Experimentation on Networks

Simon Board Moritz Meyer-ter-Vehn

UCLA

October 20, 2022

Some Innovations Diffuse Widely

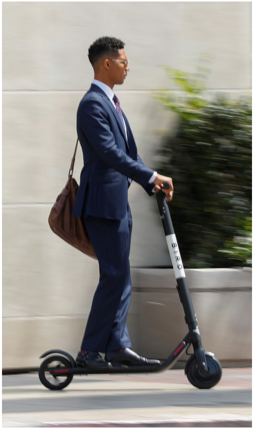

Some Innovations Diffuse Widely

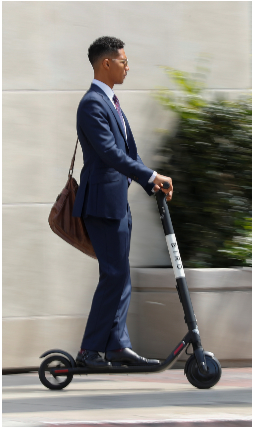

Some Innovations Diffuse Widely

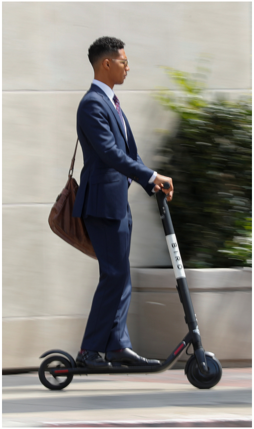

Some Innovations Diffuse Widely


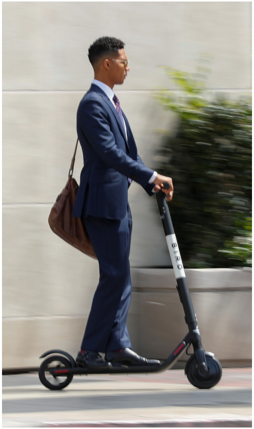

Some Innovations Diffuse Widely

Some Innovations are Abandoned Prematurely

Motivation

Discovery and diffusion are drivers of long-term growth

- Both necessary for sustained progress (Mokyr).
- But social info crowds out acquisition (Grossman-Stiglitz).

Social networks are key conduit of information

- Farmers learn about new technology from neighbors.
- Doctors learn about new drug from classmates.
- Researchers learn about innovation from colleagues.

Research questions

- Quantify the crowding out of private experimentation.
- How does this depend on the network and position within it?
- Does large society learn asymptotically? Maximize welfare?

Contribution

Canonical model of strategic experimentation

- Agents learn from own and neighbors' successes.
- Widely used class of networks (clique, trees, core-periphery).
- Tight characterization of uniquish learning dynamics.

Network density

- Asymptotic information decreases in density.
- Welfare hump-shaped in density.

Network position and link types

- Core-periphery networks: Core agents work least, are best-off.
- Tree networks: Ranking of links and clusters.

Illustration of Equilibrium Dynamics in Star Network

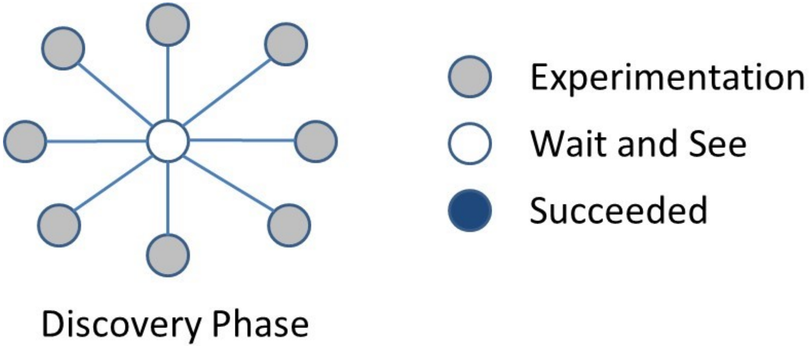

Illustration of Equilibrium Dynamics in Star Network

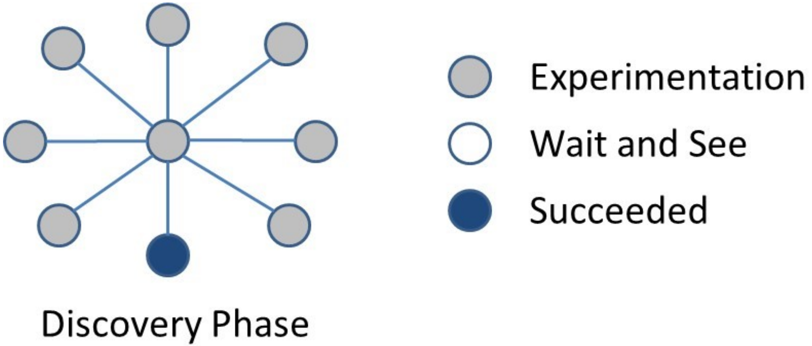

Illustration of Equilibrium Dynamics in Star Network

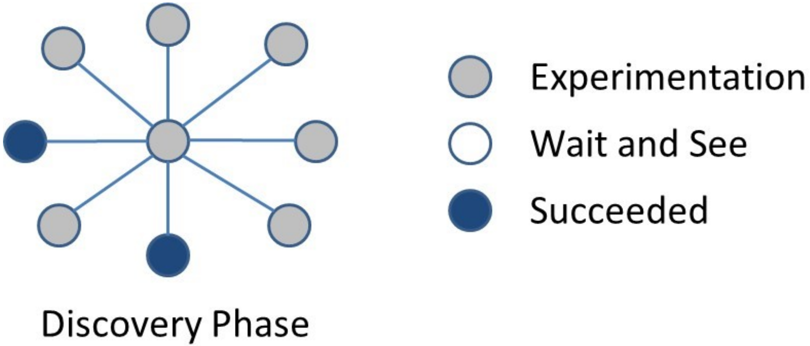

Illustration of Equilibrium Dynamics in Star Network

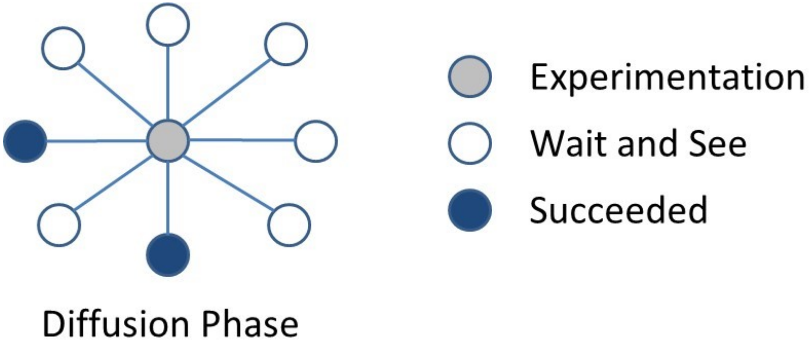

Illustration of Equilibrium Dynamics in Star Network

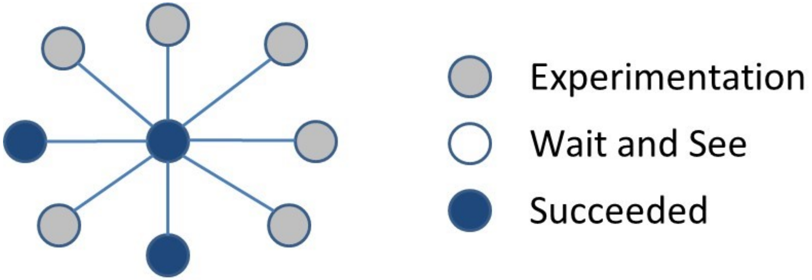

Diffusion Phase

Illustration of Equilibrium Dynamics in Star Network

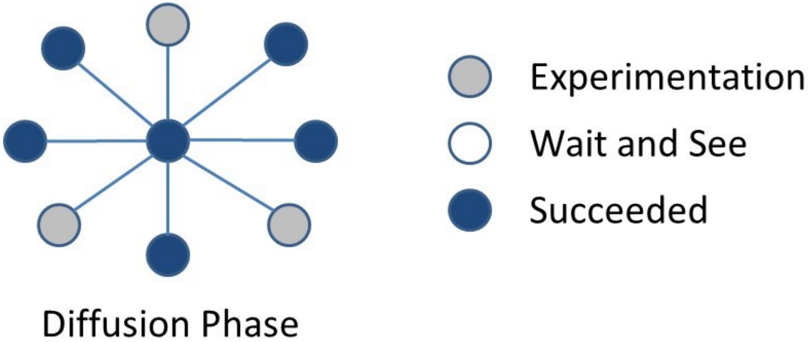

Illustration of Equilibrium Dynamics in Star Network

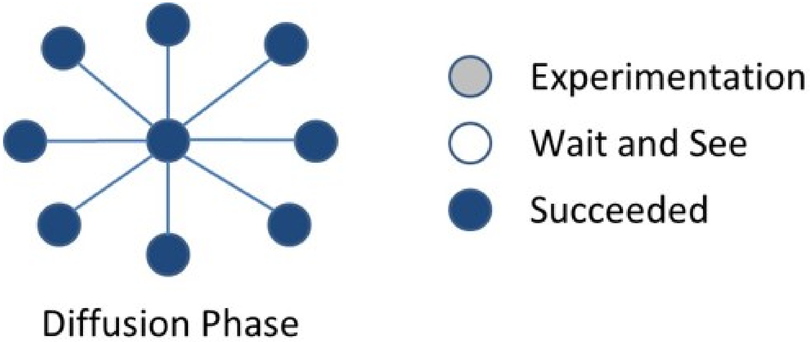

Illustration of Equilibrium Dynamics in Star Network

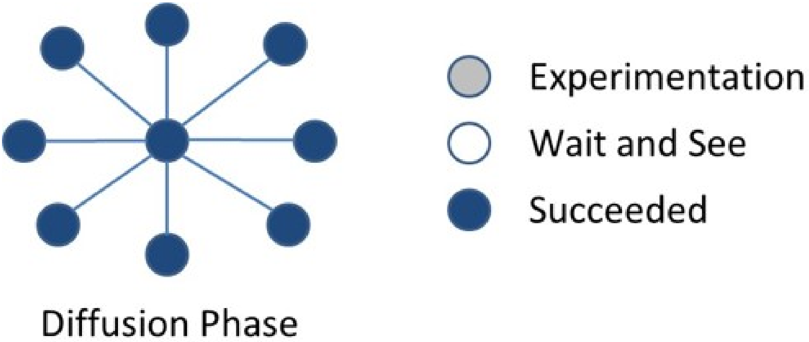

Adding a second core node...

- ... accelerates diffusion

Illustration of Equilibrium Dynamics in Star Network

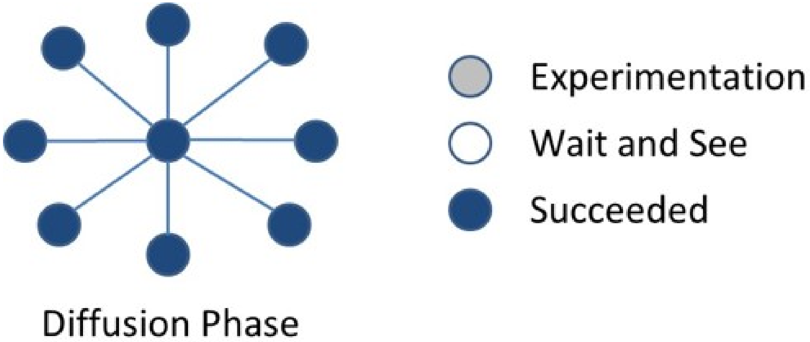

Adding a second core node...

- ... accelerates diffusion
- ... but may choke off discovery

Literature

Learning from repeated interactions on networks

- Bala, Goyal (1998), Sadler (2020) ...
- Mossel, Sly, Tamuz (2015), & Muller-Frank (2020)
- Rosenberg, Solan, and Vieille (2009), Camargo (2014)
- Sahlish (2015)

Strategic experimentation on clique

- Keller, Rady, Cripps (2005) - observed actions
- Bonatti, Horner (2011) - common payoffs
- Bonatti, Horner (2017) - bad news

Empirics

- Hodgson (2021) - UK offshore drilling
- Steck (2018) - North Dakota fracking

The Equilibrium Problem

Our agents are forward-looking

- "if agents were not myopic, their incentives for strategic behavior (such as free riding) would also interact with the imperfect monitoring of the rest of society in very complex ways." (Bala, Goyal 1998)
- Perfect good news reduces i's problem to optimal cutoff τ_i.

Our agents are fully Bayesian

- "when agents attempt to infer information in society as a whole, they must take into account that other agents simultaneously make similar inferences, and make choices based upon these inferences." (Bala, Goyal 1998)
- Focus on simple, yet interesting classes of networks.

Model

Model

Network

- Agents $i = 1, ..., I$ in (possibly) random network G.
- Realized network g; neighbors $N_i(g)$.

Learning Game

- State $\theta \in \{L, H\}$ with prior $\Pr(H) = p_0$.
- Private effort $A_{i,t} \in [0, 1]$ at $t \in [0, \infty)$ at flow cost c.
- Successes at $\{T_i^\iota\}$ with rate $A_{i,t}\mathbb{I}_{\{\theta=H\}}$; payoff $x > c$ to i .
- Observe own and neighbors' successes, and G but not g.

Agent i's problem

$$V_i = \max_{\{A_{i,t}\}_{t\geq 0}} E\left[x\sum_{\iota=1}^{\infty} e^{-rT_i^\iota} - c\int_0^\infty e^{-rt}A_{i,t}dt\right]$$

Best Responses

The Experimentation Problem

- First success times: T for agent i, and S for her neighbors.
- Continuation value after success: $y = (x - c)/r$.
- Before T, S, *experimentation* $a_t^\emptyset$ and *social learning*

$$b_t := E^H \left[\sum_{j \in N(G)} A_{j,t} \middle| T, S > t \right]$$

Belief dynamics

$$p_t = P^\emptyset \left(\int_0^t (a_s^\emptyset + b_s) ds \right) \quad \text{where} \quad P^\emptyset(z) := \frac{p_0 e^{-z}}{p_0 e^{-z} + (1 - p_0)}$$

Experimentation problem: $V = \max_{\{a_t^\emptyset\}} \Pi(\{a_t^\emptyset\}, \{b_t\})$ with

$$\Pi(\{a_t^\emptyset\}, \{b_t\}) := \int_0^\infty \left[p_t(a_t^\emptyset(x + y) + b_t y) - a_t^\emptyset c \right] e^{-\int_0^t (r + (a_s^\emptyset + b_s) p_s) ds} dt$$

Optimality of Cutoff Strategies

Proposition 1.

A cutoff strategy $a_s^{\emptyset} = \mathbb{I}_{\{s \leq \tau\}}$ *is uniquely optimal.*

Optimality of Cutoff Strategies

Proposition 1.

A cutoff strategy $a_s^{\emptyset} = \mathbb{I}_{\{s \leq \tau\}}$ *is uniquely optimal.*

Idea: Why experiment after neighbors' failure?

- Front-loading ϵ effort from t to $t - \delta$ has two effects:

1. Raises net time value of own effort by $r\delta(p_t(x + y) - c)\epsilon > 0$
2. Raises total effort if neighbor succeeds: $p_t b_t \delta \epsilon (x - c) > 0$

In contrast, literature finds internal best-responses

- BH (11): *common* success gives extra incentives to free-ride.
- KRC (05): *observed* actions leads to discouragement effect.

Optimality of Cutoff Strategies

Proposition 1.

Cutoff strategy $a_s^{\emptyset} = \mathbb{I}_{\{s \leq \tau\}}$ *is optimal, where* τ *solves* $\psi_\tau = 0$. *Here, (terminal) experimentation incentives are*

$$\psi_\tau := \frac{\partial \Pi}{\partial \tau} := \frac{d}{d\tau} \Pi(\mathbb{I}_{\{s \leq \tau\}}, \{b_s\})$$

Optimality of Cutoff Strategies

Proposition 1.

Cutoff strategy $a_s^{\emptyset} = \mathbb{I}_{\{s \leq \tau\}}$ *is optimal, where* τ *solves* $\psi_\tau = 0$*. Here, (terminal) experimentation incentives are*

$$\psi_\tau := \frac{\partial \Pi}{\partial \tau} \propto p_\tau \left(x + \left(r \int_\tau^\infty e^{-\int_\tau^s (r+b_u)du} ds \right) y \right) - c$$

Optimality of Cutoff Strategies

Proposition 1.

Cutoff strategy $a_s^{\emptyset} = \mathbb{I}_{\{s \leq \tau\}}$ *is optimal, where* τ *solves* $\psi_\tau = 0$. *Here, (terminal) experimentation incentives are*

$$\psi_\tau := \frac{\partial \Pi}{\partial \tau} \propto p_\tau \left(x + \left(r \int_\tau^\infty e^{-\int_\tau^s (r+b_u)du} ds \right) y \right) - c$$

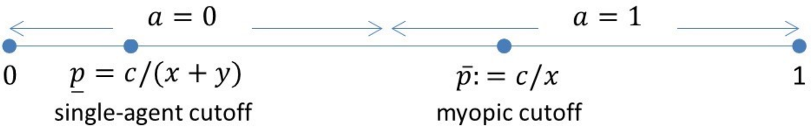

Optimality of Cutoff Strategies

Proposition 1.

Cutoff strategy $a_s^{\emptyset} = \mathbb{I}_{\{s \leq \tau\}}$ *is optimal, where* τ *solves* $\psi_\tau = 0$.
Here, (terminal) experimentation incentives are

$$\psi_\tau := \frac{\partial \Pi}{\partial \tau} \propto p_\tau \left(x + \left(r \int_\tau^\infty e^{-\int_\tau^s (r+b_u)du} ds \right) y \right) - c$$

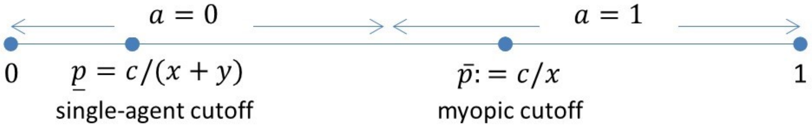

Comparative Statics: A rise in social learning $\{B_t\} := \{\int_0^t b_s\}$

- Increases value V.
- Decreases incentives ψ_τ and experimentation τ.

The Value of Learning Later

Lemma.
Value is fct. of optimal $\tau > 0$ and pre-cutoff learning B_τ

$$V = \mathcal{V}(\tau, B_\tau)$$

$\mathcal{V}$ falls in τ and in B_τ.

"Proof" that (τ, B_τ) is sufficient statistic for $\{B_t\}$

- Pre-cutoff $\{B_s\}_{s \leq t}$ suff. stat. for B_τ: agent works anyway.
- Post-cutoff $\{B_s\}_{s \geq t}$ matters via $V_\tau = P^\emptyset(\tau + B_\tau)(x + y) - c$.

The Value of Learning Later

Lemma.
Value is fct. of optimal $\tau > 0$ and pre-cutoff learning B_τ

$$V = \mathcal{V}(\tau, B_\tau)$$

$\mathcal{V}$ falls in τ and in B_τ.

"Proof" that $\mathcal{V}$ *falls* in B_τ

- Front-loading social learning: $b_{\tau-\epsilon} \uparrow$ and $b_{\tau+\epsilon} \downarrow$ ...
- ... lowers incentives ψ_τ because $|\frac{\partial \psi_\tau}{\partial b_{\tau-\epsilon}}| > |\frac{\partial \psi_\tau}{\partial b_{\tau+\epsilon}}|$.
- ... lowers optimal cutoff τ (first-order).
- ... also raises value V (second-order).
- For τ fixed, $b_{\tau-\epsilon} \uparrow$ requires $b_{\tau+\epsilon} \Downarrow$ so $V \downarrow$

The Value of Learning Later

Lemma.
Value is fct. of optimal $\tau > 0$ and pre-cutoff learning B_τ

$$V = \mathcal{V}(\tau, B_\tau)$$

$\mathcal{V}$ *falls in τ and in B_τ.*

Shirking beats working: $\mathcal{V}$ maximized by $\tau = B_\tau = 0$.

Illustration of Equilibrium Analysis

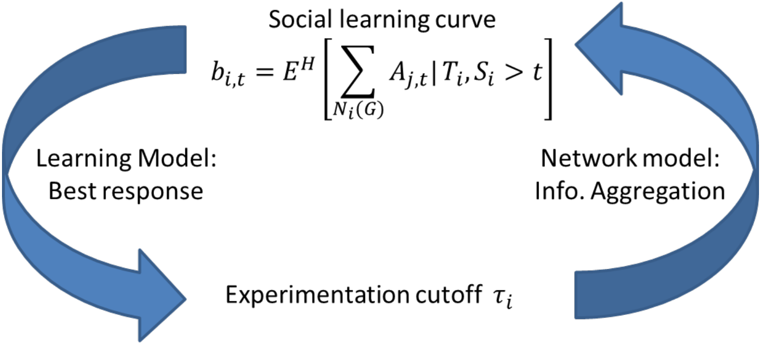

Equilibrium Existence and Uniquishness

Proposition 2.

Equilibrium exists (Brouwer).

Uniqueness for strongly symmetric networks

- Let $g_{i \leftrightarrow j}$ be network g with i, j "switched"
- Let $G_{i \leftrightarrow j}$ be network with $\Pr^{G_{i \leftrightarrow j}}(g) = \Pr^{G}(g_{i \leftrightarrow j})$

Proposition 3.

If $G = G_{i \leftrightarrow j}$, then $\tau_i = \tau_j$ in every equilibrium.

Corollary.

If $G = G_{i \leftrightarrow j}$ $\forall i, j$, equilibrium is unique and symmetric, $\tau_i \equiv \tau$.

Example Networks

A Dense Example: The Clique (Complete Network)

Lemma.

There is a unique equilibrium; cutoffs are $\tau_j \equiv \bar{\tau}/I$.

Step 1: Complete Crowding Out: $\sum_j \tau_j = \bar{\tau}$

- Common belief p_t decreases to $\underline{p} = P^{\emptyset}(\bar{\tau})$, and stays there.
  - $a^{\emptyset}_{j,t} = 0$ for all j once $p_t \leq \underline{p}$.
  - $a^{\emptyset}_{j,t} = 1$ for some j as long as $p_t > \underline{p}$.

Step 2: Uniqueness

- Agent i indifferent about experimentation at $\max\{\tau_j\}$.
- Frontloading incentives: i experiments until $\max\{\tau_j\}$.

A Dense Example: The Clique (Complete Network)

Lemma.
There is a unique equilibrium; cutoffs are $\tau_j \equiv \bar{\tau}/I$.

Step 1: Complete Crowding Out: $\sum_j \tau_j = \bar{\tau}$

- Common belief p_t decreases to $\underline{p} = P^{\emptyset}(\bar{\tau})$, and stays there.
  - $a^{\emptyset}_{j,t} = 0$ for all j once $p_t \leq \underline{p}$.
  - $a^{\emptyset}_{j,t} = 1$ for some j as long as $p_t > \underline{p}$.

Step 2: Uniqueness

- Agent i indifferent about experimentation at $\max\{\tau_j\}$.
- Frontloading incentives: i experiments until $\max\{\tau_j\}$.

Contrast to static public good games (Samuelson 1954)

- Any $\{\tau_j\}$ with $\sum_j \tau_j = \bar{\tau}$ is an equilibrium.

A Sparse Example: The Infinite Directed Line

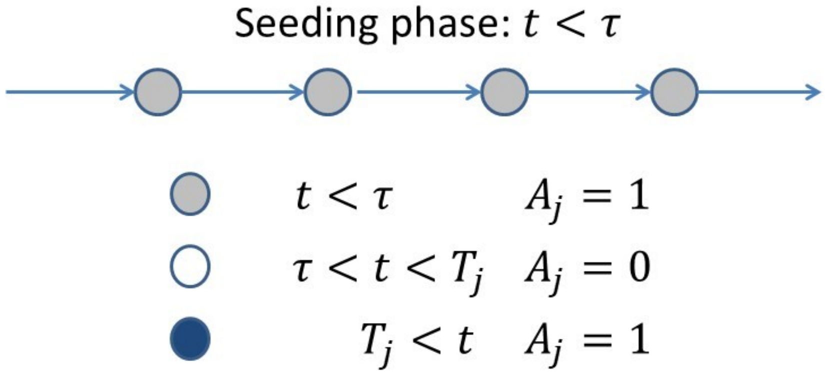

A Sparse Example: The Infinite Directed Line

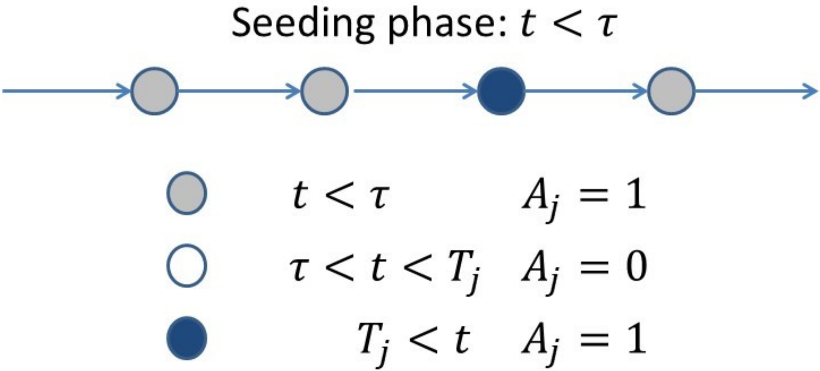

A Sparse Example: The Infinite Directed Line

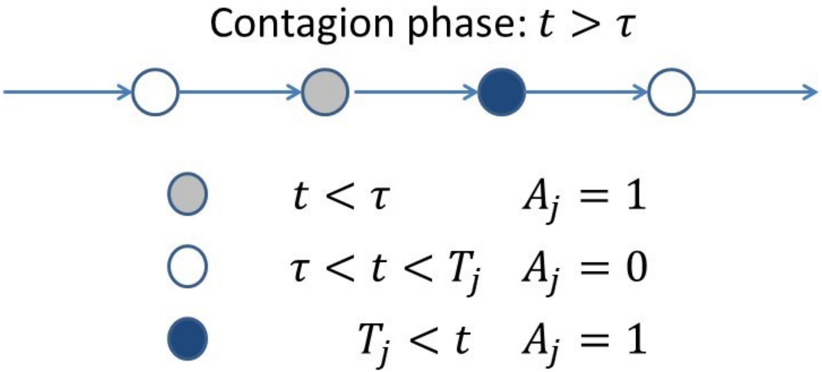

A Sparse Example: The Infinite Directed Line

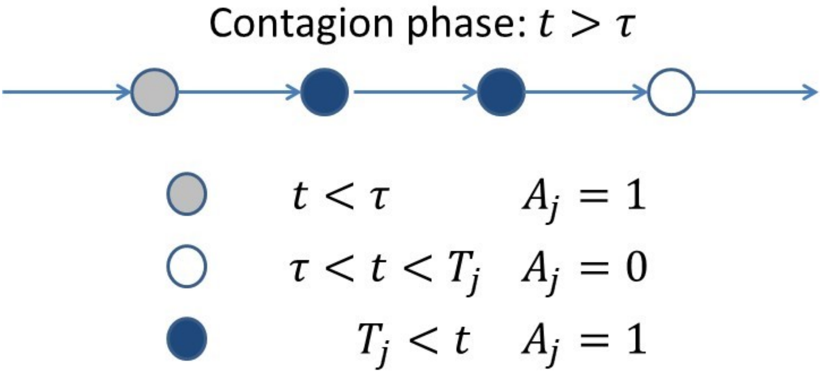

A Sparse Example: The Infinite Directed Line

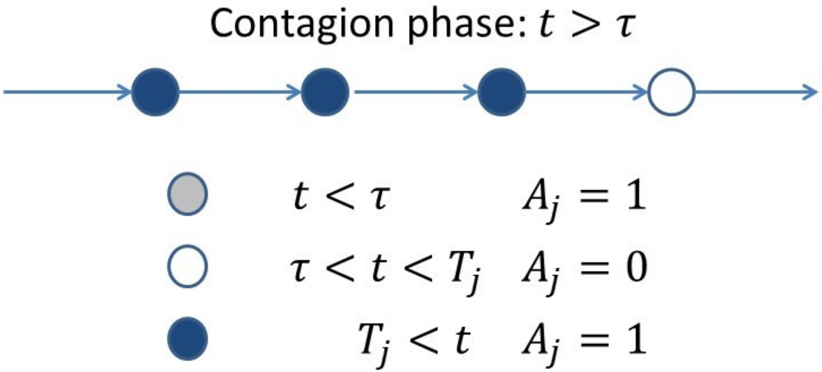

A Sparse Example: The Infinite Directed Line

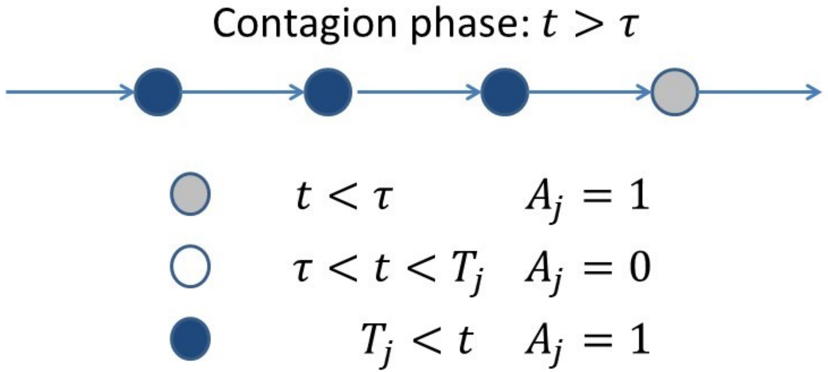

A Sparse Example: The Infinite Directed Line

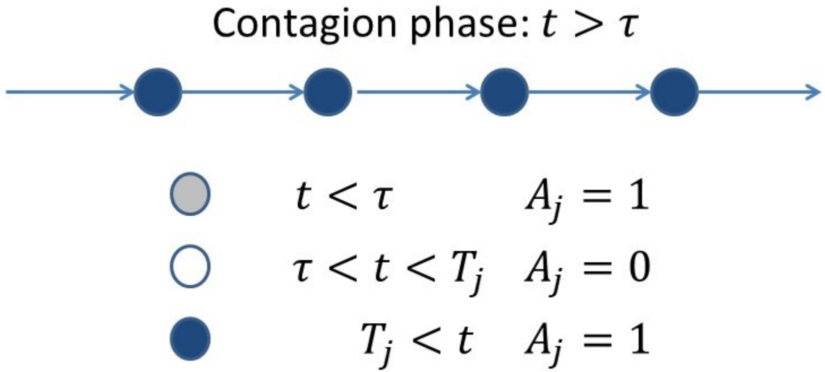

Equilibrium in the Infinite Directed Line

- Assume symmetric cutoffs τ, and focus on $t > \tau$
- Expected effort of i's neighbor j

$$a_t = E^H[A_{j,t}|t < T_i, S_i] = \Pr{}^H(T_k < t|t < T_i, S_i)$$

Equilibrium in the Infinite Directed Line

- Assume symmetric cutoffs τ, and focus on $t > \tau$
- Expected effort of i's neighbor j

$$a_t = E^H[A_{j,t}|t < T_j] = \Pr{}^H(T_k < t|t < T_j)$$

Equilibrium in the Infinite Directed Line

- Assume symmetric cutoffs τ, and focus on $t > \tau$
- Expected effort of i's neighbor j

$$a_t = E^H[A_{j,t} | t < T_j] = 1 - \Pr^H(t < T_k | t < T_j)$$

Bayes' rule

$$\Pr^H(t < T_k | t < T_j) = \frac{\Pr^H(t < T_j | t < T_k) \Pr^H(t < T_k)}{\Pr^H(t < T_j)}$$

Equilibrium in the Infinite Directed Line

- Assume symmetric cutoffs τ, and focus on $t > \tau$
- Expected effort of i's neighbor j

$$a_t = E^H[A_{j,t}|t < T_j] = 1 - \Pr^H(t < T_k|t < T_j)$$

Bayes' rule

$$\Pr^H(t < T_k|t < T_j) = \Pr^H(t < T_j|t < T_k) = e^{-\tau}$$

Equilibrium in the Infinite Directed Line

- Assume symmetric cutoffs τ, and focus on $t > \tau$
- Expected effort of i's neighbor j

$$a_t = E^H[A_{j,t}|t < T_j] = 1 - \Pr{}^H(t < T_k|t < T_j) = 1 - e^{-\tau}$$

is constant in τ

- $\Pr{}^H(T_k < t)$ increasing in t.
- Counteracted by bad news $t < T_j$.

Equilibrium in the Infinite Directed Line

- Assume symmetric cutoffs τ, and focus on $t > \tau$
- Expected effort of i's neighbor j

$$a_t = E^H[A_{j,t}|t < T_j] = 1 - \Pr{}^H(t < T_k|t < T_j) = 1 - e^{-\tau}$$

Equilibrium

$$P^{\emptyset}(2\tau)\left(x + \frac{r}{r + (1 - e^{-\tau})}y\right) = c$$

Information, Welfare and their Benchmarks

Asymptotic Information

$$B := \inf_i B_{i,\infty}$$

- Asymptotic learning, $B = \infty$, achieved by infinite line.

Information, Welfare and their Benchmarks

Asymptotic Information

$$B := \inf_i B_{i,\infty}$$

- Asymptotic learning, $B = \infty$, achieved by infinite line.

Second Best (Rawlsian) Welfare

$$\inf_i V_i \leq p_0 y$$

- Value of instantly learning θ.
- Approximated by I-cliques with (non-eq.) cutoffs $\tau_I = 1/\sqrt{I}$.

Information, Welfare and their Benchmarks

Asymptotic Information

$$B := \inf_i B_{i,\infty}$$

- Asymptotic learning, $B = \infty$, achieved by infinite line.

Second Best Welfare

$$\inf_i V_i \leq \min\{p_0 y, p_0(x+y) - c\} =: V^*$$

- When $p_0 \geq \bar{p}$, then $V^* = p_0 y$.
- When $p_0 \leq \bar{p}$, then $V^* = \mathcal{V}(0,0) = p_0(x+y) - c$.
- Will see networks that approach V^*.

Welfare as a function of the prior p_0

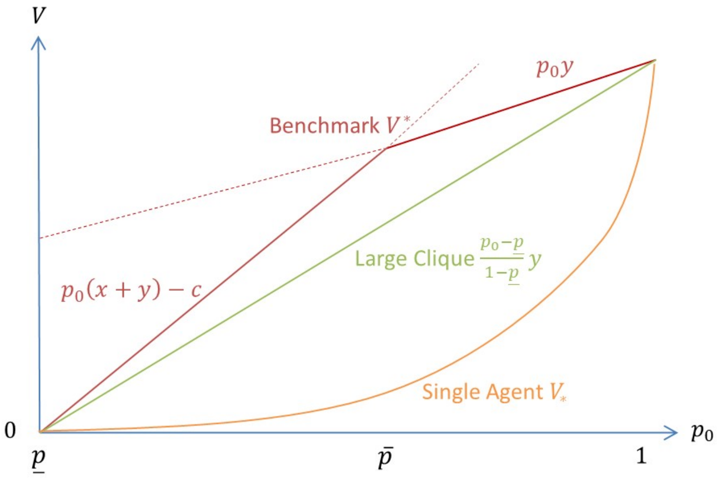

Information and Welfare in Clique and Line

Clique — Too little information

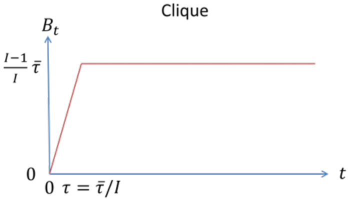

Infinite directed line — Too late information

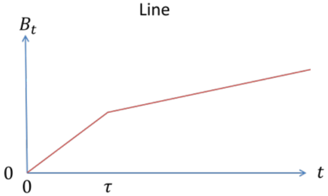

Information and Welfare in Clique and Line

Clique — Too little information

- One-for-one crowding out of private experimentation.
- No asymptotic learning: $(I-1)(\bar{\tau}/I) \to \bar{\tau} < \infty$.
- Welfare below second-best, $V_I \to \mathcal{V}(0, \bar{\tau}) =: V^C < V^*$.

Infinite directed line — Too late information

- Distance between agents mitigates free-riding.
- Asymptotic learning: $\sum_{j \geq i} \vec{\tau}^{(1)} = \infty$.
- Welfare below second-best: $\vec{\tau}^{(1)} > 0$ so $\mathcal{V}(\vec{\tau}^{(1)}, \vec{\tau}^{(1)}) < V^*$.

Large (Regular) Random Networks

Illustration: Regular Random Network $I = 200, n = 6$

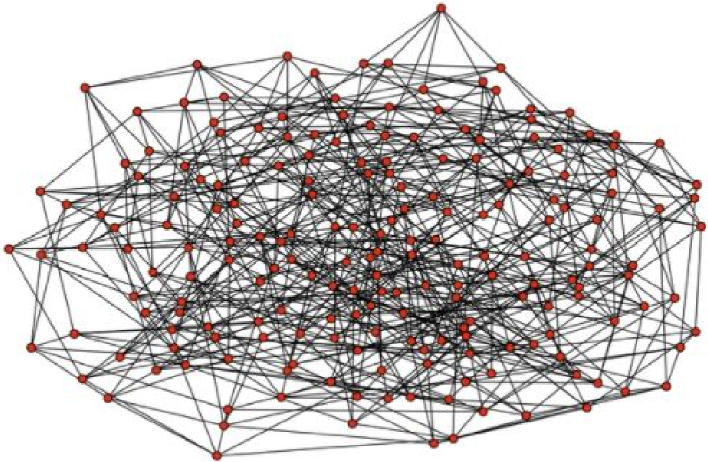

Large Regular Random Networks

The Configuration Model

- I agents with n^I link stubs.
- Randomly connect pairs of stubs into undirected links.
- Unique symmetric equilibrium; cutoff τ^I and value V^I.
- Consider sequences of such networks $\{n^I\}$ as $I \to \infty$.

Network density

- Tree: $n^I \equiv n$; Clique: $n^I/I \to \infty$.
- Network density measures

$$\nu := \lim n^I \qquad \lambda := \lim(n^I/\log I) \qquad \rho := 1 - e^{-\lim(n^I/I)}.$$

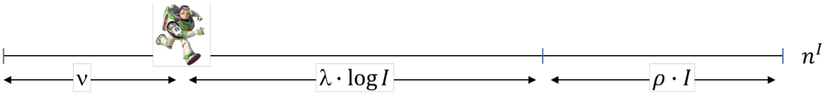

Main Result (Random Networks)

Theorem 1.

Information B falls in density; attains ∞ iff $\lambda \leq 1/\sigma^*$ and $\rho = 0$.
Welfare V is hump-shaped; attains V^* iff $\nu = \infty$ and $\rho = 0$.

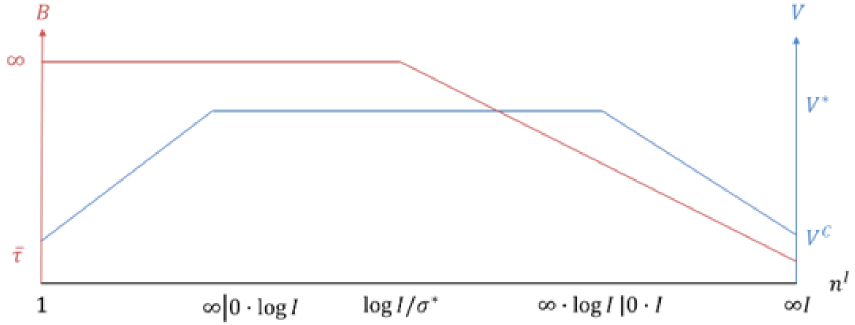

Main Result (Random Networks)

Theorem 1.
Information B falls in density; attains ∞ iff $\lambda \leq 1/\sigma^*$ and $\rho = 0$.
Welfare V is hump-shaped; attains V^* iff $\nu = \infty$ and $\rho = 0$.

Indifference when learning state at time-σ^*

$$p_0\left(x + (1 - e^{-r\sigma^*})y\right) = c \qquad (\text{or } \sigma^* = 0 \text{ if } p_0 \geq \bar{p})$$

Tension between Learning and Welfare: Network must be...

- ... sparse to sustain information generation.
- ... dense for fast diffusion and welfare benchmark.

Social Learning Curves

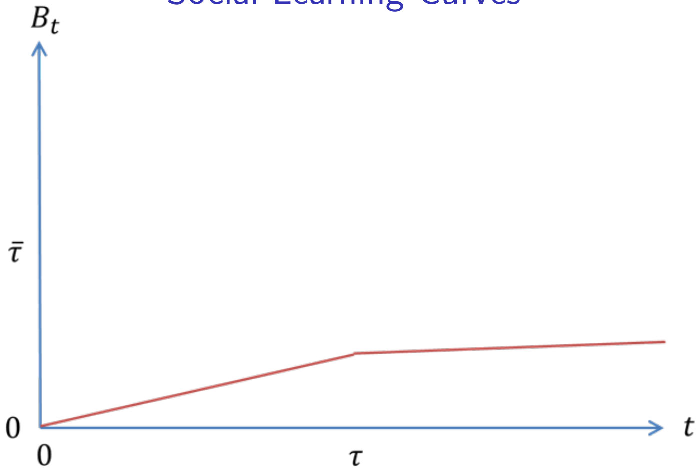

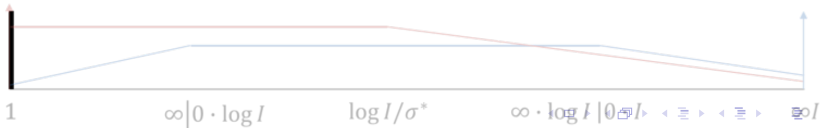

Social Learning Curves

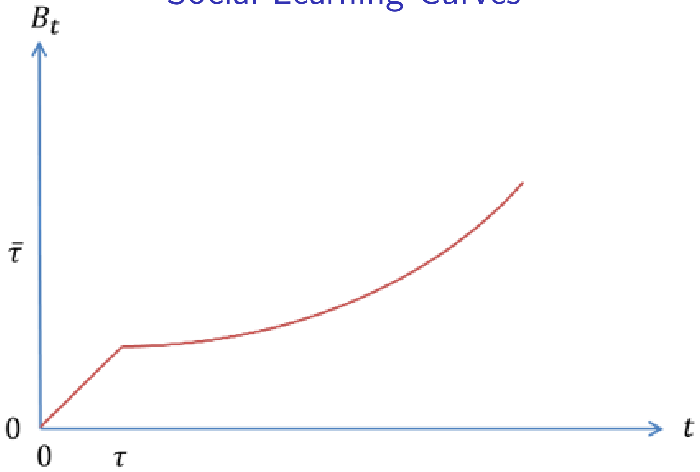

Social Learning Curves

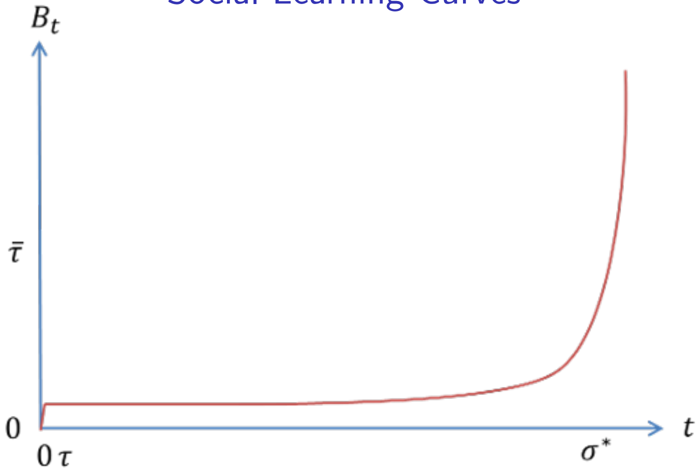

Social Learning Curves

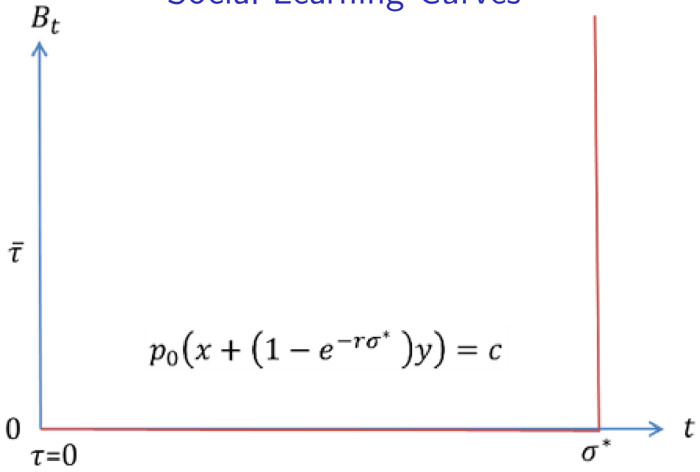

Social Learning Curves

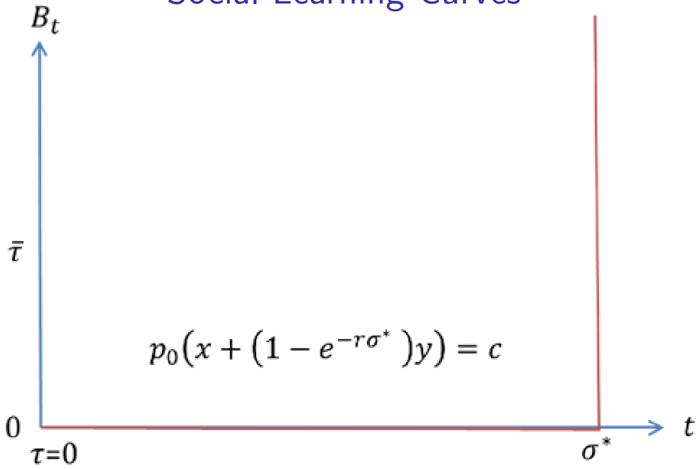

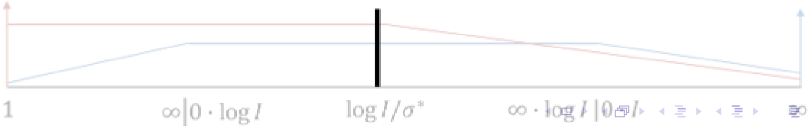

Social Learning Curves

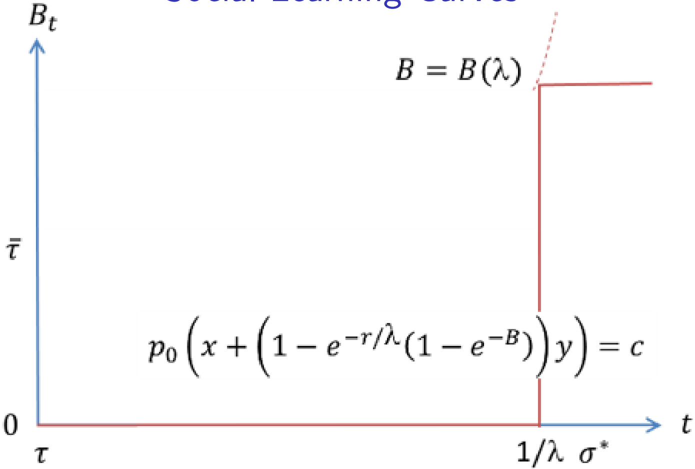

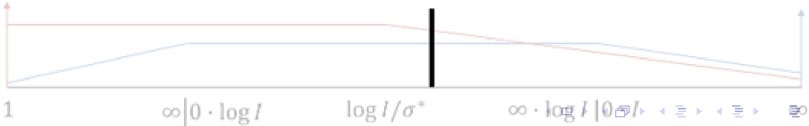

Social Learning Curves

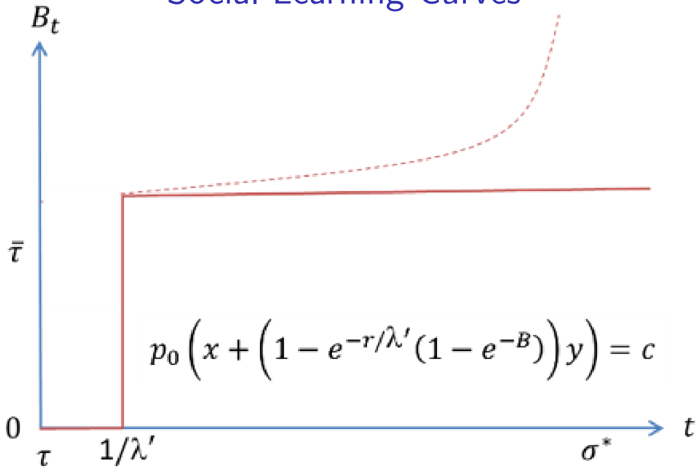

Social Learning Curves

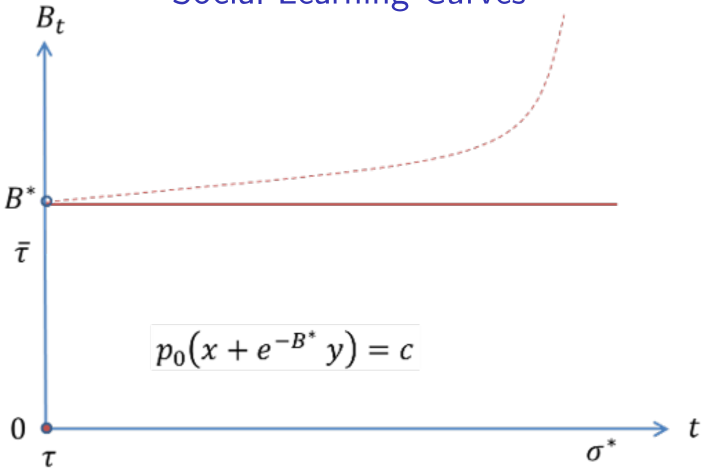

Social Learning Curves

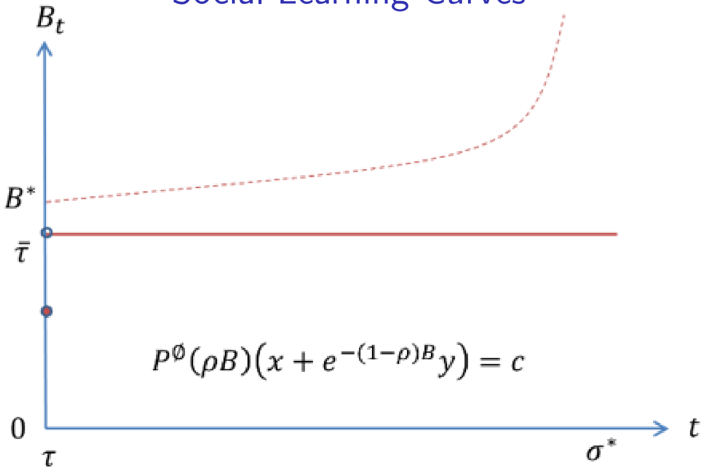

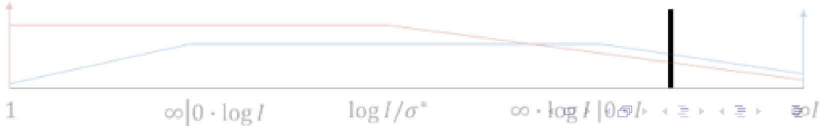

Social Learning Curves

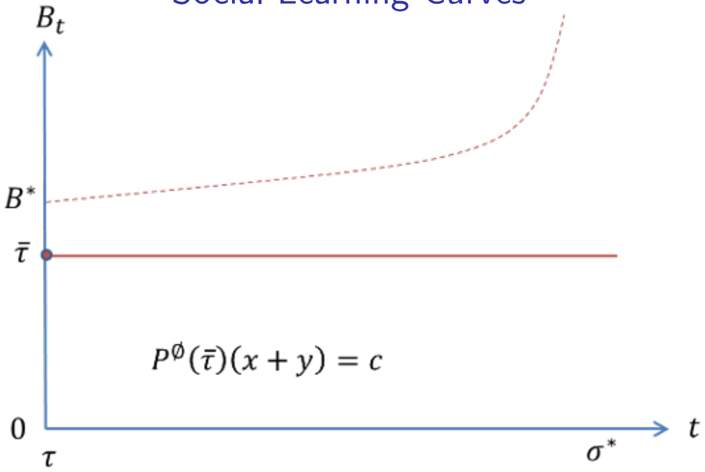

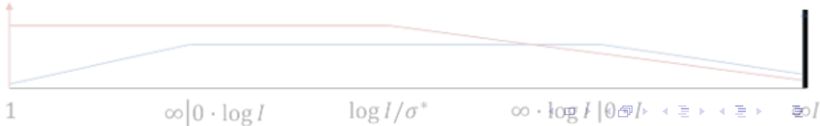

Proof Idea

Tsunami Lemma.
Assume $n^I \to \infty$. *Social learning happens at* $\sigma = \lim \frac{-\log \tau^I}{n^I}$:

$$\lim B_t^I = \begin{cases} 0 & \textit{for all } t < \sigma \\ B & \textit{for all } t > \sigma \end{cases}$$

If $B < \infty$, *then* $\sigma = \lim \frac{\log I}{n^I} = 1/\lambda$.

Proof idea

- Exposed agents

$$E_t^I := \mathbb{E}[\#\{i : S_i^I \leq t\}] \approx I\tau^I e^{n^I t}$$

until $E_\sigma^I \approx \epsilon I$, and then everyone is exposed at once.

Core-Periphery Networks

Core-Periphery Networks

- K fully connected core agents k, and $L = I - K$ peripherals ℓ.
- Examples: $K = 1$ star network; $K = 2$

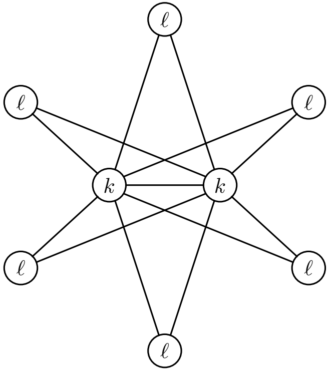

- Common in finance models, with dealer banks in core.
- Arise endogenously in network formation model.

The Effect of Network Position on Experimentation

Lemma.

- *Equilibrium is symmetric, characterized by two cutoffs* τ_k, τ_ℓ.
- *Core agents experiment less,* $\tau_k < \tau_\ell$*, are better off,* $V_k > V_\ell$.

Proof idea

- Social information crowds out experimentation incentives.
- Core agents k have more social information.

Large Core-Periphery Networks

Networks with K^I core agents as $I \to \infty$

- Asymptotic information $B := \lim B^I_{\ell,\infty}$
- Welfare $V := \lim V^I_\ell$

Network density

- Star: $K^I \equiv 1$; Clique: $K^I = I$.
- Network density measures

$$\kappa := \lim K^I \qquad\qquad \rho := \lim K^I / I.$$

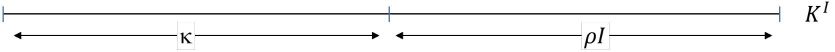

Main Result (Core-Periphery

Theorem 2.

Information B falls in density; attains ∞ iff $\kappa \leq \kappa^*$ and $\rho = 0$.
Welfare V is hump-shaped; attains V^* iff $\kappa \geq \kappa^*$ and $\rho = 0$.

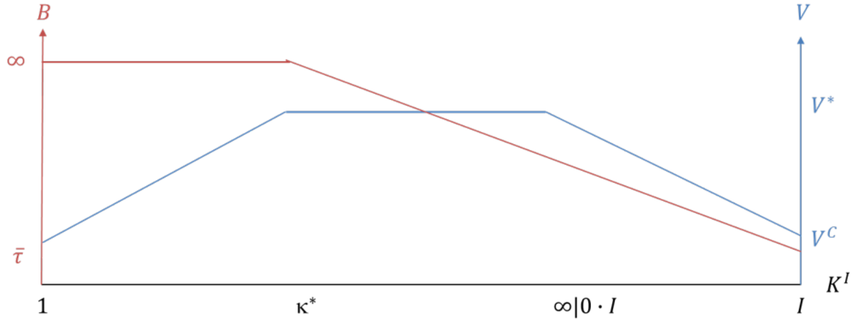

Main Result (Core-Periphery

Theorem 2.
Information B falls in density; attains ∞ iff $\kappa \leq \kappa^*$ and $\rho = 0$.
Welfare V is hump-shaped; attains V^* iff $\kappa \geq \kappa^*$ and $\rho = 0$.

Indifference with κ^* core agents fully experimenting

$$p_0 \left(x + \frac{r}{r + \kappa^*} y \right) = c \qquad \text{(or } \kappa^* = \infty \text{ if } p_0 \geq \bar{p}\text{)}$$

Learning and Welfare Benchmarks Mutually Exclusive

- Asymptotic learning: Sustain peripherals' experimentation.
- Welfare benchmark: Peripherals' experimentation must vanish.

Social Learning Curves

$$P^{\emptyset}((\kappa+1)\,\tau)\left(x+\frac{r}{r+\kappa}y\right)=c$$

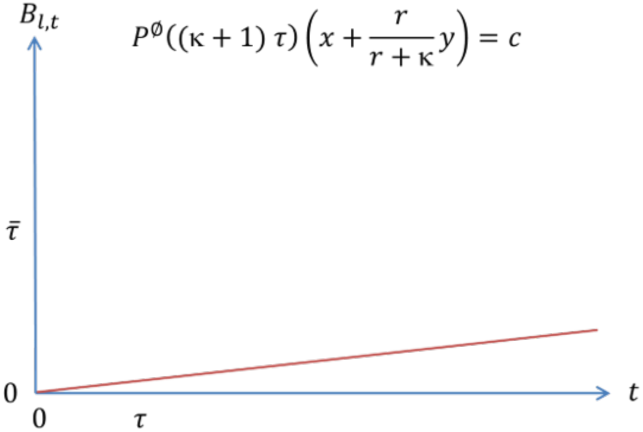

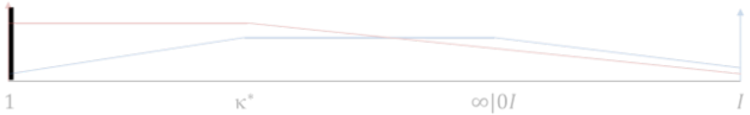

Social Learning Curves

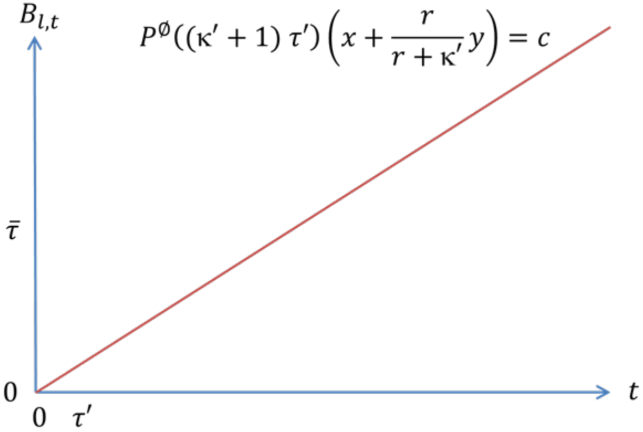

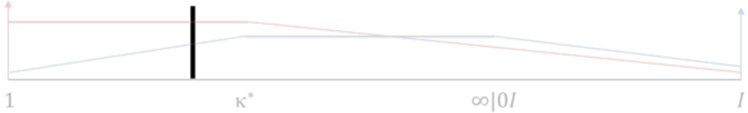

Social Learning Curves

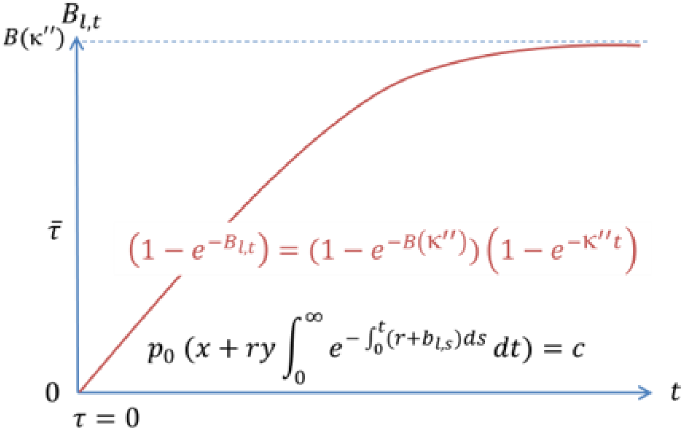

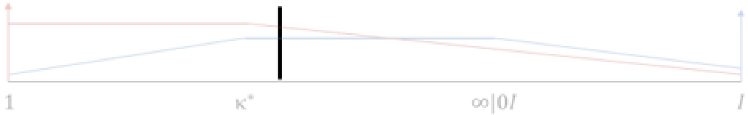

Social Learning Curves

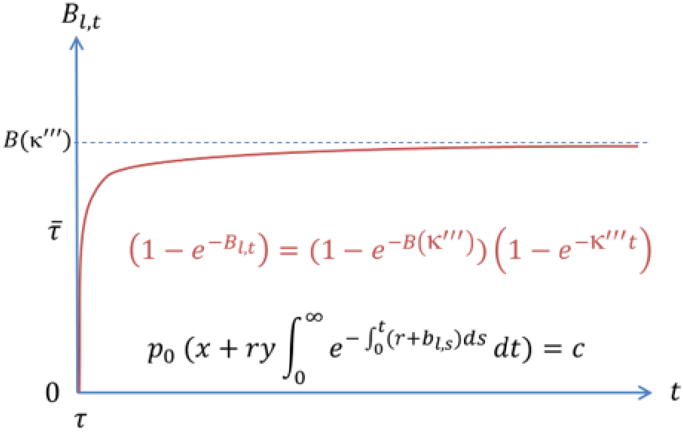

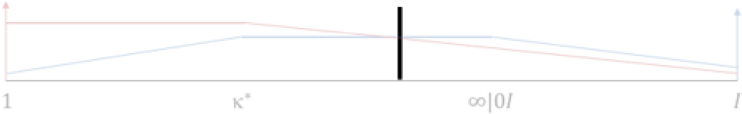

Social Learning Curves

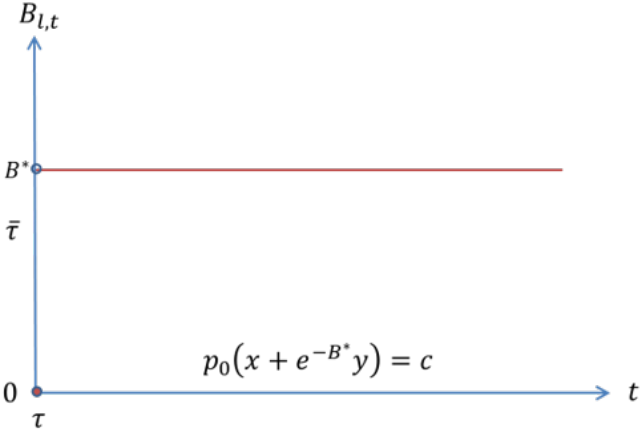

Social Learning Curves

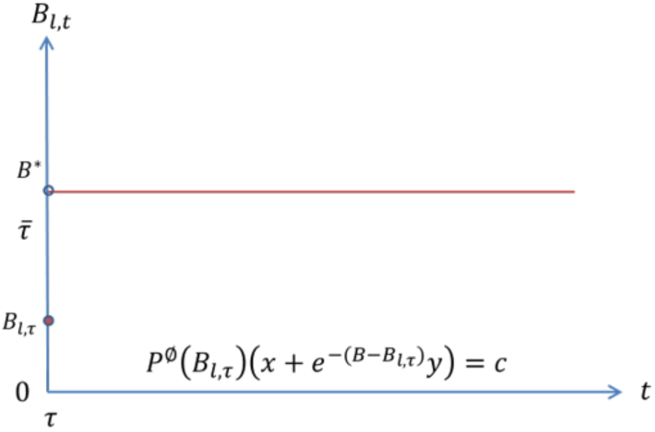

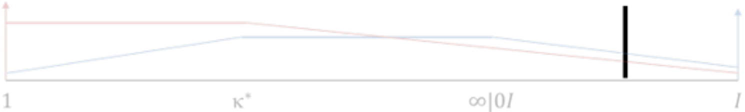

Social Learning Curves

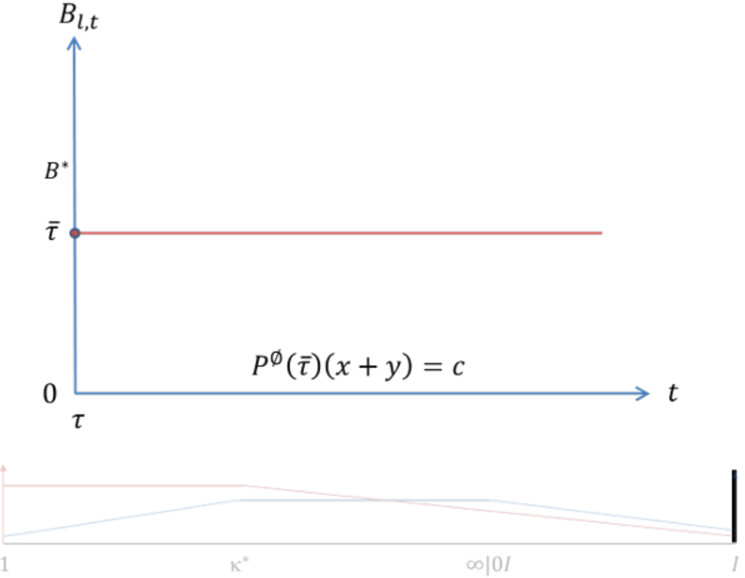

Alternative Core-Periphery Networks

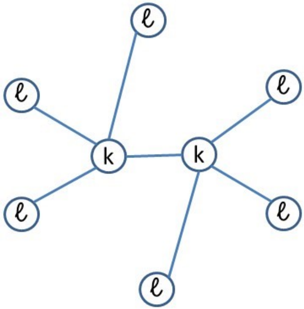

Large periphery $L(k) \to \infty$

- Equilibrium similar to star network: $\tau_\ell > 0 = \tau_k$.

Types of Links

(Infinite Regular) Trees $\mathcal{T}^{(n)}$

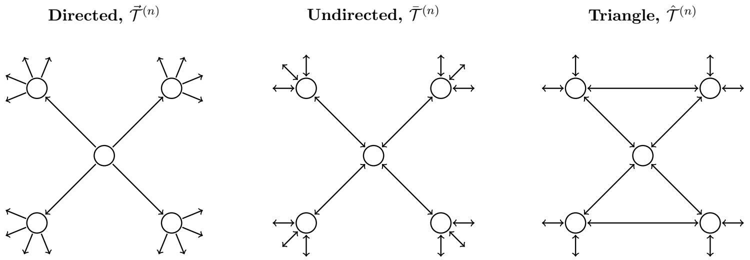

Motivation

- Directed (Twitter), undirected (LinkedIn), cluster (Facebook)
- Approximate large random network (Sadler '20, BMtV '21)
- Highly tractable because of independence across neighbors

How does social learning depend on type of tree?

(Infinite Regular) Trees $\mathcal{T}^{(n)}$

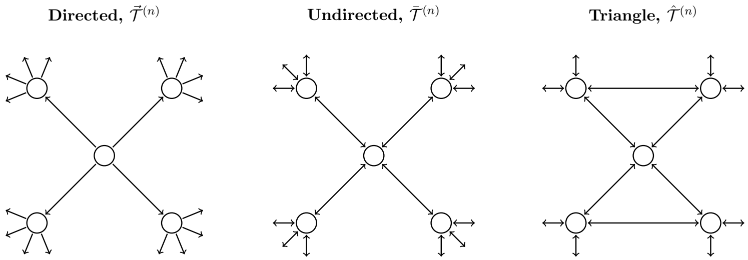

Theorem 3.

$$\vec{V}^{(n)} > \bar{V}^{(n)} > \hat{V}^{(n)}$$
$$\vec{V}^{(n)} < \bar{V}^{(n+1)} < \hat{V}^{(n+2)}$$

Recall: The Directed Line $\vec{\mathcal{T}}^{(1)}$

- Assume symmetric cutoffs τ, and focus on $t > \tau$
- Expected effort of i's neighbor j (conditional on $\theta = H$)

$$a_t = E[A_{j,t}|t < T_j] = 1 - \Pr(t < T_k|t < T_j) = 1 - e^{-\tau}$$

Equilibrium

$$P^{\emptyset}(2\tau)\left(x + \frac{r}{r + (1 - e^{-\tau})}y\right) = c$$

Contrast: The Undirected Line $\bar{\mathcal{T}}^{(2)}$

- Assume symmetric cutoffs τ, and focus on $t > \tau$
- In expectation $E^{-i}[\cdot]$ over $\{T_j\}_{j \neq i}$

$$a_t = E^{-i}[A_{j,t}|t < T_j] = 1 - \Pr{}^{-i}(t < T_k|t < T_j)$$

Bayes' rule

$$\Pr{}^{-i}(t < T_k|t < T_j) = \frac{\Pr^{-i}(t < T_j, T_k)}{\Pr^{-i}(t < T_j)} = \frac{\exp(-\tau - \int_0^t a_s ds)}{\exp(-\int_0^t a_s ds)} = e^{-\tau}$$

Equilibrium

$$P^{\emptyset}(3\tau)\left(x + \frac{r}{r + 2(1 - e^{-\tau})}y\right) = c$$

Summary: Ranking Values across Regular Trees

Theorem 3.

$$\begin{array}{ccccc} \vec{V}^{(n)} & > & \bar{V}^{(n)} & > & \hat{V}^{(n)} \\ \vec{V}^{(n)} & < & \bar{V}^{(n+1)} & < & \hat{V}^{(n+2)} \end{array}$$

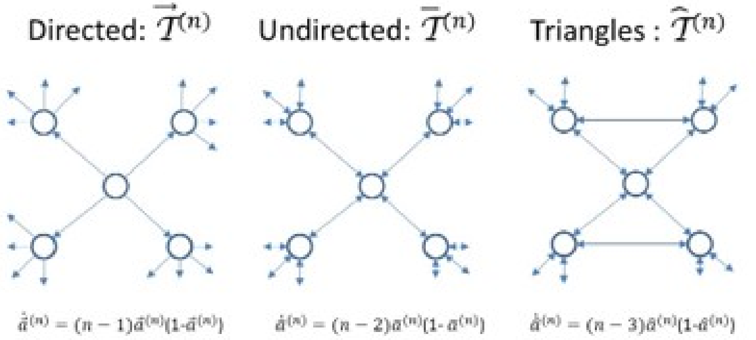

Networks, Schmetworks...

Model

- Agent $i \in [0,1]$ samples n others at every time t, iid.
- Observe only current successes of links.

Difference to fixed networks: Sampling avoids redundancy.

- Neighbor's expected effort with random sampling

$$\tilde{a}_t = E[A_{j,t}|t < T_i, S_i] = 1 - \Pr[t < T_j, S_j]$$

evolves according to $\dot{\tilde{a}} = n\tilde{a}(1 - \tilde{a})$ at $t > \tau$.

- Learn more from sampled neighbor: $\tilde{a}_t > \vec{a}_t$.

Theorem 3.

$$\vec{V}^{(n+1)} > \tilde{V}^{(n)} > \vec{V}^{(n)}$$

Large Random Networks Approximate Trees

- (Undirected) tree $\bar{\mathcal{T}} = \bar{\mathcal{T}}^{(n)}$ with "equilibrium" $\bar{\tau}$, $\{\bar{a}_t\}$.
- Random networks $n^I \equiv n$ with equilibria $\bar{\tau}^I$, $\{\bar{b}_t^I\}$.

Proposition 4.

$$\bar{\tau}^I \to \bar{\tau} \qquad \bar{b}_t^I \to n\bar{a}_t \quad \forall t \geq 0$$

Proof Idea: As $I \to \infty$

- Random networks locally approximate trees.
- ODEs for $E^H[A_{j,t}^I | t < T_i, S_i]$ converge.

Analogue constructions for $\vec{\mathcal{T}}$ and $\hat{\mathcal{T}}$.

Conclusion

An equilibrium model of experimentation on networks

- Forward-looking, optimizing, fully Bayesian agents.
- Perfect good news learning leads to cutoff strategies.

Findings

- Asymptotic information decreases in density.
- Welfare hump-shaped in density.
- Effects of network position and link type.

Next steps

- More realistic networks
- Network formation
- Planner's problem