation



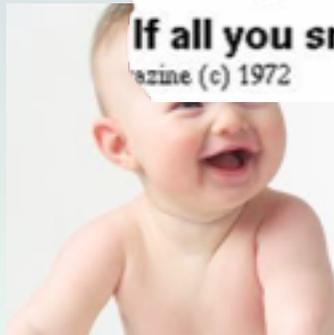
Good reputation



Seniority Rule?



No reputation



Low reputation



"If all you smart cookies agree, who am I to dissent?"  
Magazine (c) 1972

# Extension: Panel of Experts

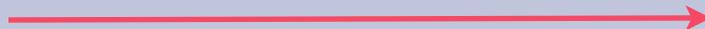
High reputation



Good reputation



Seniority Rule?



*For 94% of problems, the group's final answer was the first answer suggested, and people with dominant personalities tend to speak first and most forcefully...*

Anderson & Kilduff, Berkeley Hass School, 2009.



No reputation



Low reputation



# Extension: Panel of Experts

High reputation



Good reputation



Seniority Rule?



*For 94% of problems, the group's final answer was the first answer suggested, and people with dominant personalities tend to speak first and most forcefully...*

Anderson & Kilduff, Berkeley Hass School, 2009.



No reputation



Low reputation



Ottaviani, Sorensen, 2001. *Information Aggregation in Debate: Who should speak first?*

# SUMMARY OF PART 1

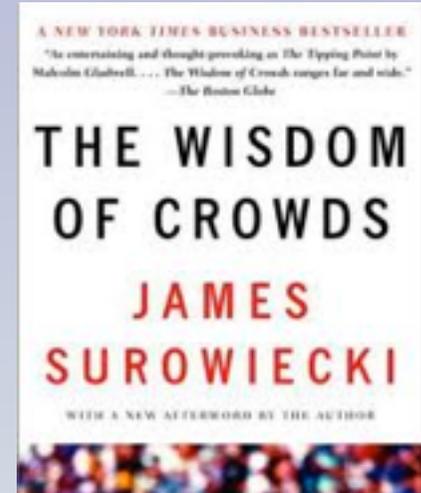
Social learning can result in herding and information cascades. Individuals end up blindly imitating others. ***Groupthink***

- Group behavior may not be wise.
- Crowds reduce diversity & are misleading

Data incest results in bias

**Datasets are non-informative.**

***Rational inattention models*** (Sims): ability of the human to absorb information is modeled via the information theoretic capacity of a communication channel



# SUMMARY OF PART 1

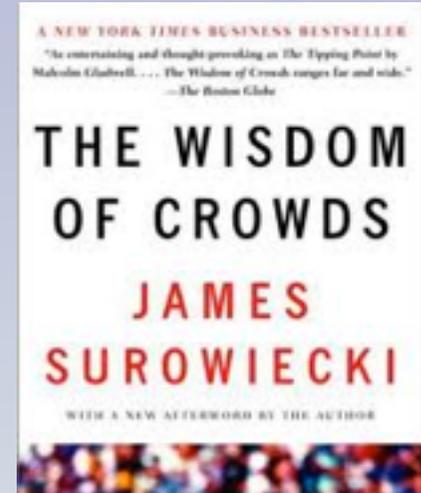
Social learning can result in herding and information cascades. Individuals end up blindly imitating others. ***Groupthink***

- Group behavior may not be wise.
- Crowds reduce diversity & are misleading

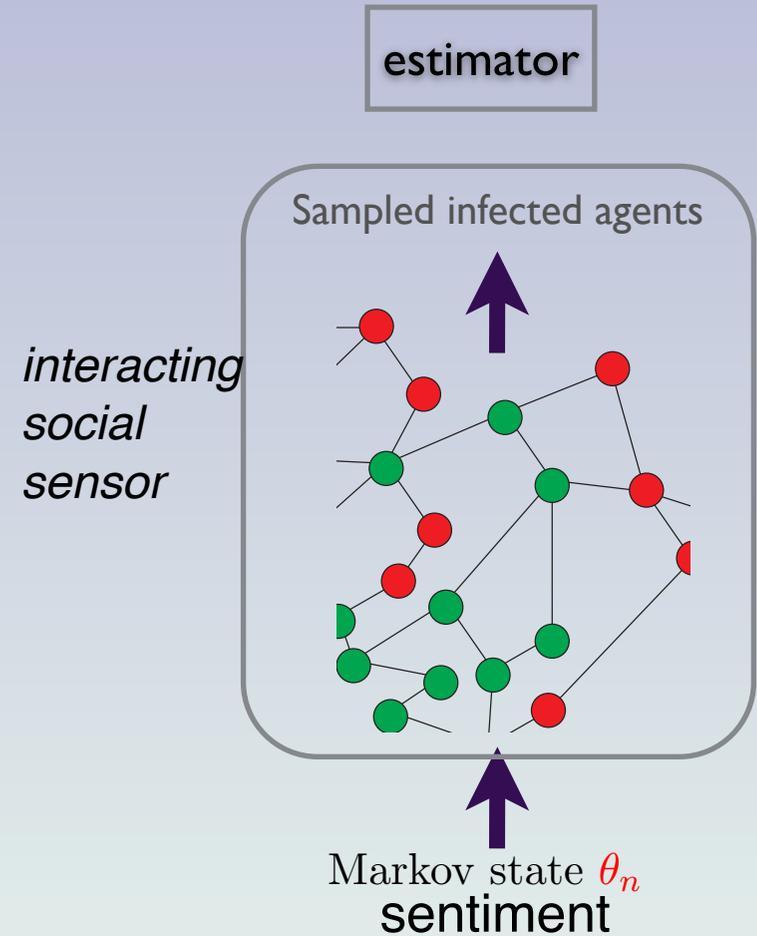
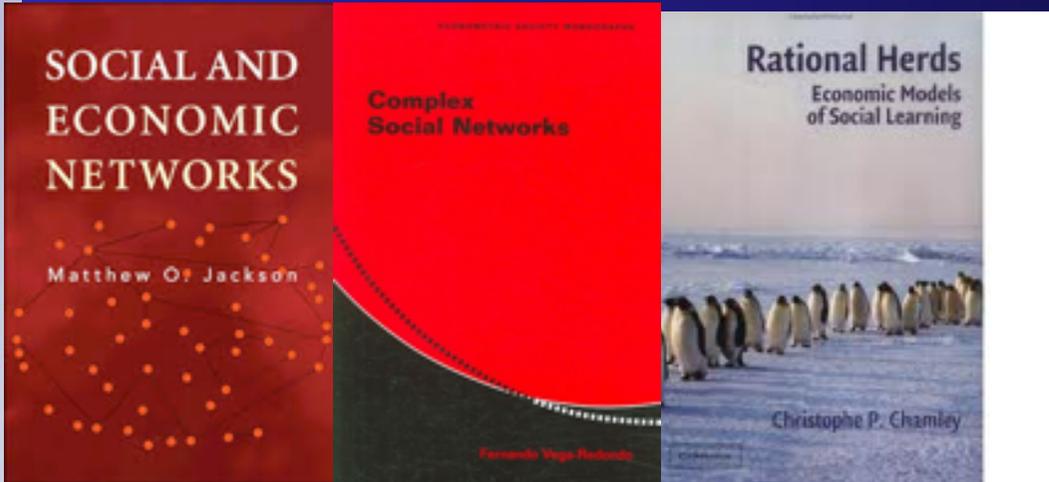
Data incest results in bias  
**Datasets are non-informative.**

***Rational inattention models*** (Sims): ability of the human to absorb information is modeled via the information theoretic capacity of a communication channel

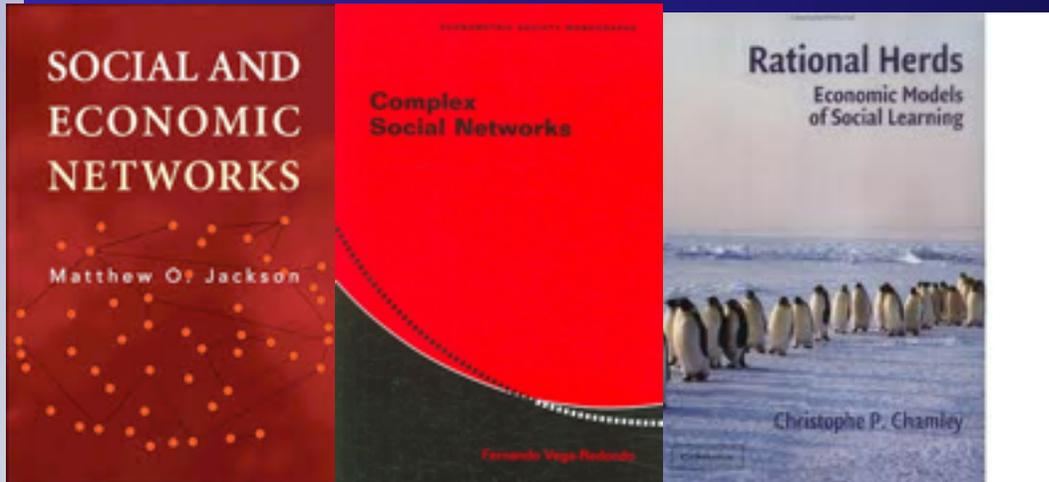
In complex settings, herding can result in interesting behavior:  
***Nobody goes there anymore ... it is always too crowded (Yogi Berra)***



# PART 2: INFORMATION DIFFUSION IN SOCIAL NETWORKS



# PART 2: INFORMATION DIFFUSION IN SOCIAL NETWORKS



## Why?

Twitter mood predicts stock market, *J. Computational Sci*, 2011.

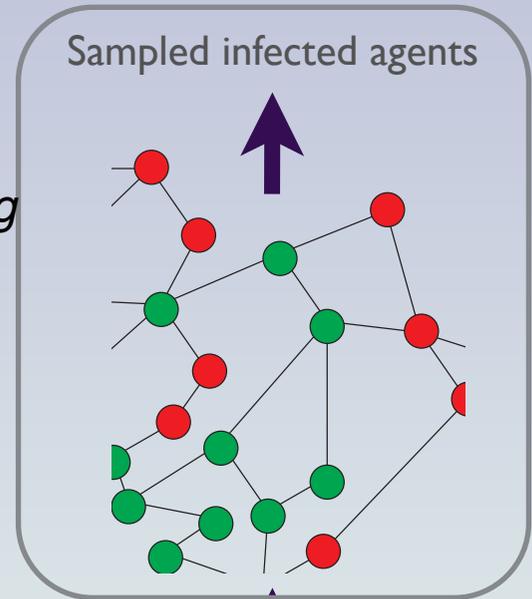
53% of people on Twitter recommend companies/products in tweets; 48% delivering on their intention to buy the product. ROI Research for Performance, 2010.

Consumer reviews are trusted nearly 12 times more than descriptions from manufacturers. eMarketer, 2010.

myYearbook: 81% of respondents recede advice from friends related to product purchase, 74% found advice to be influential in decision to buy. ClickZ, 2010.

interacting  
social  
sensor

estimator

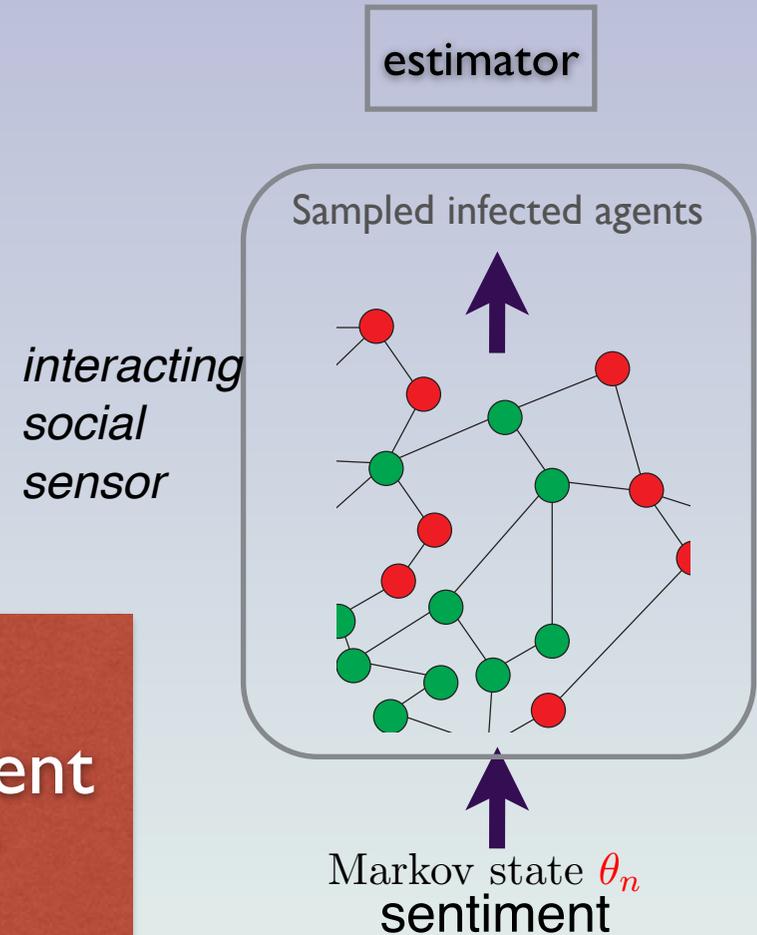


Sampled infected agents

Markov state  $\theta_n$   
sentiment



# PART 2: INFORMATION DIFFUSION IN SOCIAL NETWORKS



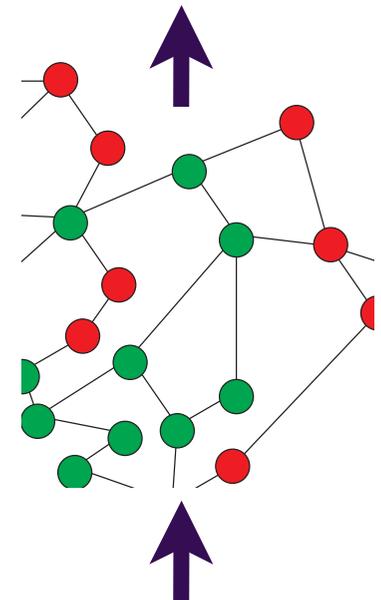
## OUTLINE for Part 2 (very brief)

- Mean Field Dynamics for Sentiment
- How to Sample Social Network?

*How to build a tractable model for information flow in large scale social networks to estimate sentiment?*

estimator

Sampled infected agents



Markov state  $\theta_n$   
sentiment

# MEAN FIELD DYNAMICS FOR SENSING

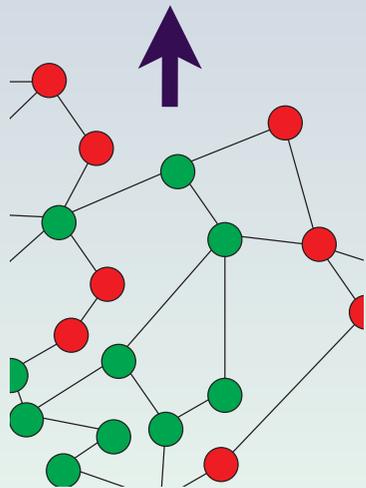
Information diffusion in random graph is asymptotically equivalent to ordinary differential (difference) equation.  
Estimating sentiment is a Bayesian filtering problem

estimator

$$\frac{d\rho(k)}{dt} = F(\rho(k), \rho(k-), \theta_t), \quad \theta_t \sim Q$$

*fraction of infected nodes of degree k*

Sampled infected agents

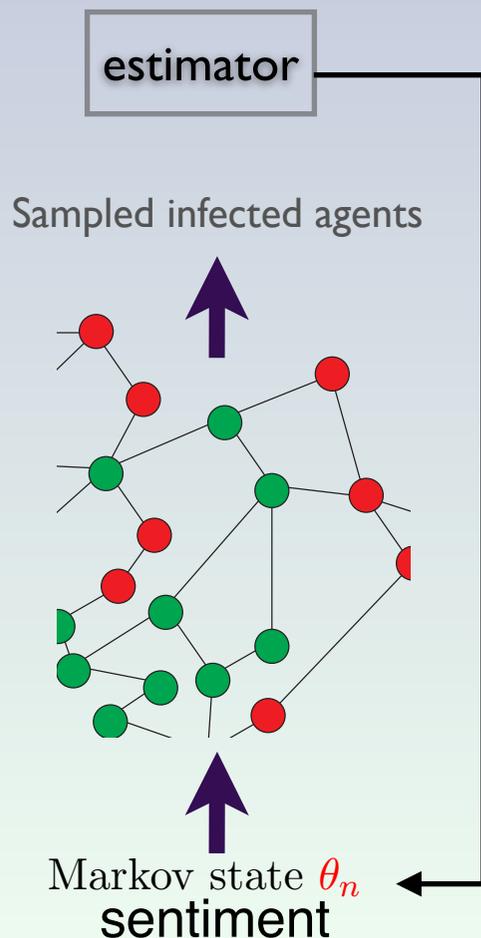


Measurement process:  $Y_t = \int_0^t \lambda(\theta_t) dt + w_t$   
*sampling network*

Markov state  $\theta_n$   
sentiment

# MEAN FIELD DYNAMICS FOR SENSING

Information diffusion in random graph is asymptotically equivalent to ordinary differential (difference) equation.  
Estimating sentiment is a Bayesian filtering problem



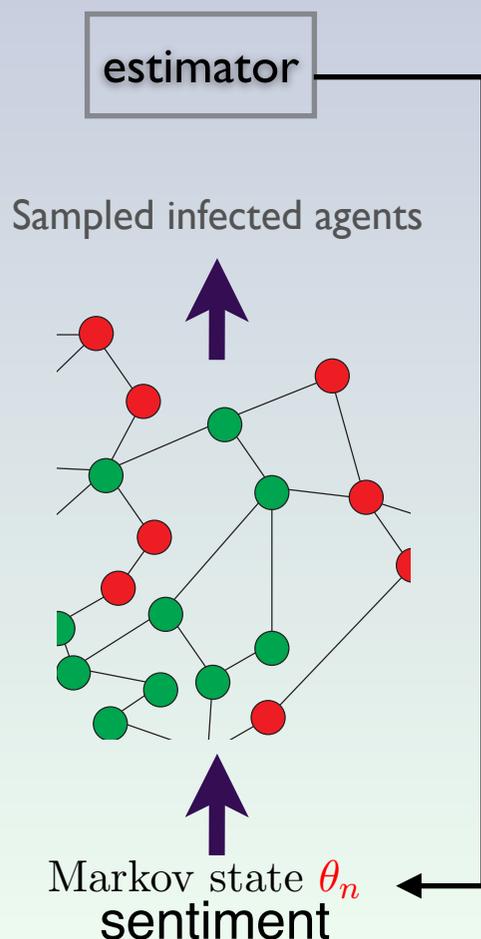
$$\frac{d\rho(k)}{dt} = F(\rho(k), \rho(k-), \theta_t), \quad \theta_t \sim Q$$

*fraction of infected nodes of degree k*

Measurement process:  $Y_t = \int_0^t \lambda(\theta_t) dt + w_t$   
*sampling network*

# MEAN FIELD DYNAMICS FOR SENSING

Information diffusion in random graph is asymptotically equivalent to ordinary differential (difference) equation.  
Estimating sentiment is a Bayesian filtering problem



$$\frac{d\rho(k)}{dt} = F(\rho(k), \rho(k-), \theta_t), \quad \theta_t \sim Q$$

*fraction of infected nodes of degree  $k$*

Measurement process:  $Y_t = \int_0^t \lambda(\theta_t) dt + w_t$   
*sampling network*

## References:

- Benaim, Econometrica, 2003
- Pastor-Satorras, Epidemic spreading in scale free networks, Physical Review Letters, 2001
- D. Lopez-Pintado, Diffusion in complex social networks, Games & Economic Behavior, 2008.
- Sun, Modeling Contagion Through Facebook News Feed, AAAI Conf Social Media, 2009
- Sakaki, Earthquake shakes twitter users: Real time event detection using social sensors, 2010.

# HOW TO SAMPLE SOCIAL NETWORK

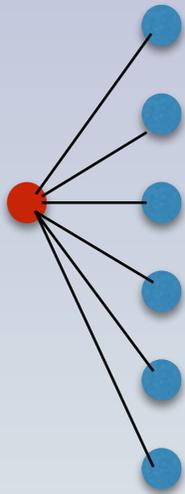


# HOW TO SAMPLE SOCIAL NETWORK

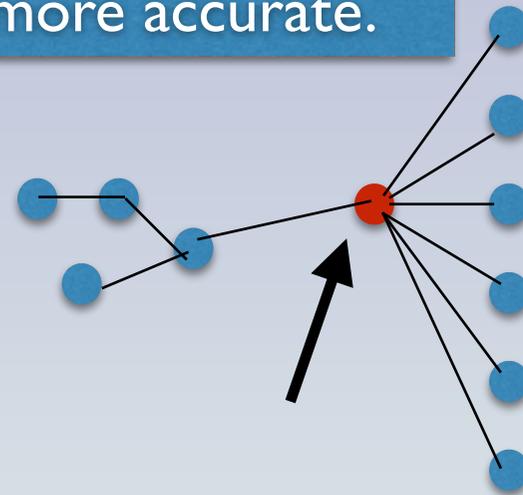
- Intent Polling: who will you vote for?
  - Expectation Polling: who do you think will win?
- Intuitively: expectation polling is more accurate.

# HOW TO SAMPLE SOCIAL NETWORK

- Intent Polling: who will you vote for?
  - Expectation Polling: who do you think will win?
- Intuitively: expectation polling is more accurate.



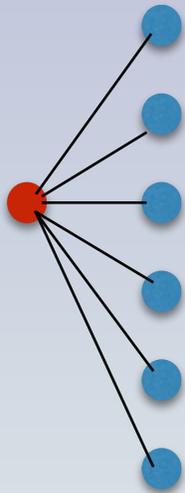
Expectation polling can be biased



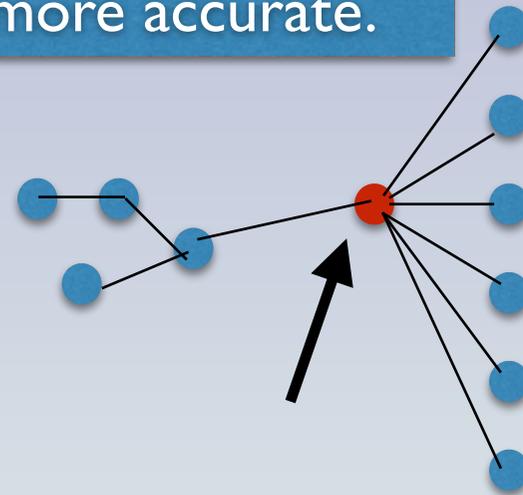
Expectation polling can have higher variance

# HOW TO SAMPLE SOCIAL NETWORK

- Intent Polling: who will you vote for?
  - Expectation Polling: who do you think will win?
- Intuitively: expectation polling is more accurate.



Expectation polling can be biased

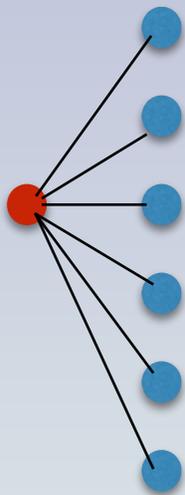


Expectation polling can have higher variance

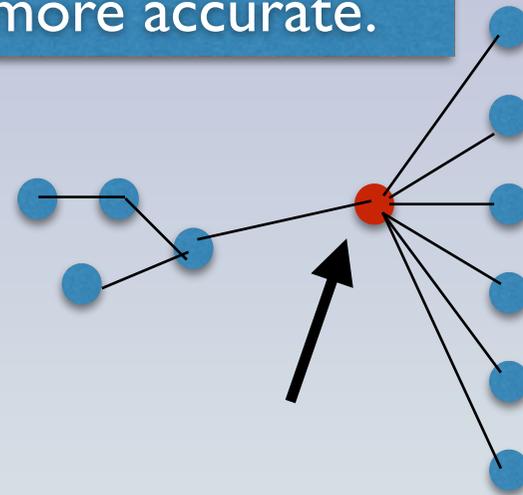
*weight samples inversely proportional to their degree - then unbiased*

# HOW TO SAMPLE SOCIAL NETWORK

- Intent Polling: who will you vote for?
- Expectation Polling: who do you think will win?  
Intuitively: expectation polling is more accurate.



Expectation polling can be biased



Expectation polling can have higher variance

*weight samples inversely proportional to their degree - then unbiased*

- Mean number of friends are smaller than mean number of friend of friends (Feld 1991 - friendship paradox).
- Respondent Driven Sampling: snowball MCMC sampling method for marginalized populations in social networks.  
US Centers for Disease Control and Prevention: HIV drug users.

# SUMMARY

*Part I. Social Learning.* Learn from observations, past decisions and others [Psychology, Economics, CS,EE]  
Herding and Data Incest yield non-informative datasets.

*This presentation is highly simplified and omits several important areas:*

*Dynamic Coherent Risk Measures*

*Homophily vs Contagion*

*Dynamics of viral marketing:*

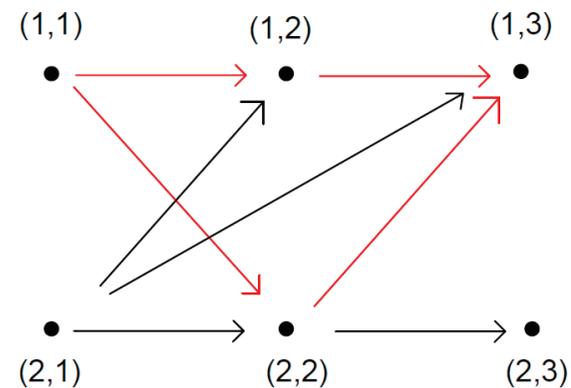
*Revealed Preferences: Are humans utility maximizers?*



# SUMMARY

*Part I. Social Learning.* Learn from observations, past decisions and others [Psychology, Economics, CS,EE]  
Herding and Data Incest yield non-informative datasets.

Abusive re-use of information



*This presentation is highly simplified and omits several important areas:*

*Dynamic Coherent Risk Measures*

*Homophily vs Contagion*

*Dynamics of viral marketing:*

*Revealed Preferences: Are humans utility maximizers?*

# SUMMARY

**Part 1. Social Learning.** Learn from observations, past decisions and others [Psychology, Economics, CS,EE]  
Herding and Data Incest yield non-informative datasets.

**Part 2. Diffusion in Large Scale Networks.**  
Mean field dynamics: infected degree distribution satisfies differential equation  
Sampling: Expectation Polling, Respondent Driven Sampling

*This presentation is highly simplified and omits several important areas:*

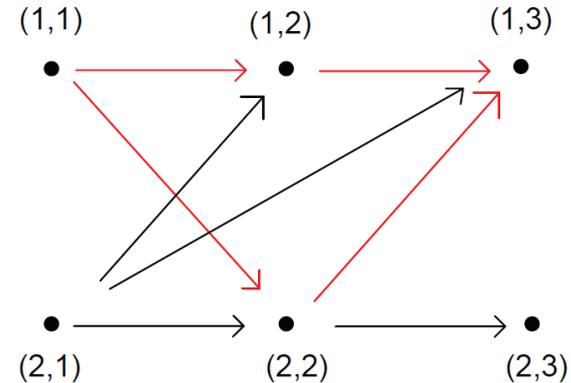
*Dynamic Coherent Risk Measures*

*Homophily vs Contagion*

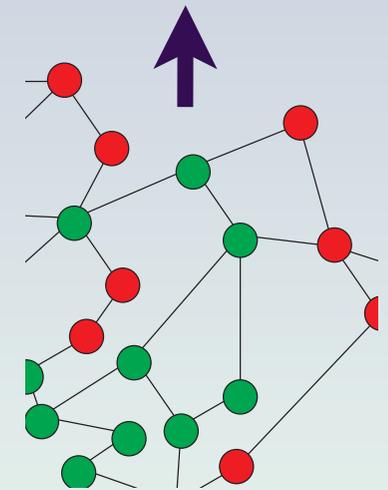
*Dynamics of viral marketing:*

*Revealed Preferences: Are humans utility maximizers?*

Abusive re-use of information



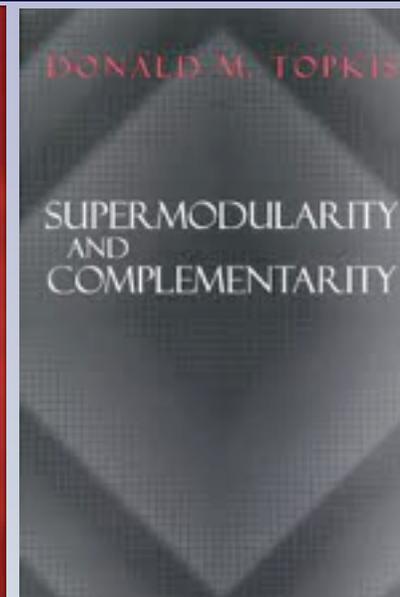
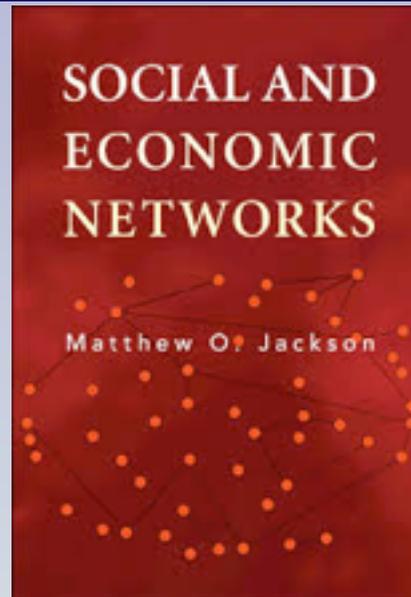
Sampled infected agents



Markov state  $\theta_n$   
sentiment

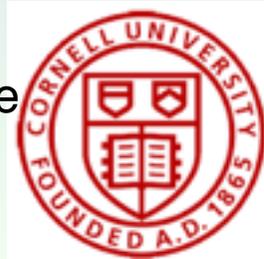


# REFERENCES

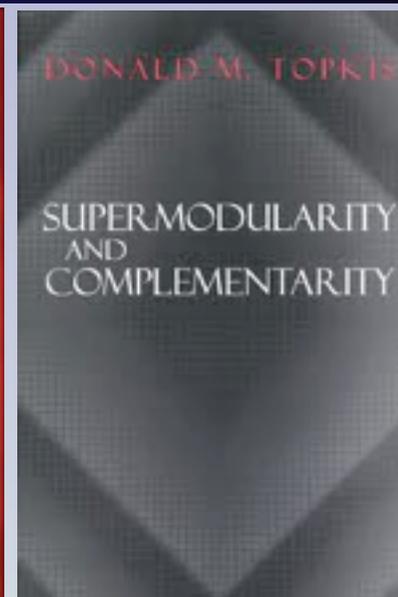
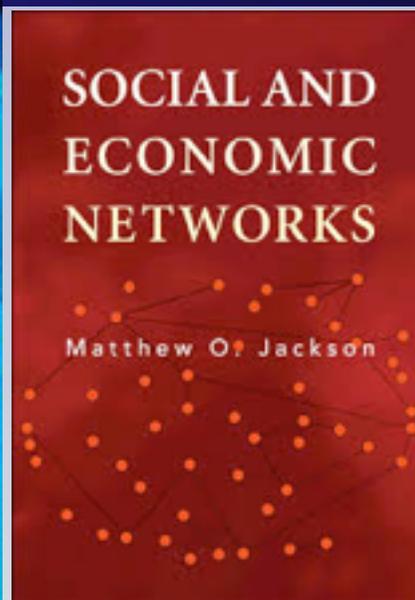
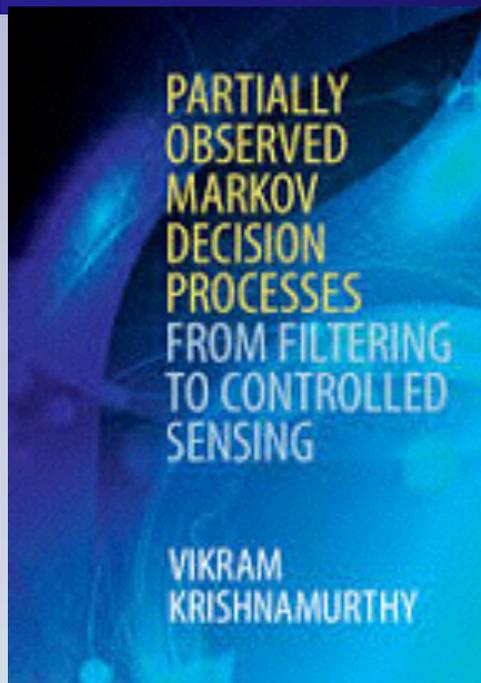
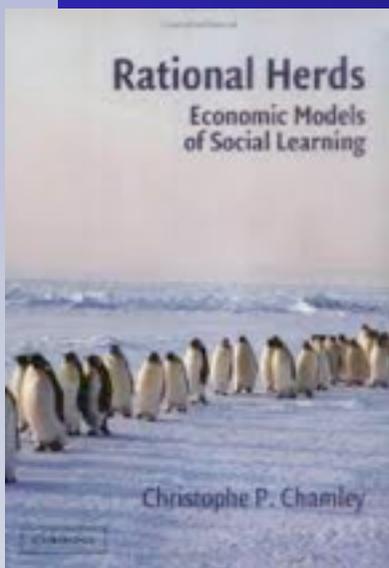


This talk

- Krishnamurthy, Namvar and Hamdi, *Interactive Sensing and Decision Making in Social Networks*, 2014 (monograph in *Foundations & Trends in SP*)
- Krishnamurthy and Hoiles, Social Learning, Data Incest and Revealed Preferences, *IEEE Journal Computational Social Systems*, 2015
- Krishnamurthy, Quickest Detection POMDPs with Social Learning, *IEEE Trans Information Theory*, 2012
- Krishnamurthy & Yin, Tracking the Degree Distribution of Random Graphs, *IEEE Trans Information Theory*, 2014.
- Krishnamurthy, Bhatt, Sequential Detection of Market Shocks with Risk-Averse CVaR social sensors, *IEEE Journal Selected Topics Signal Proc*, 2016



# REFERENCES



This talk

- Krishnamurthy, Namvar and Hamdi, *Interactive Sensing and Decision Making in Social Networks*, 2014 (monograph in *Foundations & Trends in SP*)
- Krishnamurthy and Hoiles, Social Learning, Data Incest and Revealed Preferences, *IEEE Journal Computational Social Systems*, 2015
- Krishnamurthy, Quickest Detection POMDPs with Social Learning, *IEEE Trans Information Theory*, 2012
- Krishnamurthy & Yin, Tracking the Degree Distribution of Random Graphs, *IEEE Trans Information Theory*, 2014.
- Krishnamurthy, Bhatt, Sequential Detection of Market Shocks with Risk-Averse CVaR social sensors, *IEEE Journal Selected Topics Signal Proc*, 2016

